Deniable Attribute Based Encryption for Branching Programs from LWE

Daniel Apon[1(✉)], Xiong Fan[2], and Feng-Hao Liu[3]

[1] University of Maryland, College Park, USA
dapon@cs.umd.edu
[2] Cornell University, Ithaca, USA
xfan@cs.cornell.edu
[3] Florida Atlantic University, Boca Raton, USA
fenghao.liu@fau.edu

Abstract. Deniable encryption (Canetti et al. CRYPTO '97) is an intriguing primitive that provides a security guarantee against not only eavesdropping attacks as required by semantic security, but also stronger coercion attacks performed after the fact. The concept of deniability has later demonstrated useful and powerful in many other contexts, such as leakage resilience, adaptive security of protocols, and security against selective opening attacks. Despite its conceptual usefulness, our understanding of how to construct deniable primitives under standard assumptions is restricted.

In particular from standard lattice assumptions, i.e. Learning with Errors (LWE), we have only flexibly and non-negligible advantage deniable public-key encryption schemes, whereas with the much stronger assumption of indistinguishable obfuscation, we can obtain at least fully sender-deniable PKE and computation. How to achieve deniability for other more advanced encryption schemes under standard assumptions remains an interesting open question.

In this work, we construct a flexibly bi-deniable Attribute-Based Encryption (ABE) scheme for all polynomial-size Branching Programs from LWE. Our techniques involve new ways of manipulating Gaussian noise that may be of independent interest, and lead to a significantly sharper analysis of noise growth in Dual Regev type encryption schemes. We hope these ideas give insight into achieving deniability and related properties for further, advanced cryptographic systems from lattice assumptions.

1 Introduction

Deniable encryption, introduced by Canetti et al. [14] at CRYPTO 1997, is an intriguing primitive that allows Alice to privately communicate with Bob in a way that resists not only eavesdropping attacks as required by semantic security, but also stronger *coercion attacks* performed after the fact. An eavesdropper Eve stages a cocercion attack by additionally approaching Alice (or Bob, or both) *after* a ciphertext is transmitted and demanding to see all secret information:

the plaintext, the random coins used by Alice for encryption, and any private keys held by Bob (or Alice) related to the ciphertext. In particular, Eve can use this information to "fully unroll" the *exact transcript* of some deterministic decryption procedure purportedly computed by Bob, as well as verify that the exact coins and decrypted plaintext in fact produce the coerced ciphertext. A secure deniable encryption scheme should maintain privacy of the sensitive data originally communicated between Alice and Bob under the coerced ciphertext (instead substituting a benign yet *convincing* plaintext in the view of Eve), even in the face of such a revealing attack and even if Alice and Bob may not interact during the coercion phase.

Historically, deniable encryption schemes have been challenging to construct. Under standard assumptions, Canetti et al. [14] constructed a sender-deniable[1] PKE where the distinguishing advantage between real and fake openings is an inverse polynomial depending on the public key size. But it was not until 2011 that O'Neill, Peikert, and Waters [35] proposed the first constructions of bi-deniable PKE with *negligible* deniability distinguishing advantage: from simulatable PKE generically, as well as from Learning with Errors (LWE [36]) directly.

Concurrently, Bendlin et al. [8] showed an inherent limitation: any non-interactive public-key encryption scheme may be receiver-deniable (resp. bi-deniable) only with *non-negligible* $\Omega(1/\mathsf{size}(\mathsf{pk}))$ distinguishing advantage in the deniability experiment. Indeed, O'Neill et al. [35] bypass the impossibility result of [8] by working in the so-called *flexible*[2] model of deniability. In the flexible of deniability, private keys sk are distributed by a central key authority. In the event that Bob is coerced to reveal a key sk that decrypts chosen ciphertext ct^*, the key authority distributes a *faking key* fk to Bob, which Bob can use to generate a fake key sk^* (designed to behave identically to sk except on ciphertext ct^*). If this step is allowed, then O'Neill et al. demonstrate that for their constructions, Eve has at most negligible advantage in distinguishing whether Bob revealed an honest sk or fake sk^*.

A major breakthrough in deniable encryption arrived with the work of Sahai and Waters [38], who proposed the first sender-deniable PKE with negligible distinguishing advantage from indistinguishability obfuscation ($i\mathcal{O}$) for P/poly [22]. The concept of deniability has been demonstrated useful in the contexts of leakage resilience [20], adaptive security for protocols, and as well as deniable computation (or algorithms) [16,19,23]. In addition to coercion resistance, a bi-deniable encryption scheme is a non-committing encryption scheme [15], as well as a scheme secure under selective opening (SOA) attacks [7], which are of independent theoretical interest.

Very recently, De Caro, Iovino, and O'Neill [17] gave various constructions of deniable *functional* encryption. First, they show a generic transformation of any

---

[1] We differentiate between sender-, receiver-, and bi-deniable schemes. A bi-deniable scheme is both sender- and receiver-deniable.

[2] We borrow the name "flexible" from Boneh, Lewi, and Wu [10] as the original term "multi-distributional" of O'Neill et al. [35] is used to define a slightly different security property in the recent work by De Caro et al. [17] than we achieve here.

IND-secure FE scheme for circuits into a flexibly receiver-deniable FE for circuits. Second, they give a direct construction of receiver-deniable FE for Boolean formulae from bilinear maps. Further, in the stronger *multi-distributional* model of deniable functional encryption – where there are special "deniable" set-up and encryption algorithms in addition to the plain ones, and where under coercion, it may non-interactively be made to seem as only the normal algorithms were used – De Caro et al. [17] construct receiver-deniable FE for circuits under the additional (powerful) assumption of different-inputs obfuscation ($di\mathcal{O}$).

De Caro et al. [17] also show (loosely speaking) that any receiver-deniable FE implies SIM-secure FE for the same functionality. Following [17], we also emphasize that deniability for functional encryption is a **strictly stronger** property than SIM security, since fixed coerced ciphertexts must decrypt correctly and benignly *in the real world*. Finally, we mention that in concurrent work, Apon, Fan, and Liu, in an unpublished work [5], construct flexibly bi-deniable inner product encryption from standard *lattice* assumptions. This work generalizes and thus subsumes the prior results of [5].

Despite the apparent theoretical utility in understanding the extent to which cryptographic constructions are deniable, our current knowledge of constructing such schemes from standard lattice assumptions is still limited. From LWE, we have only flexible and non-negligible advantage deniable encryption schemes (or IPE from [5]), whereas with the much more powerful assumption of indistinguishability obfuscation ($i\mathcal{O}$), we can obtain at least fully-secure sender-deniable PKE and computation [16,19,23], or as mentioned above even a multi-distributional receiver-deniable FE for all circuits from the even stronger assumption of $di\mathcal{O}$.

1.1 Our Contributions

In this work, we further narrow this gap by investigating a richer primitive – attribute-based encryption (ABE) [9,28,30,31] – *without* the use of obfuscation as a black box primitive. We hope that the techniques developed in this work can further shed light on deniability for even richer schemes such as functional encryption [9,12,22,29] under standard assumptions.

– Our main contribution is the construction of a flexibly bi-deniable ABE for poly-sized branching programs (which can compute NC1 via Barrington's theorem [6]) from the standard Learning with Errors assumption [36].

Theorem 1.1. *Under the standard* LWE *assumption, there is a flexibly bi-deniable attribute-based encryption scheme for all poly-size branching programs.*

Recall that in an attribute-based encryption (ABE) scheme for a family of functions $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$, every secret key sk_f is associated with a predicate $f \in \mathcal{F}$, and every ciphertext $\mathsf{ct}_{\boldsymbol{x}}$ is associated with an attribute $\boldsymbol{x} \in \mathcal{X}$. A ciphertext $\mathsf{ct}_{\boldsymbol{x}}$ can be decrypted by a given secret key sk_f to its payload message m only when $f(x) = 0 \in \mathcal{Y}$. Informally, the typical security notion for an ABE scheme is *collusion resistance*, which means no collection of keys can provide information on a ciphertext's message, if the individual keys are not authorized to decrypt

the ciphertext in the first place. Intuitively, a bi-deniable ABE must provide both collusion and coercion resistance.

Other contributions of this work can be summarized as:

– A new form of the Extended Learning with Errors (eLWE) assumption [2,13, 35], which is convenient in the context of Dual Regev type ABE/FE schemes that apply the Leftover Hash Lemma [21] in their security proofs.
– An explicit, tightened noise growth analysis for lattice-based ABE for branching programs. Prior work used the loose l_∞ norm to give a rough upper bound, which is technically insufficient to achieve deniability using our proof techniques. (We require matching upper *and lower* bounds on post-evaluation noise sizes.)

The eLWE assumption above is roughly the standard LWE assumption, but where the distinguisher also receives "hints" on the LWE sample's noise vector e in the form of inner products, i.e. distributions $\{\mathbf{A}, \boldsymbol{b} = \mathbf{A}^T \boldsymbol{s} + \boldsymbol{e}, \boldsymbol{z}, \langle \boldsymbol{z}, \boldsymbol{e} \rangle\}$ where (intuitively) $\boldsymbol{z}$ is a decryption key in the real system (which are denoted $\boldsymbol{r}$ elsewhere). Our contribution here is a new reduction from the standard LWE assumption to our correlated variant of extended-LWE, eLWE$^+$, where the adversary requests arbitrary correlations (expressed as a matrix $\mathbf{R}$) between the hints, in the case of a prime poly-size modulus with noise-less hints. We show this by extending the LWE to eLWE reduction of Alperin-Sheriff and Peikert [2] to our setting.

1.2   Our Approach

At a high level, our work begins with the ABE for branching programs of Gorbunov and Vinayagamurthy [30]. We will augment the basic ABE-BP = (Setup, Keygen, Enc, Dec) with an additional suite of algorithms (DenSetup, DenEnc, SendFake, RecFake) to form our flexibly bi-deniable ABE-BP. Doing so requires careful attention to the setting of parameters, as we explain in the sequel.

We remark now that – due to reasons related to the delicateness of our parameter setting – the ABE scheme of [30] is *particularly suited* to being made bi-deniable, as compared to similar schemes such as the ABE for arithmetic circuits of Boneh et al. [9]. We will explain this in what follows as well.

Intuition for Our New Deniability Mechanism. As in the work of O'Neill et al. [35], our approach to bi-deniability relies primarily on a curious property of Dual Regev type [24] secret keys: by correctness of any such scheme, each key $\boldsymbol{r}$ is guaranteed to behave as intended for some $1 - \mathsf{negl}(n)$ fraction of the possible random coins used to encrypt, but system parameters may be set so that each key is also guaranteed to be *faulty* (i.e. fail to decrypt) on some $\mathsf{negl}(n)$ fraction of the possible encryption randomness. More concretely, each secret key vector $\boldsymbol{r}$ in lattice-based schemes is sampled from an m-dimensional Gaussian distribution, as is the error term $\boldsymbol{e}$ (for LWE public key $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$). For every fixed $\boldsymbol{r}$, with overwhelming probability over the choice of $\boldsymbol{e}$, the vectors $\boldsymbol{r}, \boldsymbol{e} \in \mathbb{Z}_q^m$

will point in highly uncorrelated directions in m-space. However, if the vector $\boldsymbol{r}$ and $\boldsymbol{e}$ happen to point in similar directions, the error magnitude will be (loosely) *squared* during decryption.

Our scheme is based around the idea that a receiver, coerced on honest key-ciphertext pair $(\boldsymbol{r}, \mathsf{ct}^*)$, can use the key authority's *faking key* fk to learn the precise error vector $\boldsymbol{e}^*$ used to construct ct^*. Given $\boldsymbol{e}^*, \boldsymbol{r}$, and fk, the receiver re-samples a fresh secret key $\boldsymbol{r}^*$ that is functionally-equivalent to the honest key $\boldsymbol{r}$, except that $\boldsymbol{r}^*$ is strongly correlated with the vector $\boldsymbol{e}^*$ in ct^*. When the coercer then attempts to decrypt the challenge ciphertext ct^* using $\boldsymbol{r}^*$, the magnitude of decryption error will artificially grow and cause the decryption to output the value we want to deny to. Yet, when the coercer attempts to decrypt any other independently-sampled ciphertext ct, decryption will succeed with overwhelming probability under $\boldsymbol{r}^*$ if it would have under $\boldsymbol{r}$.

We emphasize that to properly show coercion resistance (when extending this intuition to the case of Dual Regev ABE instead of Dual Regev PKE), this behavior of $\boldsymbol{r}^*$ should hold *even when* ct *and* ct^* *embed the same attribute* $\boldsymbol{x}$. (Indeed, the majority of our effort is devoted to ensuring this simple geometric intuition allows a valid instantiation of the denying algorithms (DenSetup, DenEnc, SendFake, RecFake) without "damaging" the basic operation of (Setup, Keygen, Enc, Dec) in the underlying ABE scheme.)

Then, given the ability to "artificially blow-up" the decryption procedure of a specific key on a ciphertext-by-ciphertext basis, we can employ an idea originally due to Canetti et al. [14] of *translucent sets*, but generalized to the setting of ABE instead of PKE, to construct our new, flexibly bi-deniable ABE-BP scheme out of the framework provided by the "plain" SIM-secure ABE-BP scheme of [30].

Highlights of the Gorbunov-Vinayagamurthy Scheme. In the ABE for (width 5) branching programs of [30], bits a are "LWE-encoded" by the vector

$$\psi_{\mathbf{A}, \boldsymbol{s}, a} = \boldsymbol{s}^T (\mathbf{A} + a \cdot \mathbf{G}) + \boldsymbol{e} \in \mathbb{Z}_q^m$$

where $\mathbf{G}$ is the gadget matrix [34].

The ciphertext ct encrypting message μ under BP-input $\boldsymbol{x}$ is given by

$$\mathsf{ct} = (\psi_0, \psi^c, \{\psi_i\}_{i \in [\ell]}, \{\psi_{0,i}\}_{i \in [5]}, c),$$

and is composed of a Dual Regev ct-pair of vectors (ψ_0, c) encrypting the ciphertext's message μ, an encoding ψ^c representing the (freshly randomized) encoding of the constant 1, five encodings $\{\psi_{0,i}\}_{i \in [5]}$ representing a (freshly randomized) encoding of the initial state of a width-5, length-ℓ branching program BP, and ℓ encodings $\{\psi_i\}_{i \in [\ell]}$ – one for each step of the branching program's evaluation, storing a constant-sized *permutation matrix* associated with the i-th level of BP. Note that each "LWE encoding" ψ is performed under a distinct public key matrix $\mathbf{A}, \mathbf{A}^c, \{\mathbf{A}_i\}$, or $\{\mathbf{A}_{0,i}\}$ respectively.

The (key-homomorphic) evaluation procedure takes as input a ciphertext $\mathsf{ct} = (\psi_0, \psi^c, \{\psi_i\}, \{\psi_{0,i}\}, c)$ and the public key $\mathsf{pk} = (\mathbf{A}, \mathbf{A}^c, \{\mathbf{A}_i\}, \{\mathbf{A}_{0,i}\})$, as well as the cleartext branching program description BP and the BP-input $\boldsymbol{x}$.

It produces the evaluated public key $\mathbf{V}_{\mathsf{BP}}$ and the evaluated encoding $\psi_{\mathsf{BP}(\boldsymbol{x})}$. Given a short secret key *vector* $\boldsymbol{r} \in \mathbb{Z}^{2m}$ matching (some public coset $\boldsymbol{u}$ of) the lattice generated by $[\mathbf{A}|\mathbf{V}_{\mathsf{BP}}] \in \mathbb{Z}^{n \times 2m}$, the encoding *vector* $\psi_{\mathsf{BP}(\boldsymbol{x})}$ (whose Dual Regev encoding-components (ψ_0, c) also match coset $\boldsymbol{u}$) can be decrypted to the message μ if and only if $\mathsf{BP}(\boldsymbol{x}) = \mathsf{accept} = \mathbf{0}$.

On the Necessity of Exact Noise Control. In order to push the intuition for our deniability mechanism through for an ABE of the above form, we must overcome a number of technical hurdles.

The major challenge is an implicit technical requirement to *very tightly control* the precise noise magnitude of evaluated ciphertexts. In previous functional (and homomorphic) encryption schemes from lattices, the emphasis is placed on upper bounding evaluated noise terms, to ensure that they do not grow too large and cause decryption to fail. Moreover, security (typically) holds for any ciphertext noise level at or above the starting ciphertexts' noises. In short, noise growth during evaluation is nearly always undesirable.

As with previous schemes, we too must upper bound the noise growth of evaluated ciphertexts in order to ensure basic correctness of our ABE. But unlike previous schemes, we must take the step of also (carefully) *lower bounding* the noise growth during the branching program evaluation (which technically motivates deviating from the l_∞ norm of prior analyses). This is due to the fact, highlighted above, that producing directional alignment between a key and error term can at most *square* the noise present during decryption. Since coercion resistance requires that it must always be possible to deny any ciphertext originally intended for any honest key, it must be that, with overwhelming probability, every honest key and every honest ciphertext produce evaluated error that is no less than the square root of the maximum noise threshold tolerated.

In a little more detail – as we will later demonstrate in Sect. 4 – in dimension m there is precisely an expected $\mathsf{poly}(m)$ *gap* in magnitude between the inner products of **(i)** two relatively *orthogonal* key/error vectors $\boldsymbol{r}, \boldsymbol{e}_{\mathsf{BP}(\boldsymbol{x})}$, and **(ii)** two highly *correlated* key/error vectors $\boldsymbol{r}^*, \boldsymbol{e}_{\mathsf{BP}(\boldsymbol{x})}$. The ability to deny is based around our ability to design $\boldsymbol{r}^*$ that are statistically indistinguishable from $\boldsymbol{r}$ in the attacker's view, but where $\boldsymbol{r}^*$ "punctures out" decryptions of ciphertexts with error vectors pointing in the *direction* of $\boldsymbol{e}_{\mathsf{BP}(\boldsymbol{x})}$ in m-space (error-vector directions are unique to each honest ct with overwhelming probability).

Crucially, this approach **generically forces** the use of a *polynomial-sized modulus q* in the scheme.[3] In particular, when error vectors $\boldsymbol{e}$ may (potentially) grow to be some superpolynomial magnitude in the dimension m of the public/secret keys, we totally lose any efficiently testable notion of "error vector orientation in m-space" for the purposes of Dual Regev type decryption.

Further, in order to "correctly trace and distinguish" different orientations throughout the computation of an arbitrary branching program BP, we are

---

[3] One consequence of a poly-size modulus requirement is that the fully key-homomorphic scheme of Boneh et al. [9], taken verbatim, can only be denied for up to $\mathsf{NC0}$ functions using our approach. Past this, attempts to produce fake keys in an identical manner to this work may be detected by a statistical test under coercion.

required to make careful use of *multi-dimensional Gaussian* distributions. These are sampled using covariance matrices $\mathbf{Q} \in \mathbb{Z}^{m \times m}$ that allow us to succinctly describe the underlying, geometric *randomized rotation* action on error vector *orientations* in m-space with each arithmetic operation of the BP evaluation in the overall ABE-BP scheme. (We use the geometrically-inspired term "rotation matrix" to describe our low-norm matrices $\mathbf{R}$ for this reason.)

An additional subtlety in our new noise analysis is that we require the individual multiplications of the ct evaluation procedure to have *independently sampled* error vectors in each operand-encoding – and thus be "independently oriented" – in order for the overall analysis to go through correctly. (While there could in principle be some way around this technical obstacle in the analysis, we were unable to find one.) This appears to a priori exclude a straightforward denying procedure for *all circuits* [9], where a gate's input wires' preceding sub-circuits may have cross-wires between them. But it naturally permits denying *branching program computations*, where at the i-th time-step, an i-th *independently generated* ct-component is merged into an accumulated BP state, as with [30].

Finally, we mention that an inherent limitation in the techniques of Apon et al. [5], used to construct (the weaker notion of) flexibly bi-deniable inner product encryption from LWE, is bypassed in the current work at the cost of supporting only BP computations of an a-priori bounded length ℓ. Namely, it was the case in [5] that the *length ℓ* of the attribute vector $\boldsymbol{w}$ had to be "traded off" against the dimension m of the public/secret keys. We suppress the details, other than to point out that this issue can be resolved by artificially boosting the magnitude of the low-norm matrices used to generate error terms in fresh ciphertexts from $\{-1, 1\}$ up to $\{-\Theta(m\ell), \Theta(m\ell)\}$-valued matrices. This, of course, requires knowing the length ℓ of the branching program up front. (Intuitively, this technical change as compared to [5] allows for a sharp *inductive lower bound* on the *minimum* noise growth across all possible function-input pairs that might be evaluated in a given instance of our bi-deniable ABE-BP scheme.)

1.3 Future Directions

The next, most natural question is whether bi-deniable functional encryption can be built out of similar techniques (from only LWE), perhaps by leveraging our bi-deniable ABE for NC1 computations as a building block. We briefly sketch one possible approach and the obstacles encountered. Recall that Goldwasser et al. [27] show to transform the combination of **(i)** any ABE for a circuit family $\mathcal{C}$, **(ii)** fully homomorphic encryption, and **(iii)** a randomized encoding scheme (such as Yao's garbled circuits) into a 1-key (resp. *bounded collusion*) SIM-secure functional encryption scheme for $\mathcal{C}$.

If we instantiate the Goldwasser et al. transformation with our deniable ABE, we get a functional encryption scheme for NC1. We can then boost functional encryption for shallow circuits to functional encryption for all circuits using the "trojan method" of Ananth et al. [4]. As it turns out, it is easy to directly prove *flexible receiver-deniability* of the final scheme, independently of but matching the generic results of De Caro et al. [17] for receiver-deniable FE.

Unfortunately, we do not know how to prove (even, flexible) *sender-deniability* of this final scheme. Roughly speaking, the problem is that each ciphertext's attribute in such a scheme contains an FHE ciphertext $\mathsf{ct_{FHE}}$ for its attribute, and this attribute leaks to the attacker (resp. cocercer) on decryptions that succeed. In particular, there is nothing stopping the coercer from demanding that the sender also provide randomness r_S that *opens the attribute's FHE ciphertext.*

We speculate that a possible way around this obstacle would be to use an *adaptively-secure* homomorphic encryption scheme for $\mathsf{NC1}$ computations. Note that adaptively-secure FHE is known to be impossible for circuits with $\omega(\log(n))$ depth due to a counting argument lower bound by Katz, Thiruvengadam, and Zhou [32], but this leaves open the possibility of an $\mathsf{NC1}$-homomorphic encryption scheme with the necessary properties to re-obtain (flexible) sender deniability for lattice-based FE. We leave this as an intriguing open problem for future work.

2   Preliminaries

Notations. Let PPT denote probabilistic polynomial time. We use bold upper-case letters to denote matrices, and bold lowercase letters to denote vectors, where vectors are by default column vectors throughout the paper. We let λ be the security parameter, $[n]$ denote the set $\{1, ..., n\}$, and $|\boldsymbol{t}|$ denote the number of bits in a string or vector $\boldsymbol{t}$. We denote the i-th bit value of a string $\boldsymbol{s}$ by $\boldsymbol{s}[i]$. We use $[\cdot|\cdot]$ to denote the concatenation of vectors or matrices, and $|| \cdot ||$ to denote the norm of vectors or matrices respectively. We use the ℓ_2 norm for all vectors unless explicitly stated otherwise.

We present necessary background knowledge of branching programs and lattices (such as the LWE assumption and lattice sampling algorithms) in full version.

Randomness Extraction. We will use the following lemma to argue the indistinghishability of two different distributions, which is a generalization of the leftover hash lemma proposed by Dodis et al. [21].

Lemma 2.1 ([1]).   *Suppose that $m > (n + 1)\log q + w(\log n)$. Let $\mathbf{R} \in \{-1, 1\}^{m \times k}$ be chosen uniformly at random for some polynomial $k = k(n)$. Let $\mathbf{A}, \mathbf{B}$ be matrix chosen randomly from $\mathbb{Z}_q^{n \times m}, \mathbb{Z}_q^{n \times k}$ respectively. Then, for all vectors $\boldsymbol{w} \in \mathbb{Z}^m$, the two following distributions are statistically close:*

$$(\mathbf{A}, \mathbf{A}\mathbf{R}, \boldsymbol{w}^T\mathbf{R}) \approx (\mathbf{A}, \mathbf{B}, \boldsymbol{w}^T\mathbf{R})$$

Learning With Errors. The LWE problem was introduced by Regev [36], who showed that solving it *on the average* is as hard as (quantumly) solving several standard lattice problems *in the worst case.*

Definition 2.2 (LWE). *For an integer $q = q(n) \geq 2$, and an error distribution $\chi = \chi(n)$ over $\mathbb{Z}_q$, the learning with errors problem $\mathsf{LWE}_{n,m,q,\chi}$ is to distinguish between the following pairs of distributions:*

$$\{\mathbf{A}, \boldsymbol{b} = \mathbf{A}^T \boldsymbol{s} + \boldsymbol{e}\} \ and \ \{\mathbf{A}, \boldsymbol{u}\}$$

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

Trapdoors and Sampling Algorithms. We will use the algorithms TrapGen, SampleLeft, SampleRight, ExtBasis, Invert first proposed in [1,18,24,34] to sample short vectors from specified lattices. For details of these sampling algorithms, please refer to the full version.

3 New Definitions and Tools

In this section, we first describe our new notion of flexibly bi-deniable ABE, which is a natural generalization of the flexibly bi-deniable PKE of [35]. Then we define the notion of a flexibly attribute-based bi-translucent set (AB-BTS), which generalizes the idea of bi-translucent set (BTS) in the work [35]. Using a similar argument as in the work [35], we can show that an AB-BTS suffices to construct bi-deniable ABE. In the last part of this section, we define a new assumption called Extended LWE Plus, and show its hardness by giving a reduction from the standard LWE problem.

3.1 Flexibly Bi-Deniable ABE: Syntax and Deniability Definition

A flexibly bi-deniable key-policy attribute based encryption for a class of Boolean circuits $\mathcal{C} : \{0,1\}^\ell \to \{0,1\}$ consists a tuple of PPT algorithms $\Pi =$ (Setup, Keygen, Enc, Dec, DenSetup, DenEnc, SendFake, RecFake). We describe them in detail as follows:

Setup(1^λ): On input the security parameter λ, the setup algorithm outputs public parameters pp and master secret key msk.

Keygen(msk, f): On input the master secret key msk and a function $f \in \mathcal{C}$, it outputs a secret key sk_f.

Enc(pp, $\boldsymbol{x}, \mu; r_S$): On input the public parameter pp, an attribute/message pair $(\boldsymbol{x}, \mu)$ and randomness r_S, it outputs a ciphertext $c_{\boldsymbol{x}}$.

Dec($\mathsf{sk}_f, c_{\boldsymbol{x}}$): On input the secret key sk_f and a ciphertext $c_{\boldsymbol{x}}$, it outputs the corresponding plaintext μ if $f(\boldsymbol{x}) = 0$; otherwise, it outputs $\perp$.

DenSetup(1^λ): On input the security parameter λ, the deniable setup algorithm outputs pubic parameters pp, master secret key msk and faking key fk.

DenEnc(pp, $\boldsymbol{x}, \mu; r_S$): On input the public parameter pp, an attribute/message pair $(\boldsymbol{x}, \mu)$ and randomness r_S, it outputs a ciphertext $c_{\boldsymbol{x}}$.

SendFake(pp, r_S, μ, μ'): On input public parameters pp, original random coins r_S, message μ of DenEnc and desired message μ', it outputs a faked random coin r_S'.

RecFake(pp, fk, $c_{\boldsymbol{x}}, f, \mu'$): On input public parameters pp, faking key fk, a ciphertext $c_{\boldsymbol{x}}$, a function $f \in \mathcal{C}$, and desired message μ', the receiver faking algorithm outputs a faked secret key sk_f'.

Correctness. We say the flexibly bi-deniable ABE scheme described above is correct, if for any $(\mathsf{msk}, \mathsf{pp}) \leftarrow \mathsf{S}(1^\lambda)$, where $\mathsf{S} \in \{\mathsf{Setup}, \mathsf{DenSetup}\}$, any message μ, function $f \in \mathcal{C}$, and any attribute vector $\boldsymbol{x}$ where $f(\boldsymbol{x}) = 0$, we have $\mathsf{Dec}(\mathsf{sk}_f, c_{\boldsymbol{x}}) = \mu$, where $\mathsf{sk}_f \leftarrow \mathsf{Keygen}(\mathsf{msk}, f)$ and $c_{\boldsymbol{x}} \leftarrow \mathsf{E}(\mathsf{pp}, \boldsymbol{x}, \mu; r_S)$ where $\mathsf{E} \in (\mathsf{Enc}, \mathsf{DenEnc})$.

Bi-Deniability Definition. Let μ, μ' be two arbitrary messages, not necessarily different. We propose the bi-deniability definition by describing real experiment $\mathbf{Expt}_{\mathcal{A},\mu,\mu'}^{\mathsf{Real}}(1^\lambda)$ and faking experiment $\mathbf{Expt}_{\mathcal{A},\mu,\mu'}^{\mathsf{Fake}}(1^\lambda)$ regarding adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ as shown in Fig. 1:

where $\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)$ returns a secret key $\mathsf{sk}_f \leftarrow \mathsf{Keygen}(\mathsf{msk}, f)$ if $f(\boldsymbol{x}^*) \neq 0$ and $\perp$ otherwise.

1. $(\boldsymbol{x}^*, \mathsf{state}_1) \leftarrow \mathcal{A}_1(\lambda)$
2. $(\mathsf{pp}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda)$
3. $c'_{\boldsymbol{x}^*} \leftarrow \mathsf{Enc}(\mathsf{pp}, \boldsymbol{x}^*, \mu; r_S)$
4. $(f^*, \mathsf{state}_2) \leftarrow \mathcal{A}_2^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{pp}, \mathsf{state}_1, c'_{\boldsymbol{x}^*})$
5. $\mathsf{sk}_{f^*} \leftarrow \mathsf{Keygen}(\mathsf{msk}, f^*)$
6. $b \leftarrow \mathcal{A}_3^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{sk}_{f^*}, \mathsf{state}_2, r_S)$
7. Output $b \in \{0, 1\}$
   (a) $\mathbf{Expt}_{\mathcal{A}}^{\mathsf{Real}}(1^\lambda)$

1. $(\boldsymbol{x}^*, \mathsf{state}_1) \leftarrow \mathcal{A}_1(\lambda)$
2. $(\mathsf{pp}, \mathsf{msk}, \mathsf{fk}) \leftarrow \mathsf{DenSetup}(1^\lambda)$
3. $c'_{\boldsymbol{x}^*} \leftarrow \mathsf{DenEnc}(\mathsf{pp}, \boldsymbol{x}^*, \mu'; r_S)$
4. $(f^*, \mathsf{state}_2) \leftarrow \mathcal{A}_2^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{pp}, \mathsf{state}_1, c'_{\boldsymbol{x}^*})$
5. $r'_S \leftarrow \mathsf{SendFake}(\mathsf{pp}, \mu, \mu', r_S)$
6. $\mathsf{sk}_{f^*} \leftarrow \mathsf{RecFake}(\mathsf{pp}, \mathsf{fk}, c'_{\boldsymbol{x}}, \boldsymbol{v}^*, \mu')$
7. $b \leftarrow \mathcal{A}_3^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{sk}_{f^*}, \mathsf{state}_2, r'_S)$
8. Output $b \in \{0, 1\}$
   (b) $\mathbf{Expt}_{\mathcal{A}}^{\mathsf{Fake}}(1^\lambda)$

Fig. 1. Security experiments for bi-deniable ABE

Definition 3.1 (Flexibly Bi-Deniable ABE). *An ABE scheme Π is bi-deniable if for any two messages μ, μ', any probabilistic polynomial-time adversaries $\mathcal{A}$ where $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$, there is a negligible function $\mathsf{negl}(\lambda)$ such that*

$$\mathbf{Adv}_{\mathcal{A},\mu,\mu'}^{\Pi}(1^\lambda) = |\mathbf{Pr}[\mathbf{Expt}_{\mathcal{A},\mu,\mu'}^{\mathsf{Real}}(1^\lambda) = 1] - \mathbf{Pr}[\mathbf{Expt}_{\mathcal{A},\mu,\mu'}^{\mathsf{Fake}}(1^\lambda) = 1]| \leq \mathsf{negl}(\lambda)$$

3.2 Attribute Based Bitranslucent Set Scheme

In this section, we define the notion of a *Attribute Based Bitranslucent Set* (AB-BTS), which is an extension of bitranslucent sets (BTS) as defined by O'Neill et al. in [35]. Our new notion permits a more fine-grained degree of access control, where pseudorandom samples and secret keys are associated with attributes $\boldsymbol{x}$, and the testing algorithm can successfully distinguish a pseudorandom sample from a truly random one if and only if the attribute of the sample is accepted under a given secret key's policy f – i.e. when $f(\boldsymbol{x}) = \mathbf{0}$. This concept is reminiscent of *attribute-based encryption* (ABE), and in fact, we will show in the sequel how to construct a flexibly bi-deniable ABE from an AB-BTS. This is analogous to the construction of a flexibly bi-deniable PKE from O'Neill et al.'s BTS. We present the formal definition below.

Let $\mathcal{F}$ be some family of functions. An attribute based bitranslucent set (AB-BTS) scheme for $\mathcal{F}$ consists of the following algorithms:

Setup(1^λ): On input the security parameter, the normal setup algorithm outputs a public parameter pp and master secret key msk.

DenSetup(1^λ): On input the security parameter, the deniable setup algorithm outputs a public parameter pp, master secret key msk and faking key fk.

Keygen(msk, f): On input the master secret key msk and a function $f \in \mathcal{F}$, the key generation algorithm outputs a secret key sk_f.

P - and U-samplers SampleP(pp, $\boldsymbol{x}; r_S$) and SampleU(pp, $\boldsymbol{x}; r_S$) output some $\boldsymbol{c}$.

TestP($\mathsf{sk}_f, \boldsymbol{c_x}$): On input a secret key sk_f and a ciphertext $\boldsymbol{c_x}$, the P-tester algorithm outputs 1 (accepts) or 0 (rejects).

FakeSCoins(pp, r_S): On input a public parameters pp and randomness r_S, the sender-faker algorithm outputs randomness r_S^*.

FakeRCoins(pp, fk, $\boldsymbol{c_x}, f$): On input a public parameters pp, the faking key fk, a ciphertext $\boldsymbol{c_x}$ and a function $f \in \mathcal{F}$, the receiver-faker algorithm outputs a faked secret key sk_f'.

Definition 3.2 (AB-BTS). *We say a scheme $\Pi =$ (Setup, DenSetup, Keygen, SampleP, SampleU, TestP, FakeSCoins, FakeRCoins) is an AB-BTS scheme for a function family $\mathcal{F}$ if it satisfies:*

1. *(Correctness.) The following experiments accept or respectively reject with overwhelming probability over the randomness.*
   - *Let* (pp, msk) $\leftarrow$ Setup(1^λ), $f \in \mathcal{F}$, $\mathsf{sk}_f \leftarrow$ Keygen(msk, f). *If* $f(\boldsymbol{x}) = 0$ *and* $\boldsymbol{c_x} \leftarrow$ SampleP(pp, $\boldsymbol{x}; r_S$), *then* TestP($\mathsf{sk}_f, \boldsymbol{c_x}$) = 1; *otherwise,* TestP($\mathsf{sk}_f, \boldsymbol{c_x}$) = 0.
   - *Let* (pp, msk) $\leftarrow$ Setup(1^λ), $f \in \mathcal{F}$, $\mathsf{sk}_f \leftarrow$ Keygen(msk, f), $\boldsymbol{c} \leftarrow$ SampleU(pp; r_S). *Then* TestP($\mathsf{sk}_f, \boldsymbol{c}$) = 0.
2. *(Indistinguishable public parameters.) The public parameters pp generated by the two setup algorithms* (pp, msk) $\leftarrow$ Setup(1^λ) *and* (pp, msk, fk) $\leftarrow$ DenSetup(1^λ) *should be indistinguishable.*
3. *(Selective bi-deniability.) Let $\mathcal{F}$ be a family of functions. We define the following two experiments: the real experiment $\mathbf{Expt}_{\mathcal{A},\mathcal{F}}^{\mathsf{Real}}(1^\lambda)$ and the faking experiment $\mathbf{Expt}_{\mathcal{A},\mathcal{F}}^{\mathsf{Fake}}(1^\lambda)$ regarding an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ as shown in Fig. 2:*
   where KG(msk, $\boldsymbol{x}^*, \cdot$) *returns a secret key $\mathsf{sk}_f \leftarrow$ Keygen(msk, f) if $f \in \mathcal{F}$ and $f(\boldsymbol{x}^*) \neq 0$; it returns $\perp$ otherwise. We also require that $f^* \in \mathcal{F}$.*
   We say the scheme is selectively bi-deniable for $\mathcal{F}$, if for any probabilistic polynomial-time adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$, there is a negligible function $\mathsf{negl}(\lambda)$ such that

$$\mathbf{Adv}_{\mathcal{A}}^{\Pi}(1^\lambda) = |\mathbf{Pr}[\mathbf{Expt}_{\mathcal{A},\mathcal{F}}^{\mathsf{Real}}(1^\lambda) = 1] - \mathbf{Pr}[\mathbf{Expt}_{\mathcal{A},\mathcal{F}}^{\mathsf{Fake}}(1^\lambda) = 1]| \leq \mathsf{negl}(\lambda)$$

Remark 3.3. *Correctness for the faking algorithms is implied by the bi-deniability property. In particular, with overwhelming probability over the overall randomness, the following holds: let* (pp, msk, fk) $\leftarrow$ DenSetup(1^λ), $f \in \mathcal{F}$, $\mathsf{sk}_f \leftarrow$ Keygen(msk, f), $\boldsymbol{x}$ *be a string and* $\boldsymbol{c_x} \leftarrow$ SampleP(pp, $x; r_S$), *then*

(a) $(f^*, \boldsymbol{x}^*, \mathsf{state}_1) \leftarrow \mathcal{A}_1(\lambda)$
(b) $(\mathsf{pp}, \mathsf{msk}, \mathsf{fk}) \leftarrow \mathsf{DenSetup}(1^\lambda)$
(c) $\boldsymbol{c} \leftarrow \mathsf{SampleU}(\mathsf{pp}, \boldsymbol{x}^*; r_S)$
(d) $\mathsf{state}_2 \leftarrow \mathcal{A}_2^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{pp}, \mathsf{state}_1, \boldsymbol{c})$
(e) $\mathsf{sk}_{f^*} \leftarrow \mathsf{Keygen}(\mathsf{msk}, f^*)$
(f) $b \leftarrow \mathcal{A}_3^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{sk}_{f^*}, c, \mathsf{state}_2, r_S)$
(g) Output $b \in \{0, 1\}$
        (a) $\mathbf{Expt}_{\mathcal{A}}^{\mathsf{Real}}(1^\lambda)$

(a) $(f^*, \boldsymbol{x}^*, \mathsf{state}_1) \leftarrow \mathcal{A}_1(\lambda)$
(b) $(\mathsf{pp}, \mathsf{msk}, \mathsf{fk}) \leftarrow \mathsf{DenSetup}(1^\lambda)$
(c) $\boldsymbol{c} \leftarrow \mathsf{SampleP}(\mathsf{pp}, \boldsymbol{x}^*; r_S)$
(d) $\mathsf{state}_2 \leftarrow \mathcal{A}_2^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{pp}, \mathsf{state}_1, \boldsymbol{c})$
(e) $r_S' \leftarrow \mathsf{FakeSCoins}(\mathsf{pp}, r_S)$
(f) $\mathsf{sk}_{f^*} \leftarrow \mathsf{FakeRCoins}(\mathsf{pp}, \mathsf{fk}, \boldsymbol{c}, f^*)$
(g) $b \leftarrow \mathcal{A}_3^{\mathsf{KG}(\mathsf{msk}, \boldsymbol{x}^*, \cdot)}(\mathsf{sk}_{f^*}, c, \mathsf{state}_2, r_S')$
(h) Output $b \in \{0, 1\}$
        (b) $\mathbf{Expt}_{\mathcal{A}}^{\mathsf{Fake}}(1^\lambda)$

Fig. 2. Security experiments for AB-BTS

– $\mathsf{SampleU}(\mathsf{pp}, \boldsymbol{x}; \mathsf{FakeSCoins}(\mathsf{pp}, r_S)) = \boldsymbol{c_x}$,
– $\mathsf{TestP}(\mathsf{FakeRCoins}(\mathsf{pp}, \mathsf{fk}, \boldsymbol{c_x}, f), \boldsymbol{c_x}) = 0$
– *For any other x', let $\boldsymbol{c}' \leftarrow \mathsf{SampleP}(\mathsf{pp}, x'; r_S')$, then (with overwhelming probability) we have*

$$\mathsf{TestP}\left(\mathsf{FakeRCoins}(\mathsf{pp}, \mathsf{fk}, \boldsymbol{c_x}, f), \boldsymbol{c}'\right) = \mathsf{TestP}(\mathsf{sk}_f, c').$$

It is not hard to see that if one of these does not hold, then one can easily distinguish the real experiment from the faking experiment.

Remark 3.4. *Canetti et al. [14] gave a simple encoding technique to construct a sender-deniable encryption scheme from a translucent set. O'Neill, Peikert, and Waters [35] used a similar method to construct a flexibly bi-deniable encryption from a bi-translucent set scheme. Here we further observe that the same method as well allows us to construct a flexibly bi-deniable ABE scheme from bi-deniable AB-BTS. We present the construction in Sect. 4.4.*

3.3   Extended LWE and Our New Variant

O'Neill et al. [35] introduced the Extended LWE problem, which allows a "hint" on the error vector $\boldsymbol{x}$ to leak in form of a noisy inner product. They observe a trivial "blurring" argument shows that LWE reduces to eLWE when the hint-noise βq is superpolynomially larger than the magnitude of samples from χ, and also allows for unboundedly many *independent* hint vectors $\langle \boldsymbol{z}, \boldsymbol{x}_i \rangle$ while retaining LWE-hardness.

Definition 3.5 (Extended LWE). *For an integer $q = q(n) \geq 2$, and an error distribution $\chi = \chi(n)$ over $\mathbb{Z}_q$, the extended learning with errors problem $\mathsf{eLWE}_{n,m,q,\chi,\beta}$ is to distinguish between the following pairs of distributions:*

$$\{\mathbf{A}, \boldsymbol{b} = \mathbf{A}^T \boldsymbol{s} + \boldsymbol{e}, \boldsymbol{z}, \langle \boldsymbol{z}, \boldsymbol{b} - \boldsymbol{e} \rangle + e'\} \ and \ \{\mathbf{A}, \boldsymbol{u}, \boldsymbol{z}, \langle \boldsymbol{z}, \boldsymbol{u} - \boldsymbol{x} \rangle + e'\}$$

where $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\boldsymbol{s} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$, $\boldsymbol{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$, $\boldsymbol{e}, \boldsymbol{z} \overset{\$}{\leftarrow} \chi^m$ and $e' \overset{\$}{\leftarrow} \mathcal{D}_{\beta q}$.

Further, Alperin-Sheriff and Peikert [2] show that LWE reduces to eLWE with a polynomial modulus and no hint-noise (i.e. $\beta = 0$), even in the case of a bounded number of *independent* hints.

We introduce the following new form of extended-LWE, called eLWE$^+$, which considers leaking a pair of *correlated hints* on the same noise vector. Our security proof of the AB-BTS construction relies on this new assumption.

Definition 3.6 (Extended LWE Plus). *For integer* $q = q(n) \geq 2, m = m(n)$, *an error distribution* $\chi = \chi(n)$ *over* $\mathbb{Z}_q$, *and a matrix* $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$, *the extended learning with errors problem* eLWE$^+_{n,m,q,\chi,\beta,\mathbf{R}}$ *is to distinguish between the following pairs of distributions:*

$$\{\mathbf{A}, \boldsymbol{b} = \mathbf{A}^T \boldsymbol{s} + \boldsymbol{e}, \boldsymbol{z}_0, \boldsymbol{z}_1, \langle \boldsymbol{z}_0, \boldsymbol{b} - \boldsymbol{e} \rangle + e, \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{b} - \boldsymbol{e} \rangle + e'\} \ and$$

$$\{\mathbf{A}, \boldsymbol{u}, \boldsymbol{z}_0, \boldsymbol{z}_1, \langle \boldsymbol{z}_0, \boldsymbol{u} - \boldsymbol{e} \rangle + e, \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{u} - \boldsymbol{e} \rangle + e'\}$$

where $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\boldsymbol{s} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$, $\boldsymbol{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$, $\boldsymbol{e}, \boldsymbol{z}_0, \boldsymbol{z}_1 \overset{\$}{\leftarrow} \chi^m$ and $e, e' \overset{\$}{\leftarrow} \mathcal{D}_{\beta q}$.

Hardness of Extended-LWE$^+$. A simple observation, following prior work, is that when χ is poly(n)-bounded and the hint noise βq (and thus, modulus q) is superpolynomial in n, then LWE$_{n,m,q,\chi}$ trivially reduces to eLWE$^+_{n,m,q,\chi,\beta,\mathbf{R}}$ for every $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$ so that $\mathbf{R}\boldsymbol{z}_1$ has poly(n)-bounded norm. This is because, for any $r = \omega(\sqrt{\log n}), c \in \mathbb{Z}$, the statistical distance between $\mathcal{D}_{\mathbb{Z},r}$ and $c + \mathcal{D}_{\mathbb{Z},r}$ is at most $O(|c|/r)$.

However, our cryptosystem will require a polynomial-size modulus q. So, we next consider the case of *prime* modulus q of poly(n) size and no noise on the hints (i.e. $\beta = 0$). Following [2][4], it will be convenient to swap to the "knapsack" form of LWE, which is: given $\mathbf{H} \leftarrow \mathbb{Z}_q^{(m-n) \times m}$ and $\boldsymbol{c} \in \mathbb{Z}_q^{m-n}$, where either $\boldsymbol{c} = \mathbf{H}\boldsymbol{e}$ for $\boldsymbol{e} \leftarrow \chi^m$ or $\boldsymbol{c}$ uniformly random and independent of $\mathbf{H}$, determine which is the case (with non-negligible advantage). The "extended-plus" form of the knapsack problem also reveals a pair of hints $(\boldsymbol{z}_0, \boldsymbol{z}_1, \langle \boldsymbol{z}_0, \boldsymbol{e} \rangle, \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e} \rangle)$. Note the equivalence between LWE and knapsack-LWE is proven in [33] for $m \geq n + \omega(\log n)$.

Theorem 3.7. *For* $m \geq n + \omega(\log n)$, *for every prime* $q = \text{poly}(n)$, *for every* $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$, *and for every* $\beta \geq 0$, $\mathbf{Adv}_{\mathcal{BA}}^{\text{LWE}_{n,m,q,\chi}}(1^\lambda) \geq (1/q^2)\mathbf{Adv}_{\mathcal{A}}^{\text{eLWE}^+_{n,m,q,\chi,\beta,\mathbf{R}}}(1^\lambda)$.

---

[4] We note that a higher quality reduction from LWE to eLWE is given in [13] in the case of binary secret keys. However for our cryptosystem, it will be more convenient to have secret key coordinates in $\mathbb{Z}_q$, so we extend the reduction of [2] to eLWE$^+$ instead.

Proof. We construct an LWE to eLWE$^+$ reduction $\mathcal{B}$ as follows. $\mathcal{B}$ receives a knapsack-LWE instance $\mathbf{H} \in \mathbb{Z}_q^{(m-n) \times m}, \boldsymbol{c} \in \mathbb{Z}_q^{m-n}$. It samples $\boldsymbol{e}', \boldsymbol{z}_0, \boldsymbol{z}_1 \leftarrow \chi^m$ and uniform $\boldsymbol{v}_0, \boldsymbol{v}_1 \leftarrow \mathbb{Z}_q^{m-n}$. It chooses any $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$, then sets

$$\mathbf{H}' := \mathbf{H} - \boldsymbol{v}_0 \boldsymbol{z}_0^T - \boldsymbol{v}_1 \left(\mathbf{R}\boldsymbol{z}_1\right)^T \in \mathbb{Z}_q^{(m-n) \times m},$$
$$\boldsymbol{c}' := \boldsymbol{c} - \boldsymbol{v}_0 \cdot \langle \boldsymbol{z}_0, \boldsymbol{e}' \rangle - \boldsymbol{v}_1 \cdot \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e}' \rangle \in \mathbb{Z}_q^{m-n}.$$

It sends $(\mathbf{H}', \boldsymbol{c}', \boldsymbol{z}_0, \boldsymbol{z}_1, \langle \boldsymbol{z}_0, \boldsymbol{e}' \rangle, \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e}' \rangle)$ to the knapsack-eLWE$^+$ adversary $\mathcal{A}$, and outputs what $\mathcal{A}$ outputs.

Notice that when $\mathbf{H}, \boldsymbol{c}$ are independent and uniform, so are $\mathbf{H}', \boldsymbol{c}'$, in which case $\mathcal{B}$'s simulation is perfect.

Now, consider the case when $\mathbf{H}, \boldsymbol{c}$ are drawn from the knapsack-LWE distribution, with $\boldsymbol{c} = \mathbf{H}\boldsymbol{x}$ for $\boldsymbol{e} \leftarrow \chi^m$. In this case, $\mathbf{H}'$ is uniformly random over the choice of $\mathbf{H}$, and we have

$$\boldsymbol{c}' = \mathbf{H}\boldsymbol{x} - \boldsymbol{v}_0 \cdot \langle \boldsymbol{z}_0, \boldsymbol{e}' \rangle - \boldsymbol{v}_1 \cdot \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e}' \rangle$$
$$= \left(\mathbf{H}' + \boldsymbol{v}_0 \boldsymbol{z}_0^T + \boldsymbol{v}_1 \left(\mathbf{R}\boldsymbol{z}_1\right)^T\right) \boldsymbol{e} - \boldsymbol{v}_0 \cdot \langle \boldsymbol{z}_0, \boldsymbol{e}' \rangle - \boldsymbol{v}_1 \cdot \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e}' \rangle$$
$$= \mathbf{H}'\boldsymbol{e} + \boldsymbol{v}_0 \cdot \langle \boldsymbol{z}_0, \boldsymbol{e} - \boldsymbol{e}' \rangle + \boldsymbol{v}_1 \cdot \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e} - \boldsymbol{e}' \rangle.$$

Define the event $E = [E_0 \wedge E_1]$ as

$$E_0 \stackrel{\text{def}}{=} \left[\langle \boldsymbol{z}_0, \boldsymbol{e} \rangle = \langle \boldsymbol{z}_0, \boldsymbol{e}' \rangle\right],$$
$$E_1 \stackrel{\text{def}}{=} \left[\langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e} \rangle = \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e}' \rangle\right].$$

If event E occurs, then the reduction $\mathcal{B}$ perfectly simulates a pseudorandom instance of knapsack-eLWE$^+$ to $\mathcal{A}$, as then $\boldsymbol{v}_0 \cdot \langle \boldsymbol{z}_0, \boldsymbol{e} - \boldsymbol{e}' \rangle + \boldsymbol{v}_1 \cdot \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e} - \boldsymbol{e}' \rangle$ vanishes, leaving $\boldsymbol{c}' = \mathbf{H}'\boldsymbol{e}$ for $\mathbf{H}' \leftarrow \mathbb{Z}_q^{(m-n) \times m}$ and $\boldsymbol{e} \leftarrow \chi^m$ as required. Otherwise since q is prime, the reduction $\mathcal{B}$ (incorrectly) simulates an independent and uniform instance of knapsack-eLWE$^+$ to $\mathcal{A}$, as then either one of $\boldsymbol{v}_0 \cdot \langle \boldsymbol{z}_0, \boldsymbol{e} - \boldsymbol{e}' \rangle$ or $\boldsymbol{v}_1 \cdot \langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e} - \boldsymbol{e}' \rangle$ does not vanish, implying that $\boldsymbol{c}'$ is uniform in $\mathbb{Z}_q^{m-n}$ over the choice of $\boldsymbol{v}_0$ (resp. $\boldsymbol{v}_1$) alone, independent of the choices of $\mathbf{H}'$ and $\boldsymbol{x}$.

It remains to analyze the probability that event E occurs. Because $\boldsymbol{e}$ and $\boldsymbol{e}'$ are i.i.d., we may define the random variable $\mathcal{Z}_0$ that takes values $\langle \boldsymbol{z}_0, \boldsymbol{e}^* \rangle \in \mathbb{Z}_q$ and the random variable $\mathcal{Z}_1$ that takes values $\langle \mathbf{R}\boldsymbol{z}_1, \boldsymbol{e}^* \rangle \in \mathbb{Z}_q$ jointly over choice of $\boldsymbol{e}^* \leftarrow \chi^m$, and analyze their collision probabilities independently. Since the collision probability of *any* random variable $\mathcal{Z}$ is at least $1/|\mathsf{Supp}(\mathcal{Z})|$, we have that $\Pr[E] \geq \min CP[\mathcal{Z}_0] \cdot \min CP[\mathcal{Z}_1] = 1/q^2 = 1/\mathsf{poly}(n)$, and the theorem follows. $\qquad\square$

4   Flexibly Bi-Deniable Attribute-Based Encryption (ABE) for Branching Programs

In this section, we present our flexibly bi-deniable ABE for bounded-length Branching Program. We organize our approach into the following three steps:

(1) first, we recall the encoding scheme proposed in the SIM-secure ABE-BP of [30]; **(2)** Then, we present our flexibly bi-deniable attribute bi-translucent set (AB-BTS) scheme, as was defined in Definition 3.2. Our AB-BTS construction uses the ideas of Gorbunov and Vinayagamurthy [30], with essential modifications that allow us to tightly upper and lower bound evaluated noise terms. As discussed in the Introduction, this tighter analysis plays a key role in proving bi-deniability. **(3)** Finally, we show how to obtain the desired bi-deniable ABE scheme from our AB-BTS. As pointed out by Canetti et al. [14] and O'Neill et al. [35], a bitranslucent set scheme implies flexibly bi-deniable PKE. We observe that the same idea generalizes to the case of an AB-BTS scheme and flexibly bi-deniable ABE in a straightforward manner.

4.1   Encoding Schemes for Branching Programs

Basic Homomorphic Encoding. Before proceeding to the public key evaluation algorithm, we first described basic homomorphic addition and multiplication over public keys and encoded ciphertexts based on the techniques in [3,9,25].

Definition 4.1 (LWE Encoding). *For any matrix $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, we define an LWE encoding of a bit $a \in \{0, 1\}$ with respect to a public key $\mathbf{A}$ and randomness $\boldsymbol{s} \leftarrow \mathbb{Z}_q^n$ as*

$$\psi_{\mathbf{A}, \boldsymbol{s}, a} = \boldsymbol{s}^T(\mathbf{A} + a \cdot \mathbf{G}) + \boldsymbol{e} \in \mathbb{Z}_q^m$$

for error vector $\boldsymbol{e} \leftarrow \chi^m$ and the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$.

In our construction, all LWE encodings will be encoded using the same LWE secret $\boldsymbol{s}$, thus for simplicity, we will simply refer to such an encoding as $\psi_{\mathbf{A}, a}$.

For homomorphic addition, the addition algorithm takes as input two encodings $\psi_{\mathbf{A}, a}, \psi_{\mathbf{A}', a'}$, and outputs the sum of them. Let $\mathbf{A}^+ = \mathbf{A} + \mathbf{A}'$ and $a^+ = a + a'$

$$\mathsf{Add}(\psi_{\mathbf{A}, a}, \psi_{\mathbf{A}', a'}) = \psi_{\mathbf{A}, a} + \psi_{\mathbf{A}', a'} = \psi_{\mathbf{A}^+, a^+}$$

For homomorphic multiplication, the multiplication algorithm takes as input two encodings $\psi_{\mathbf{A}, a}, \psi_{\mathbf{A}', a}$, and outputs an encoding $\psi_{\mathbf{A}^\times, a^\times}$, where $\mathbf{A}^\times = -\mathbf{A}\mathbf{G}^{-1}(\mathbf{A}')$ and $a^\times = aa'$.

$$\mathsf{Mult}(\psi_{\mathbf{A}, a}, \psi_{\mathbf{A}', a'}) = -\psi \cdot \mathbf{G}^{-1}(\mathbf{A}') + a \cdot \psi' = \psi_{\mathbf{A}^\times, a^\times}$$

Public Key Evaluation Algorithm. Following the notation in [30], we define a public evaluation algorithm $\mathsf{Eval}_{\mathsf{pk}}$. The algorithm takes as input a description of the branching program BP, a collection of public keys $\{\mathbf{A}_i\}_{i \in [\ell]}$ (one for each attribute bit x_i), a collection of public keys $\mathbf{V}_{0,i}$ for initial state vector and an auxiliary matrix $\mathbf{A}^c$, and outputs an evaluated public key corresponding to the branching program BP.

$$\mathbf{V}_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{pk}}(\mathsf{BP}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \mathbf{A}^c)$$

where the auxiliary matrix $\mathbf{A}^c$ are used to encoded constant 1 for each input wire. We also define matrix $\mathbf{A}'_i = \mathbf{A}^c - \mathbf{A}_i$ as a public key used to encode

$1 - x_i$. By the definition of branching programs, the output $\mathbf{V}_{\mathsf{BP}} \in \mathbb{Z}_q^{n \times m}$ is the homomorphically generated public key $\mathbf{V}_{L,1}$ at position 1 of the state vector for the L-th step of the branching program evaluation.

Recall that in the definition of branching programs, BP is represented by the tuple $\{\mathsf{var}(t), \{\gamma_{t,i,0}, \gamma_{t,i,1}\}_{i \in [5]}\}$ for $t \in [L]$, and the initial state vector is set to be $\boldsymbol{v}_0 = (1,0,0,0,0)$. Further, for $t \in [L]$, the computation is performed as $\boldsymbol{v}_t[i] = \boldsymbol{v}_{t-1}[\gamma_{t,i,0}](1 - x_{\mathsf{var}(t)}) + \boldsymbol{v}_{t-1}[\gamma_{t,i,1}] \cdot x_{\mathsf{var}(t)}$. It is important for the security proof (among other reasons) that the evaluated state vector in each step is independent of the attribute vector.

Encoding Evaluation Algorithm. We define an encoding evaluation algorithm $\mathsf{Eval}_{\mathsf{ct}}$ that takes as input the description of a branching program BP, an attribute vector $\boldsymbol{x}$, a set of encodings for the attribute $\{\mathbf{A}_i, \psi_i := \psi_{\mathbf{A}_i, x_i}\}_{i \in [\ell]}$, encodings of the initial state vector $\{\mathbf{V}_{0,i}, \psi_{0,i} := \psi_{\mathbf{V}_{0,i} \boldsymbol{v}_0[i]}\}_{i \in [5]}$ and an encoding of a constant 1, i.e. $\psi^c := \psi_{\mathbf{A}^c, 1}$. The algorithm $\mathsf{Eval}_{\mathsf{ct}}$ outputs an encoding of the result $y := \mathsf{BP}(\boldsymbol{x})$ with respect to the homomorphically derived public key $\mathbf{V}_{\mathsf{BP}} := \mathbf{V}_{L,1}$

$$\psi_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{ct}}(\mathsf{BP}, \boldsymbol{x}, \{\mathbf{A}_i, \psi_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}, \psi_{0,i}\}_{i \in [5]}, \{\mathbf{A}^c, \psi^c\})$$

As mentioned above, in branching program computation, for $t \in [L]$, we have for all $i \in [5]$

$$\boldsymbol{v}_t[i] = \boldsymbol{v}_{t-1}[\gamma_{t,i,0}](1 - x_{\mathsf{var}(t)}) + \boldsymbol{v}_{t-1}[\gamma_{t,i,1}] \cdot x_{\mathsf{var}(t)}$$

The evaluation algorithm proceeds inductively to update the encoding of the state vector for each step of the branching program. Next, we need to instantiate this inductive computation using the homomorphic operations described above, i.e. $\mathsf{Add}, \mathsf{Mult}$. Following the notation used in [30], we define $\psi_i' := \psi_{\mathbf{A}_i', (1-x_i)} = \boldsymbol{s}^T(\mathbf{A}_i' + (1 - x_i)\mathbf{G}) + \boldsymbol{e}_i'$, where $\mathbf{A}_i' = \mathbf{A}^c - \mathbf{A}_i$, to denote the encoding of $1 - x_i$. This encoding can be computed using $\mathsf{Add}(\psi_{\mathbf{A}^c, 1}, -\psi_{\mathbf{A}_i, x_i})$. Then assuming at time $t - 1 \in [L]$ we hold encodings of the state vector $\{\psi_{\mathbf{V}_{t-1,i}, \boldsymbol{v}_{t-1}[i]}\}_{i \in [5]}$. For $i \in [5]$, we compute the encodings of new state values as

$$\psi_{i,t} = \mathsf{Add}(\mathsf{Mult}(\psi_{\mathsf{var}(t)}', \psi_{t-1,\gamma_0}), \mathsf{Mult}(\psi_{\mathsf{var}(t)}, \psi_{t-1,\gamma_1}))$$

where $\gamma_0 := \gamma_{t,i,0}$ and $\gamma_1 := \gamma_{t,i,1}$. We omit the correctness proof of the encoding here, which is presented in [30].

The above algorithms suffice for us to describe our construction. To analyze the scheme, we need to encode a simulated public key and evaluate the simulated key. Due to lack of space, we present details on these simulated encodings and evaluations, plus some useful lemmas to bound the norm of errors in the simulation, in full version. As mentioned in the Introduction, these new bounds play a critical, technical role in our formal proof of flexible bi-deniability.

4.2 Construction of Flexibly Bi-Deniable ABE for Branching Programs

In this part, we present our flexibly bi-deniable AB-BTS scheme for bounded-length Branching Programs. We use a semantically-secure public key encryption $\Pi = (\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}')$ with message space $\mathcal{M}_\Pi = \mathbb{Z}_q^{m \times m}$ and ciphertext space $\mathcal{C}_\Pi$. For a family of branching programs of length bounded by L and input space $\{0, 1\}^\ell$, the description of $\mathsf{BiDenAB\text{-}BTS} = (\mathsf{Setup}, \mathsf{DenSetup}, \mathsf{Keygen}, \mathsf{SampleP}, \mathsf{SampleU}, \mathsf{TestP}, \mathsf{FakeRCoins}, \mathsf{FakeSCoins})$ are as follows:

- $\mathsf{Setup}(1^\lambda, 1^L, 1^\ell)$: On input the security parameter λ, the length of the branching program L and length of the attribute vector ℓ,
  1. Set the LWE dimension be $n = n(\lambda)$, modulus $q = q(n, L)$. Choose Gaussian distribution parameter $s = s(n)$. Let $\mathsf{params} = (n, q, m, s)$.
  2. Sample one random matrix associated with its trapdoor as
  $$(\mathbf{A}, \mathbf{T_A}) \leftarrow \mathsf{TrapGen}(q, n, m)$$
  3. Choose $\ell + 6$ random matrices $\{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \mathbf{A}^c$ from $\mathbb{Z}_q^{n \times m}$.
  4. Choose a random vector $\boldsymbol{u} \in \mathbb{Z}_q^n$.
  5. Compute a public/secret key pair $(\mathsf{pk}', \mathsf{sk}')$ for a semantically secure public key encryption $(\mathsf{pk}', \mathsf{sk}') \leftarrow \mathsf{Gen}'(1^\lambda)$
  6. Output the public parameter pp and master secret key msk as
  $$\mathsf{pp} = (\mathsf{params}, \mathbf{A}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \mathbf{A}^c, \boldsymbol{u}, \mathsf{pk}'), \quad \mathsf{msk} = (\mathbf{T_A}, \mathsf{sk}')$$

- $\mathsf{DenSetup}(1^\lambda, 1^L, 1^\ell)$: On input the security parameter λ, the length of branching program L and length of attribute vector ℓ, the deniable setup algorithm runs the same computation as setup algorithm, and outputs
  $$\mathsf{pp} = (\mathsf{params}, \mathbf{A}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \mathbf{A}^c, \boldsymbol{u}, \mathsf{pk}'), \quad \mathsf{msk} = (\mathbf{T_A}, \mathsf{sk}') \quad \mathsf{fk} = (\mathbf{T_A}, \mathsf{sk}')$$

- $\mathsf{Keygen}(\mathsf{msk}, \mathsf{BP})$: On input the master secret key msk and the description of a branching program BP, $\mathsf{BP} = (\boldsymbol{v}_0, \{\mathsf{var}(t), \{\gamma_{t,i,0}, \gamma_{t,i,1}\}_{i \in [5]}\}_{t \in [L]})$.
  1. Homomorphically compute a public matrix with respect to the branching program BP: $\mathbf{V}_\mathsf{BP} \leftarrow \mathsf{Eval}_\mathsf{pk}(\mathsf{BP}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \mathbf{A}^c)$.
  2. Sample a low norm vector $\boldsymbol{r}_\mathsf{BP} \in \mathbb{Z}_q^{2m}$, using
  $$\boldsymbol{r}_\mathsf{BP} \leftarrow \mathsf{SampleLeft}(\mathbf{A}, \mathbf{T_A}, (\mathbf{V}_\mathsf{BP} + \mathbf{G}), \boldsymbol{u}, sq)$$
  such that $\boldsymbol{r}_\mathsf{BP}^T \cdot [\mathbf{A} | \mathbf{V}_\mathsf{BP} + \mathbf{G}] = \boldsymbol{u}$.
  3. Output the secret key sk_BP for branching program as $\mathsf{sk}_\mathsf{BP} = (\boldsymbol{r}_\mathsf{BP}, \mathsf{BP})$.
- $\mathsf{SampleP}(\mathsf{pp}, \boldsymbol{x})$: On input public parameters pp and attribute $\boldsymbol{x}$,
  1. Choose an LWE secret $\boldsymbol{s} \in \mathbb{Z}_q^n$ uniformly at random.
  2. Choose noise vector $\boldsymbol{e} \leftarrow \mathcal{D}_{\mathbb{Z}_q^m, \alpha}$, and compute $\psi_0 = \boldsymbol{s}^T \mathbf{A} + \boldsymbol{e}$.
  3. Choose one random matrices $\mathbf{R}^c \leftarrow \{-1, 1\}^{m \times m}$, and let $\boldsymbol{e}^c = \boldsymbol{e}^T \mathbf{R}^c$. Compute an encoding of constant 1: $\psi^c = \boldsymbol{s}^T(\mathbf{A}^c + \mathbf{G}) + \boldsymbol{e}^c$.
  4. Encode each bit $i \in [\ell]$ of the attribute vector:

(a) Choose a random matrix $\mathbf{R}_i \leftarrow \{-1, 1\}^{m \times m}$, and let $\boldsymbol{e}_i = \boldsymbol{e}^T \mathbf{R}_i$.

(b) Compute $\psi_i = \boldsymbol{s}^T (\mathbf{A}_i + x_i \mathbf{G}) + \boldsymbol{e}_i$.

5. Encode the initial state vector $\boldsymbol{v}_0 = (1, 0, 0, 0, 0)$, for $i \in [5]$

   (a) Choose a random matrix $\mathbf{R}'_{0,i} \leftarrow \{-1, 1\}^{m \times m}$, and let $\mathbf{R}_{0,i} = \eta \mathbf{R}'_{0,i}, \boldsymbol{e}_{0,i} = \boldsymbol{e}^T \mathbf{R}_{0,i}$, where the noise scaling parameter η is set in Sect. 4.3.

   (b) Compute $\psi_{0,i} = \boldsymbol{s}^T (\mathbf{V}_{0,i} + \boldsymbol{v}_0[i]\mathbf{G}) + \boldsymbol{e}_{0,i}$.

6. Compute $c = \boldsymbol{s}^T \boldsymbol{u} + e$, where $e \leftarrow \mathcal{D}_{\mathbb{Z}_q, s}$

7. Use PKE to encrypt randomly chosen matrices $\mathbf{R}^c, \{\mathbf{R}_i\}_{i \in [\ell]}$ and $\{\mathbf{R}_{0,i}\}_{i \in [5]}$:

$$\mathbf{T}_i \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{R}_i), \mathbf{T}^c \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{R}^c), \mathbf{T}_{0,i} \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{R}_{0,i})$$

8. Output the ciphertext

$$\mathsf{ct}_{\boldsymbol{x}} = (\boldsymbol{x}, \psi_0, \{\psi_i\}_{i \in [\ell]}, \psi^c, \{\psi_{0,i}\}_{i \in [5]}, c, \{\mathbf{T}_i\}_{i \in [\ell]}, \mathbf{T}^c, \{\mathbf{T}_{0,i}\}_{i \in [5]})$$

– $\mathsf{SampleU}(\mathsf{pp}, \boldsymbol{x})$: Output a uniformly random vector $\mathsf{ct} \in \mathbb{Z}_q^m \times \mathbb{Z}_q^{\ell m} \times \mathbb{Z}_q^{\ell m} \times \mathbb{Z}_q^{5m} \times \mathbb{Z}_q \times \mathcal{C}_\Pi^\ell \times \mathcal{C}_\Pi \times \mathcal{C}_\Pi^5$.

– $\mathsf{TestP}(\mathsf{sk}_{\mathsf{BP}}, \mathsf{ct}_{\boldsymbol{x}})$: On input the secret key $\mathsf{sk}_{\mathsf{BP}}$ for a branching program BP and a ciphertext associated with attribute $\boldsymbol{x}$, if $\mathsf{BP}(\boldsymbol{x}) = 0$, output $\bot$, otherwise,

1. Homomorphically compute the evaluated ciphertext of result $\mathsf{BP}(\boldsymbol{x})$

$$\psi_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{ct}}(\mathsf{BP}, \boldsymbol{x}, \{\mathbf{A}_i, \psi_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}, \psi_{0,i}\}_{i \in [5]}, \{\mathbf{A}_i^c, \psi_i^c\}_{i \in [\ell]})$$

2. Then compute $\phi = [\psi_0 | \psi_{\mathsf{BP}}]^T \cdot \boldsymbol{r}_{\mathsf{BP}}$. Accept $\mathsf{ct}_{\boldsymbol{x}}$ as a P-sample if $|c - \phi| < 1/4$, otherwise reject.

– $\mathsf{FakeSCoins}(r_S)$: Simply output the P-sample $\mathbf{c}$ as the randomness r_S^* that would cause $\mathsf{SampleU}$ to output $\boldsymbol{c}_{\boldsymbol{x}}$.

– $\mathsf{FakeRCoins}(\mathsf{pp}, \mathsf{fk}, \mathsf{ct}_{\boldsymbol{x}}, \mathsf{BP})$: On input the public parameters pp, the faking key fk, a ciphertext $\mathsf{ct}_{\boldsymbol{x}}$ and description of a branching program BP

1. If $\mathsf{BP}(\boldsymbol{x}) \neq 0$, then output $\mathsf{sk}_f \leftarrow \mathsf{Keygen}(\mathsf{fk}, \mathsf{BP})$.

2. Otherwise, parse ciphertext $\mathsf{ct}_{\boldsymbol{x}}$ as

$$\mathsf{ct}_{\boldsymbol{x}} = (\boldsymbol{x}, \psi_0, \{\psi_i\}_{i \in [\ell]}, \psi^c, \{\psi_{0,i}\}_{i \in [5]}, c, \{\mathbf{T}_i\}_{i \in [\ell]}, \mathbf{T}^c, \{\mathbf{T}_{0,i}\}_{i \in [5]})$$

Compute $\boldsymbol{e} \leftarrow \mathsf{Invert}(\mathbf{A}, \mathbf{T_A}, \psi_0)$. Then decrypt $(\{\mathbf{T}_i\}_{i \in [\ell]}, \mathbf{T}^c, \{\mathbf{T}_{0,i}\}_{i \in [5]})$ respectively using $\mathsf{Dec}(\mathsf{sk}', \cdot)$ to obtain $\{\mathbf{R}_i\}_{i \in [\ell]}, \mathbf{R}^c, \{\mathbf{R}_{0,i}\}_{i \in [5]}$. Compute evaluated error

$$\boldsymbol{e}_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{ct}}(\mathsf{BP}, \boldsymbol{x}, \{\mathbf{A}_i, \boldsymbol{e}^T \mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}, \boldsymbol{e}^T \mathbf{R}_{0,i}\}_{i \in [5]}, \{\mathbf{A}^c, \boldsymbol{e}^T \mathbf{R}^c\})$$

such that $\boldsymbol{e}_{\mathsf{BP}} = \boldsymbol{e}^T \mathbf{R}_{\mathsf{BP}}$.

3. Homomorphically compute a public matrix with respect to the branching program BP: $\mathbf{V}_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{pk}}(\mathsf{BP}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \{\mathbf{A}_i^c\}_{i \in [\ell]})$. Then sample a properly distributed secret key $\boldsymbol{r}_{\mathsf{BP}} \in \mathbb{Z}_q^{2m}$, using

$$\boldsymbol{r}_{\mathsf{BP}} \leftarrow \mathsf{SampleLeft}(\mathbf{A}, \mathbf{T_A}, (\mathbf{V}_{\mathsf{BP}} + \mathbf{G}), \boldsymbol{u}, s)$$

4. Sample correlation vector $\mathbf{y}_0 \leftarrow \mathcal{D}_{\mathbb{Z}_q^m, \beta^2 q^2 \mathbf{I}_{m \times m}}$. Then sample correlation coefficient $\mu \leftarrow \mathcal{D}_\gamma$, and set vector $\boldsymbol{y}_1 = (\mu \boldsymbol{e}_{\mathsf{BP}} + \mathcal{D}_{\mathbf{Z}^m, \mathbf{Q}})q$, where

$$\mathbf{Q} = \beta^2 \mathbf{I}_{m \times m} - \gamma^2 \alpha^2 \mathbf{R}_{\mathsf{BP}}^T \mathbf{R}_{\mathsf{BP}} \tag{1}$$

5. Let $\boldsymbol{y} = (\boldsymbol{y}_0 | \boldsymbol{y}_1)$, then sample and output the faked secret key $\mathsf{sk}_{\mathsf{BP}}^* = \mathbf{r}_{\mathsf{BP}}^*$ as $\mathbf{r}_{\mathsf{BP}}^* \leftarrow \boldsymbol{y} + \mathcal{D}_{\Lambda + \boldsymbol{r}_{\mathsf{BP}} - \boldsymbol{y}, \sqrt{s^2 - \beta^2}}$, using $\mathsf{SampleD}(\mathsf{ExtBasis}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{V}_{\mathsf{BP}} + \mathbf{G}), \boldsymbol{r}_{\mathsf{BP}} - \boldsymbol{y}, \sqrt{s^2 - \beta^2})$, where $\Lambda = \Lambda^\perp([\mathbf{A} | \mathbf{V}_{\mathsf{BP}} + \mathbf{G}])$.

The SampleP algorithm is similar to the ABE ciphertexts in the work [30], except that we add another scaling factor η to the rotation matrices $\mathbf{R}_{0,i}$'s. This allows us to both upper and lower bound the noise growth, which is essential to achieve bi-deniability. Detailed analysis can be found in full version. As we discussed in the introduction, the FakeRCoins embeds the evaluated noise into the secret key, so that it will change the decrypted value of the targeted ciphertext, but not others. Next we present the theorem we achieve and a high level ideas of the proof. We defer the formal analysis after the proof intuition.

Theorem 4.2. *Assuming the hardness of extended-$\mathsf{LWE}_{q,\beta'}$, the above algorithms form a secure attribute-based bitranslucent set schemem, as in Definition 3.2.*

Overview of Our Security Proof. At a high level, our security proof begins at the Fake experiment (cf. Definition 3.1 for a formal description), where first a ciphertext ct^* and its associated noise terms $\boldsymbol{e}^*$ are sampled, then a fake key $\boldsymbol{r}^*$ is generated that "artificially" fails to decrypt any ciphertext with noise vector (oriented close to) $\boldsymbol{e}^*$. In the end, we will arrive at the Real experiment, where an honest key $\boldsymbol{r}$ is generated that "genuinely" fails to decrypt the honestly generated, coerced ciphertext ct^*. (Multi-ct coercion security follows by a standard hybrid argument that repeatedly modifies respective $\boldsymbol{r}^*$ to $\boldsymbol{r}$ for each coerced ct^* in order.) In order to transition from Fake to Real, we move through a sequence of computationally- or statistically-indistinguishable hybrid experiments.

The first set of intermediate experiments (represented by H_1 and H_2 in our formal proof) embeds the attribute $\boldsymbol{x}$ of the challenge ciphertext ct^* in the public parameters, in a similar fashion to the beginning of every SIM-secure proof of lattice-based ABE. Indistinguishability follows via the Leftover Hash Lemma [21]. (Note that the additional hybrid in our proof is used to ensure that the random rotation matrices $\mathbf{R}$ employed by the LHL for public key embedding of $\boldsymbol{x}$ are the *exact same* matrices $\mathbf{R}$ as used to generate the noise terms of the coerced ct^*, and uses the security of any semantically-secure PKE for computational indistinguishability.)

The next set of intermediate experiments (given by $\mathsf{H}_3, \mathsf{H}_4$, and H_5 in our formal proof) perform the "main, new work" of our security proof. Specifically, they "swap the order" of the generation of the pk matrices $\{\mathbf{A}\}$, the public coset $\boldsymbol{u}$ (in the public parameters and in the coerced ciphertext), and the error vector(s) $\boldsymbol{e}$ in the coerced ciphertext components. (An additional hybrid is used to

toggle the order of a "correlation vector" $\boldsymbol{y}$ – a random, planted vector used to allow for a more modular analysis of these steps.) In each case, we give a *statistical* argument that the adversary's view in adjacent hybrids is indistinguishable or identical, using elementary properties of multi-dimensional Gaussians.

In the next step (given by H_6), we apply the eLWE^+ assumption to (roughly) change every component of the coerced ciphertext ct^* to uniform – except for the final c^* component used to blind the message μ.

In the final step (given by H_7), we transition to the Real experiment by changing the c^* component to uniform (in the presence of Dual Regev decryption under honest $\boldsymbol{z}$), using our sharper noise analysis as described above to show statistical indistinguishability of the final decryption output of $\boldsymbol{z}$ on ct^*.

Lemma 4.3. *For parameters set in Sect. 4.3, the AB-BTS defined above satisfies the correctness property in Definition 3.2.*

Proof. As we mentioned in Remark 3.3, the correctness of faking algorithms is implied by the bi-deniability property. Therefore, we only need to prove the correctness of normal decryption algorithm. For branching program BP and input $\boldsymbol{x}$, such that $\mathsf{BP}(\boldsymbol{x}) = 1$, we compute $\psi_{t,i}$ for $t \in [\ell]$ as

$$\psi_{t,i} = \mathsf{Add}(\mathsf{Mult}(\psi'_{\mathsf{var}(t)}, \psi_{t-1,\gamma_0}), \mathsf{Mult}(\psi_{\mathsf{var}(t)}, \psi_{t-1,\gamma_1}))$$

$$= \mathsf{Add}\Big([\boldsymbol{s}^T(-\mathbf{A}'_{\mathsf{var}(t)}\mathbf{G}^{-1}(\mathbf{V}_{t-1,\gamma_0}) + (\boldsymbol{v}_t[\gamma_0] \cdot (1 - x_{\mathsf{var}(t)})) \cdot \mathbf{G}) + \boldsymbol{e}_1],$$

$$([\boldsymbol{s}^T(-\mathbf{A}'_{\mathsf{var}(t)}\mathbf{G}^{-1}(\mathbf{V}_{t-1,\gamma_1}) + (\boldsymbol{v}_t[\gamma_1] \cdot x_{\mathsf{var}(t)}) \cdot \mathbf{G}) + \boldsymbol{e}_2]\Big)$$

$$= \boldsymbol{s}^T\Big[\underbrace{\big(-\mathbf{A}'_{\mathsf{var}(t)}\mathbf{G}^{-1}(\mathbf{V}_{t-1,\gamma_0}) - \mathbf{A}'_{\mathsf{var}(t)}\mathbf{G}^{-1}(\mathbf{V}_{t-1,\gamma_1})\big)}_{\mathbf{V}_{t,i}}$$

$$+ \underbrace{\big(\boldsymbol{v}_t[\gamma_0] \cdot (1 - x_{\mathsf{var}(t)}) + \boldsymbol{v}_t[\gamma_1] \cdot x_{\mathsf{var}(t)}\big)}_{\boldsymbol{v}_t[i]} \cdot \mathbf{G}\Big] + \boldsymbol{e}_{t,i}$$

At the end of the ciphertext evaluation, since $\mathsf{BP}(\boldsymbol{x}) = 1$, we can obtain $\psi_{\mathsf{BP}} = \boldsymbol{s}^T(\mathbf{V}_{\mathsf{BP}} + \mathbf{G}) + \boldsymbol{e}_{\mathsf{BP}}$, where $\boldsymbol{e}_{\mathsf{BP}} = \boldsymbol{e}^T\mathbf{R}_{\mathsf{BP}}$. Recall that the secret key $\mathsf{sk} = \boldsymbol{r}_{\mathsf{BP}}$ satisfying $[\mathbf{A}|\mathbf{V}_{\mathsf{BP}} + \mathbf{G}] \cdot \boldsymbol{r}_{\mathsf{BP}} = \boldsymbol{u}$. Then for $c - [\psi_0|\psi_{\mathsf{BP}}] \cdot \boldsymbol{r}_{\mathsf{BP}}$, it holds that

$$c - [\psi_0|\psi_{\mathsf{BP}}]^T \cdot \boldsymbol{r}_{\mathsf{BP}} = e - \boldsymbol{e}^T\mathbf{R}_{\mathsf{BP}} \cdot \boldsymbol{r}_{\mathsf{BP}}$$

Now we need to compute a bound for the final noise term. By applying analysis of norm in full version, we obtain that

$$||\boldsymbol{e}^T|| \cdot ||\mathbf{R}_{\mathsf{BP}}|| + 2m^{1.5}\ell||e|| \leq (2m^{1.5} + \eta\sqrt{m})||e|| \leq \alpha\sqrt{m}(2m^{1.5}\ell + \eta\sqrt{m}) \cdot sq\sqrt{m} \leq \frac{1}{4}$$

So by setting the parameters appropriately, as in Sect. 4.3, we have that

$$|c - [\psi_0|\psi_{\mathsf{BP}}]^T \cdot \boldsymbol{r}_{\mathsf{BP}}| \leq 1/4$$

and the lemma follows. $\square$

Lemma 4.4. *Assuming the hardness of extended-*$\mathsf{LWE}_{q,\beta'}$*, the AB-BTS scheme described above is bi-deniable as defined in Definition 3.2.*

Proof. First, we notice that because $\mathsf{SampleU}$ simply outputs its random coins as a uniformly random ct, we can use ct itself as the coins.

We prove the bi-deniability property by a sequence of hybrids H_i with details as follows:

Hybrid H_0: Hybrid H_0 is the same as the view of adversary $\mathcal{A}$ in the right-hand faking experiment in the definition of bi-deniability. We use the fact that algorithm Invert successfully recovers e from ct with overwhelming probability over all randomness in the experiment.

Hybrid H_1: In hybrid H_2, we switch the encryptions of matrices $(\{\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{R}_{0,i}\}_{i\in[5]}, \mathbf{R}^c)$ in the ciphertext to encryptions of zero.
Recall that in hybrid H_0, we encrypt the randomness matrices $(\{\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{R}_{0,i}\}_{i\in[5]}, \mathbf{R}^c)$ using semantically secure PKE Π, i.e.

$$\mathbf{T}_i \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{R}_i), \quad \mathbf{T}^c \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{R}^c), \quad \mathbf{T}_{0,i} \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{R}_{0,i})$$

In hybrid H_1, we just set

$$\mathbf{T}_i \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{0}), \quad \mathbf{T}^c \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{0}), \quad \mathbf{T}_{0,i} \leftarrow \mathsf{Enc}'(\mathsf{pk}', \mathbf{0})$$

to be encryptions of $\mathbf{0} \in \mathbb{Z}^{m\times m}$ to replace encryptions of matrices $(\{\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{R}_{0,i}\}_{i\in[5]}, \mathbf{R}^c)$.

Hybrid H_2: In hybrid H_2, we embed random matrices $(\{\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{R}_{0,i}\}_{i\in[5]}, \mathbf{R}^c)$ and challenge attribute x^* in the public parameters pp.
Recall that in hybrid H_1 the matrices $(\{\mathbf{A}_i\}_{i\in[\ell]}, \{\mathbf{V}_{0,i}\}_{i\in[5]}, \mathbf{A}^c)$ are sampled at random. In hybrid H_2, we slightly change how these matrices are generated. Let $x^* = (x_1^*, ..., x_\ell^*)$ be the challenge attribute that the adversary $\mathcal{A}$ intends to attack. We sample matrices $(\{\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{R}'_{0,i}\}_{i\in[5]}, \mathbf{R}^c)$ uniformly random from $\{-1,1\}^{m\times m}$ and set $\mathbf{R}_{0,i} = \eta\mathbf{R}'_{0,i}$, which would be used both in the generation of public parameters and challenge ciphertext. We set $(\{\mathbf{A}_i\}_{i\in[\ell]}, \{\mathbf{V}_{0,i}\}_{i\in[5]}, \mathbf{A}^c)$ respectively as

$$\mathbf{A}_i = \mathbf{A}\mathbf{R}_i - x_i^*\mathbf{G}, \quad \mathbf{V}_{0,i} = \mathbf{A}\mathbf{R}_{0,i} - v_0[i]\mathbf{G}, \quad \mathbf{A}^c = \mathbf{A}\mathbf{R}^c - \mathbf{G}$$

where $v_0 = [1, 0, 0, 0, 0]$. The rest of the hybrid remains unchanged.

Hybrid H_3: In hybrid H_3, we change the generation of matrix $\mathbf{A}$ and vector u in public parameters pp.
Let $\mathbf{A}$ be a random matrix in $\mathbb{Z}_q^{n\times m}$. The construction of matrices $(\{\mathbf{A}_i\}_{i\in[\ell]}, \{\mathbf{V}_{0,i}\}_{i\in[5]}, \mathbf{A}^c)$ remains the same, as in hybrid H_2. Sample error vectors e that would be used in algorithm $\mathsf{SampleP}$ later. Then compute the error vector

$$e_{\mathsf{BP}^*} \leftarrow \mathsf{Eval}_{\mathsf{ct}}(\mathsf{BP}, x, \{\mathbf{A}_i, e^T\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{V}_{0,i}, e^T\mathbf{R}_{0,i}\}_{i\in[5]}, \{\mathbf{A}^c, e^T\mathbf{R}^c\})$$

and choose a correlation coefficient $\mu \leftarrow \mathcal{D}_\gamma$, and set vector $y_1 = (\mu e_{\mathsf{BP}^*} + \mathcal{D}_{\mathbb{Z}^m,\mathbf{Q}})q$, where

$$\mathbf{Q} = \beta^2\mathbf{I}_{m\times m} - \gamma^2\alpha^2\mathbf{R}_{\mathsf{BP}^*}^T\mathbf{R}_{\mathsf{BP}^*}$$

Then let $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1)$, where $\boldsymbol{y}_0 \leftarrow \mathcal{D}_{\mathbf{Z}_q^m, \beta^2 q^2 \mathbf{I}_{m \times m}}$. Sample vector $\mathbf{r}_{\mathsf{BP}^*} \leftarrow \boldsymbol{y} + \mathcal{D}_{\mathbb{Z}^{2m} - \boldsymbol{y}, (s^2 - \beta^2) q^2 \mathbf{I}_{2m \times 2m}}$, and compute matrix

$$\mathbf{V}_{\mathsf{BP}^*} \leftarrow \mathsf{Eval}_{\mathsf{pk}}(\mathsf{BP}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \mathbf{A}^c)$$

Set vector $\boldsymbol{u}$ in public parameters pp as $\boldsymbol{u} = [\mathbf{A}|\mathbf{V}_{\mathsf{BP}^*}] \cdot \mathbf{r}_{\mathsf{BP}^*}$. Since $\mathbf{A}$ is a random matrix without trapdoor $\mathbf{T_A}$ to answer key queries, we will use trapdoor $\mathbf{T_G}$ to answer queries as follows. Consider a secret key query for branching program BP such that $\mathsf{BP}(\boldsymbol{x}^*) = 0$. To respond, we do the following computations:

1. First, we compute

$$\mathbf{R}_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{Sim}}(\mathsf{BP}, \boldsymbol{x}, \{\mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{R}_{0,i}\}_{i \in [5]}, \mathbf{R}^c, \mathbf{A})$$

   to obtain a low-norm matrix $\mathbf{R}_{\mathsf{BP}} \in \mathbb{Z}_q^{m \times m}$ satisfying $\mathbf{A}\mathbf{R}_{\mathsf{BP}} - \mathsf{BP}(\boldsymbol{x}^*)\mathbf{G} = \mathbf{V}_{\mathsf{BP}}$.

2. Then, we sample $\boldsymbol{r}_{\mathsf{BP}}$ using

$$\boldsymbol{r}_{\mathsf{BP}} \leftarrow \mathsf{SampleRight}(\mathbf{A}, \mathbf{G}, \mathbf{R}_{\mathsf{BP}}, \mathbf{T_G}, \boldsymbol{u}, sq)$$

   such that

$$\boldsymbol{r}_{\mathsf{BP}}^T \cdot [\mathbf{A}|\mathbf{V}_{\mathsf{BP}} + \mathbf{G}] = \boldsymbol{u}$$

   By property of algorithm $\mathsf{SampleRight}$, vector $\boldsymbol{r}_{\mathsf{BP}}$ is distributed as required.

The computation of answering P-sampler query, $\mathsf{SampleP}$ is the same as hybrid H_1 with error vectors $\boldsymbol{e}$, For faking receiver coins, $\mathsf{FakeRCoins}$, simply output the vector $\boldsymbol{r}_{\mathsf{BP}^*}$ pre-sampled in the generation of vector $\boldsymbol{u}$ before.

Hybrid H_4: In hybrid H_4, we change the generation order of vector $\boldsymbol{y}$ and error vector $\boldsymbol{e}$.

First sample vector $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1) \leftarrow \mathcal{D}_{\mathbb{Z}^{2m}, \beta^2 q^2 \mathbf{I}_{2m \times 2m}}$ and compute $\boldsymbol{r}_{\mathsf{BP}^*}$ from $\boldsymbol{y}$ as in previous hybrid. Next, we compute error term $\boldsymbol{e}$ as $\boldsymbol{e} = \nu \boldsymbol{y}_1^T \mathbf{R}_{\mathsf{BP}^*}/q + \mathcal{D}_{\mathbb{Z}^m, \mathbf{Q}'}$, where $\nu \leftarrow \mathcal{D}_\tau, \tau = \gamma \alpha^2/\beta^2$, and $\mathcal{D}_{\mathbb{Z}^m, \mathbf{Q}'}$ is sampled as $\mathbf{L}' \mathcal{D}_{\mathbf{Z}_1^m, \mathbf{I}_{m \times m}}$ for

$$\mathbf{Q}' = \mathbf{L}'\mathbf{L}'^T = \alpha^2 \mathbf{I} - \tau^2 \beta^2 \mathbf{R}_{\mathsf{BP}^*}^T \mathbf{R}_{\mathsf{BP}^*} \tag{2}$$

Additionally, we modify the challenge ciphertext to be

$$\psi_0^* = \boldsymbol{s}^T \mathbf{A}/q + \boldsymbol{e}, \quad \psi_i^* = \psi_0^{*T} \mathbf{R}_i/q, \quad \psi_{0,i}^* = \psi_0^{*T} \mathbf{R}_{0,i}/q, \quad \psi^{*c} = \psi_0^{*T} \mathbf{R}^c/q$$

and $c^* = \boldsymbol{s}^T \boldsymbol{u} + \mathcal{D}_{\mathbf{Z}^m, \alpha \mathbf{I}_{m \times m}}$.

Hybrid H_5: In hybrid H_5, we change the generation order of secret key $\boldsymbol{r}_{\mathsf{BP}^*}$ and vector $\boldsymbol{y}$.

We first sample matrix $\boldsymbol{r}_{\mathsf{BP}^*}$ from discrete Gaussian distribution $\mathcal{D}_{\mathbb{Z}^{2m}, s^2 q^2 \mathbf{I}_{2m \times 2m}}$, and set vector $\boldsymbol{u}$ in public parameters pp to be $\boldsymbol{u} = [\mathbf{A}|\mathbf{V}_{\mathsf{BP}^*}] \cdot \boldsymbol{r}_{\mathsf{BP}^*}$, where

$$\mathbf{V}_{\mathsf{BP}^*} \leftarrow \mathsf{Eval}_{\mathsf{pk}}(\mathsf{BP}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \{\mathbf{A}_i^c\}_{i \in [\ell]})$$

Then set $\boldsymbol{y} = (\boldsymbol{y}_0 | \boldsymbol{y}_1) = \boldsymbol{r}_{\mathsf{BP}^*}/2 + \mathcal{D}_{\mathbb{Z}^{2m}, (\beta^2 - s^2/4)q^2 \mathbf{I}_{2m \times 2m}}$. The remainder of the hybrid remains roughly the same. In particular, the challenge ciphertext ct^* is generated in the same manner as Hybrid H_4. We break the noise term $\boldsymbol{e}$ into two terms $\boldsymbol{e} = \boldsymbol{e}_0^{(1)} + \boldsymbol{e}_0^{(2)} + \nu \boldsymbol{y}_1^T \mathbf{R}_{\mathsf{BP}^*}/q$, where $\boldsymbol{e}_0^{(1)} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \beta' \mathbf{I}_{m \times m}}, \boldsymbol{e}_0^{(2)} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \mathbf{Q}' - \beta'^2 \mathbf{I}_{m \times m}}$ and $\beta' = \alpha/2$.

Hybrid H_6: In hybrid H_6, we change how the challenge ciphertext is generated by using the Extended-LWE$^+$ instance.

First sample uniformly random vector $\boldsymbol{b} \in \mathbb{Z}^m$ and set the challenge ciphertext as

$$\psi_0^* = \boldsymbol{b}/q + \boldsymbol{e}_0^{(2)}, \quad \psi_i^* = \psi_0^{*T} \mathbf{R}_i/q, \quad \psi_{0,i}^* = \psi_0^{*T} \mathbf{R}_{0,i}/q, \quad \psi^{*c} = \psi_0^{*T} \mathbf{R}^c/q$$

and $c^* = \boldsymbol{r}_{\mathsf{BP}^*}^T [\mathbf{I}_{m \times m} | \mathbf{R}_{\mathsf{BP}^*}] (\boldsymbol{b}/q - \boldsymbol{e}_0^{(1)}) + \mathcal{D}_{\mathbb{Z}^m, \alpha \mathbf{I}_{m \times m}}$.

Hybrid H_7: In hybrid H_7, we change the challenge ciphertext to be uniformly random.

In algorithm SampleP, sample uniformly random vectors $\mathsf{ct} \in \mathbb{Z}_q^m \times \mathbb{Z}_q^{\ell m} \times \mathbb{Z}_q^m \times \mathbb{Z}_q^{5m} \times \mathbb{Z}_q$ and outputs ct.

Claim 4.5. *Assuming the semantic security of PKE $\Pi = (\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}')$, hybrid H_0 and H_1 are computationally indistinguishable.*

Proof. Observe there is only one difference between hybrids H_0 and H_1 occurs in the challenge ciphertext, i.e. the encryption (under PKE Π) of the random matrices $\mathbf{S}_i$ are replaced by encryption of 0. If a PPT adversary $\mathcal{A}$ distinguishes between the H_0-encryptions of $(\{\mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{R}_{0,i}\}_{i \in [5]}, \{\mathbf{R}^c\})$ and the H_1-encryptions of $\boldsymbol{0}$ with non-negligible probability, then we can construct an efficient reduction $\mathcal{B}$ that uses $\mathcal{A}$ to break the semantic security of PKE Π with similar probability. $\square$

Claim 4.6. *Hybrids H_1 and H_2 are statistically indistinguishable.*

Proof. Observe the only difference between hybrids H_1 and H_2 is the generation of matrices

$$(\{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \{\mathbf{A}_i^c\}_{i \in [\ell]})$$

The random matrices $(\{\mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{R}_{0,i}\}_{i \in [5]}, \{\mathbf{R}_i^c\}_{i \in [\ell]})$ are used in the generation of public parameters pp:

$$\mathbf{A}_i = \mathbf{A}\mathbf{R}_i - x_i^* \mathbf{G}, \quad \mathbf{V}_{0,i} = \mathbf{A}\mathbf{R}_{0,i} - \boldsymbol{v}_0[i]\mathbf{G}, \quad \mathbf{A}^c = \mathbf{A}\mathbf{R}^c - \mathbf{G}$$

and the construction of errors in challenge ciphertext

$$\boldsymbol{e}_i = \boldsymbol{e}^T \mathbf{R}_i, \quad \boldsymbol{e}^c = \boldsymbol{e}^T \mathbf{R}^c, \quad \boldsymbol{e}_{0,i} = \boldsymbol{e}^T \mathbf{R}_{0,i}$$

Then by Leftover Hash Lemma 2.1, the following two distributions are statistically indistinguishable

$$(\mathbf{A}, \{\mathbf{A}\mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{A}\mathbf{R}_{0,i}\}_{i \in [5]}, \{\mathbf{A}\mathbf{R}^c\}, \widetilde{\boldsymbol{e}}) \approx (\mathbf{A}, \{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \{\mathbf{A}^c\}, \widetilde{\boldsymbol{e}})$$

where $\widetilde{\boldsymbol{e}} = (\{\boldsymbol{e}_i\}_{i \in [\ell]}, \{\boldsymbol{e}_{0,i}\}_{i \in [5]}, \{\boldsymbol{e}^c\})$. Hence, hybrid H_0 and H_1 are statistically indistinguishable. $\square$

Claim 4.7. *Hybrids* H_2 *and* H_3 *are statistically indistinguishable.*

Proof. Observe there are three differences between hybrid H_2 and H_3: The generation of matrix $\mathbf{A}$ and vector $\boldsymbol{u}$ in pp, challenge secret key $\mathsf{sk}_{\mathsf{BP}^*}$ and the computation methods to answer secret key queries. By the property of algorithm $\mathsf{TrapGen}(q, n, m)$, the distribution of matrix $\mathbf{A}$ in hybrid H_2 is statistically close to uniform distribution, from which matrix $\mathbf{A}$ in hybrid H_3 is sampled.

For secret key queries regarding branching program BP, in hybrid H_2, we sample vector $\boldsymbol{r}_{\mathsf{BP}}$, using

$$\boldsymbol{r}_{\mathsf{BP}} \leftarrow \mathsf{SampleLeft}(\mathbf{A}, \mathbf{T_A}, (\mathbf{V}_{\mathsf{BP}} + \mathbf{G}), \boldsymbol{u}, s)$$

While in hybrid H_3, we sample vector $\boldsymbol{r}_{\mathsf{BP}}$, using

$$\boldsymbol{r}_{\mathsf{BP}} \leftarrow \mathsf{SampleRight}(\mathbf{A}, \mathbf{G}, \mathbf{R}_{\mathsf{BP}}, \mathbf{T_G}, \boldsymbol{u}, sq)$$

By setting the parameters appropriately as specified in Sect. 4.3, and the properties of algorithms $\mathsf{SampleLeft}$ and $\mathsf{SampleRight}$, the answers to secret key queries are statistically close.

By Leftover Hash Lemma 2.1, the distribution $([\mathbf{A}|\mathbf{V}_{\mathsf{BP}^*}], [\mathbf{A}|\mathbf{V}_{\mathsf{BP}^*}]\cdot\boldsymbol{r}_{\mathsf{BP}^*})$ and $([\mathbf{A}|\mathbf{V}_{\mathsf{BP}^*}], \boldsymbol{u})$ are statistically close. Hence, hybrid H_2 and H_3 are statistically indistinguishable. □

Claim 4.8. *Hybrids* H_3 *and* H_4 *are statistically indistinguishable.*

Proof. The only difference between the two experiments is in the choice of $\boldsymbol{y}$ and $\boldsymbol{e}$, specifically, the choice of the $\boldsymbol{y}_1$ component of $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1)$. We will show that the joint distribution of $(\boldsymbol{e}, \boldsymbol{y}_1)$ is identically distributed in these two hybrids:

In hybrid H_3, $\boldsymbol{y}_1$ is set as $\boldsymbol{y}_1 = (\mu\boldsymbol{e}_{\mathsf{BP}^*} + \mathcal{D}_{\mathbb{Z}^m, \mathbf{Q}})q$, where $\mathbf{Q} = \beta^2 \mathbf{I}_{m \times m} - \gamma^2\alpha^2 \mathbf{R}_{\mathsf{BP}^*}^T \mathbf{R}_{\mathsf{BP}^*}$ with $\boldsymbol{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \alpha^2 \mathbf{I}_{m \times m}}$ and

$$\boldsymbol{e}_{\mathsf{BP}^*} \leftarrow \mathsf{Eval}_{\mathsf{ct}}(\mathsf{BP}, \boldsymbol{x}, \{\mathbf{A}_i, \boldsymbol{e}^T\mathbf{R}_i\}_{i\in[\ell]}, \{\mathbf{V}_{0,i}, \boldsymbol{e}^T\mathbf{R}_{0,i}\}_{i\in[5]}, \{\mathbf{A}^c, \boldsymbol{e}^T\mathbf{R}^c\})$$

Therefore, in hybrid H_3, we may write the joint distribution of $(\boldsymbol{e}, \boldsymbol{y}_1)$ as $\mathbf{T}_1 \cdot \mathcal{D}_{\mathbb{Z}^{2m}, \mathbf{I}_{2m \times 2m}}$, where $\mathbf{T}_1 \stackrel{\text{def}}{=} \begin{pmatrix} \alpha\mathbf{I}_{m \times m} & \mathbf{0}_{m \times m} \\ \gamma\alpha q\mathbf{R}_{\mathsf{BP}^*}^T & q\mathbf{L} \end{pmatrix}$ for $\mathbf{Q} = \mathbf{L}\mathbf{L}^T \in \mathbb{Z}^{m \times m}$ via the Cholesky decomposition Lemma in full version.

In hybrid H_4, vector $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1)$ is sampled as $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1) \leftarrow \mathcal{D}_{\mathbb{Z}^{2m}, \beta^2 q^2 \mathbf{I}_{2m \times 2m}}$. Then $\boldsymbol{e}$ is computed as $\boldsymbol{e} = \nu\boldsymbol{y}_1^T\mathbf{R}_{\mathsf{BP}^*}/q + \mathcal{D}_{\mathbb{Z}^m, \mathbf{Q}'}$, where $\nu \leftarrow \mathcal{D}_\tau, \tau = \gamma\alpha^2/\beta^2$, and $\mathbf{Q}' = \alpha^2\mathbf{I} - \tau^2\beta^2\mathbf{R}_{\mathsf{BP}^*}^T\mathbf{R}_{\mathsf{BP}^*}$. Then in hybrid H_4, we may write the joint distribution of $(\boldsymbol{e}, \boldsymbol{y}_1)$ as $\mathbf{T}_2 \cdot \mathcal{D}_{\mathbb{Z}^{2m}, \mathbf{I}_{2m \times 2m}}$, where $\mathbf{T}_2 \stackrel{\text{def}}{=} \begin{pmatrix} \mathbf{L}' & \tau\beta\mathbf{R}_{\mathsf{BP}^*} \\ \mathbf{0}_{m \times m} & \beta q\mathbf{I}_{m \times m} \end{pmatrix}$ for $\mathbf{Q}' = \mathbf{L}'\mathbf{L}'^T \in \mathbb{Z}^{m \times m}$ via the Cholesky decomposition Lemma in full version.

We claim equality of the following systems of equations:

$$\mathbf{T}_1\mathbf{T}_1^T = \begin{pmatrix} \alpha^2\mathbf{I}_{m \times m} & \gamma\alpha^2 q\mathbf{R}_{\mathsf{BP}^*} \\ \gamma\alpha^2 q\mathbf{R}_{\mathsf{BP}^*}^T & \gamma^2\alpha^2 q^2\mathbf{R}_{\mathsf{BP}^*}^T\mathbf{R}_{\mathsf{BP}^*} + q^2\mathbf{L}\mathbf{L}^T \end{pmatrix}$$
$$= \begin{pmatrix} \mathbf{L}'\mathbf{L}'^T + \tau^2\beta^2\mathbf{R}_{\mathsf{BP}^*}\mathbf{R}_{\mathsf{BP}^*}^T & \tau\beta^2 q\mathbf{R}_{\mathsf{BP}^*} \\ \tau\beta^2 q\mathbf{R}_{\mathsf{BP}^*}^T & \beta^2 q^2\mathbf{I}_{m \times m} \end{pmatrix} = \mathbf{T}_2\mathbf{T}_2^T.$$

This fact may be seen quadrant-wise by our choice of $\tau = \gamma\alpha^2/\beta^2$ and the settings of $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$ and $\mathbf{Q}' = \mathbf{L}'\mathbf{L}'^T$ in Eqs. (1) and (2). It then follows that $(\mathbf{T}_2^{-1}\mathbf{T}_1)(\mathbf{T}_2^{-1}\mathbf{T}_1)^T = \mathbf{I}_{2m \times 2m}$, implying $\mathbf{T}_1 = \mathbf{T}_2\mathbf{Q}^*$ for some orthogonal matrix $\mathbf{Q}^*$. Because the spherical Gaussian $\mathcal{D}_{\mathbb{Z}^{2m}, \mathbf{I}_{2m \times 2m}}$ is invariant under rigid transformations, we have $\mathbf{T}_1 \cdot \mathcal{D}_{\mathbb{Z}^{2m}, \mathbf{I}_{2m \times 2m}} = \mathbf{T}_2\mathbf{Q}^* \cdot \mathcal{D}_{\mathbb{Z}^{2m}, \mathbf{I}_{2m \times 2m}} = \mathbf{T}_2 \cdot \mathcal{D}_{\mathbb{Z}^{2m}, \mathbf{I}_{2m \times 2m}}$, and the claim follows. $\square$

Claim 4.9. *Hybrids* H_4 *and* H_5 *are statistically indistinguishable.*

Proof. Observe the main difference between hybrids H_4 and H_5 is the order of generation of vectors $\boldsymbol{y}$ and $\boldsymbol{r}_{\mathsf{BP}^*}$: In hybrid H_4, we first sample $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1) \leftarrow \mathcal{D}_{\mathbb{Z}^{2m}, \beta^2 q^2 \mathbf{I}_{2m \times 2m}}$ and set $\boldsymbol{r}_{\mathsf{BP}^*} \leftarrow \boldsymbol{y} + \mathcal{D}_{\mathbb{Z}^{2m} - \boldsymbol{y}, q^2(s^2 - \beta^2)\mathbf{I}_{2m \times 2m}}$, while in hybrid H_5, we first sample $\boldsymbol{r}_{\mathsf{BP}^*} \leftarrow \mathcal{D}_{\mathbb{Z}^{2m}, s^2 q^2 \mathbf{I}_{2m \times 2m}}$ and set $\boldsymbol{y} = (\boldsymbol{y}_0|\boldsymbol{y}_1) \leftarrow \boldsymbol{r}_{\mathsf{BP}^*}/2 + \mathcal{D}_{\mathbb{Z}^{2m}, (\beta^2 - s^2/4)q^2 \mathbf{I}_{2m \times 2m}}$. By setting parameters appropriately as in Sect. 4.3, these two distributions are statistically close. $\square$

Claim 4.10. *Assuming the hardness of extended-*$\mathsf{LWE}^+_{n,m,q,\mathcal{D}_{\mathbb{Z}^m,\beta'},\mathbf{R}}$ *for any adversarially chosen distribution over matrices* $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$*, then hybrids* H_5 *and* H_6 *are computationally indistinguishable.*

Proof. Suppose $\mathcal{A}$ has non-negligible advantage in distinguishing hybrid H_5 and H_6, then we use $\mathcal{A}$ to construct an extended-LWE^+ algorithm $\mathcal{B}$ as follows:

Invocation. $\mathcal{B}$ invokes adversary $\mathcal{A}$ to commit to a challenge attribute vector $\boldsymbol{x}^* = (x_1^*, ..., x_\ell^*)$ and challenge branching program BP^*. Then $\mathcal{B}$ generates $\mathbf{R}_{\mathsf{BP}^*}$ by first sampling $(\{\mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{R}_{0,i}\}_{i \in [5]}, \{\mathbf{R}^c\})$ as in the hybrid, and computes

$$\mathbf{R}_{\mathsf{BP}} \leftarrow \mathsf{Eval}_{\mathsf{Sim}}(\mathsf{BP}, \boldsymbol{x}, \{\mathbf{R}_i\}_{i \in [\ell]}, \{\mathbf{R}_{0,i}\}_{i \in [5]}, \{\mathbf{R}^c\}, \mathbf{A})$$

Then it receives an extended-LWE^+ instance for the matrix $\mathbf{R} = \mathbf{R}_{\mathsf{BP}^*}$ as follows:

$$\{\mathbf{A}, \boldsymbol{b} = \boldsymbol{s}^T\mathbf{A} + \boldsymbol{e}, \boldsymbol{z}_0, \boldsymbol{z}_1, \langle \boldsymbol{z}_0, \boldsymbol{b} - \boldsymbol{e} \rangle + e, \langle \boldsymbol{z}_1^T\mathbf{R}, \boldsymbol{b} - \boldsymbol{e} \rangle + e'\}$$

where $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{n \times m}$, $\boldsymbol{s} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$, $\boldsymbol{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$, $\boldsymbol{e}, \boldsymbol{z}_0, \boldsymbol{z}_1 \overset{\$}{\leftarrow} \chi^n$ and $e, e' \overset{\$}{\leftarrow} \chi$. Algorithm $\mathcal{B}$ aims to leverage adversary $\mathcal{A}$'s output to solve the extended-LWE^+ assumption.

Setup. $\mathcal{B}$ generates matrices $(\{\mathbf{A}_i\}_{i \in [\ell]}, \{\mathbf{V}_{0,i}\}_{i \in [5]}, \{\mathbf{A}^c\})$ as specified in hybrid H_1. Then, $\mathcal{B}$ sets challenge secret key $\mathsf{sk}_{\mathsf{BP}^*} = \boldsymbol{r}_{\mathsf{BP}^*} = (\boldsymbol{r}_0^*|\boldsymbol{r}_1^*) = (\boldsymbol{z}_0|\boldsymbol{z}_1)$ from extended-LWE^+ instance and computes vector $\boldsymbol{u}$ as in hybrid H_5.

Secret key queries. $\mathcal{B}$ answers adversary $\mathcal{A}$'s secret key queries as in hybrid H_2.

Challenge ciphertext. $\mathcal{B}$ answers adversary $\mathcal{A}$'s P-sample query by setting

$$\psi_0^* = \boldsymbol{b}/q + \boldsymbol{e}_0^{(2)} + \nu\boldsymbol{y}_1^T\mathbf{R}_{\mathsf{BP}^*}/q, \quad \psi_i^* = \psi_0^{*T}\mathbf{R}_i/q, \quad \psi_{0,i}^* = \psi_0^{*T}\mathbf{R}_{0,i}/q, \quad \psi^{*c} = \psi_0^{*T}\mathbf{R}^c/q$$

and $c^* = \boldsymbol{r}_{\mathsf{BP}^*}^T[\mathbf{I}_{m \times m}|\mathbf{R}_{\mathsf{BP}^*}](\boldsymbol{b}/q - \boldsymbol{e}^{(1)}) + \mathcal{D}_{\mathbb{Z}^m, \alpha\mathbf{I}_{m \times m}}$.

Faking receiver coin query. $\mathcal{B}$ answers adversary $\mathcal{A}$'s faking receiver coin query by outputting the extended-LWE instance's vector $\mathsf{sk}_{\mathsf{BP}^*} = r_{\mathsf{BP}^*}$.

Output. $\mathcal{B}$ outputs whatever $\mathcal{A}$ outputs.

We can rewrite the expression of $c^{*'}$ to be

$$
\begin{aligned}
c^{*'} &= ([\mathbf{A}^* | \mathbf{A}^* \mathbf{R}_{\mathsf{BP}^*}] (\begin{smallmatrix} z_0 \\ z_1 \end{smallmatrix}))^T s/q + \mathcal{D}_{\mathbb{Z}_1, \alpha} \\
&= ((z_0 | z_1) (\begin{smallmatrix} \mathbf{A}^{*T} \\ \mathbf{R}_{\mathsf{BP}^*}^T \mathbf{A}^{*T} \end{smallmatrix})) s/q + D_{\mathbb{Z}_1, \alpha} = z_0 \mathbf{A}^{*T} s/q + z_1 \mathbf{R}_{\mathsf{BP}^*}^T \mathbf{A}^{*T} s/q + \mathcal{D}_{\mathbb{Z}_1, \alpha} \\
&= \langle z_0, b/q - e^{(1)} \rangle + \langle z_1^T \mathbf{R}_{\mathsf{BP}^*}, b/q - e^{(1)} \rangle + \mathcal{D}_{\mathbb{Z}_1, \alpha}
\end{aligned}
$$

We can see that if the eLWE^+ instance's vector b is pseudorandom, then the distribution simulated by $\mathcal{B}$ is exactly the same as H_5. If b is truly random and independent, then the distribution simulated by $\mathcal{B}$ is exactly the same as H_6. Therefore, if $\mathcal{A}$ can distinguish H_5 from H_6 with non-negligible probability, then $\mathcal{B}$ can break the $\mathsf{eLWE}^+_{n,m,q,\mathcal{D}_{(\alpha/2)q},\alpha',\mathbf{S}_{f^*}}$ problem for some $\alpha' \geq 0$ with non-negligible probability. □

Claim 4.11. *Hybrids H_6 and H_7 are statistically indistinguishable.*

Proof. Recall the only difference between hybrids H_6 and H_7 is the generation of challenge ciphertext. In hybrid H_7, we observe if ψ_0^* is chosen from uniform distribution, then by Leftover Hash Lemma 2.1, it holds

$$
\psi_i^* = \psi_0^{*T} \mathbf{R}_i/q, \quad \psi_{0,i}^* = \psi_0^{*T} \mathbf{R}_{0,i}/q, \quad \psi^{*c} = \psi_0^{*T} \mathbf{R}^c/q
$$

is also uniformly random (in their marginal distribution). Therefore, it remains to show that c^* is still uniformly random even conditioned on fixed samples of $(\psi_0^*, \{\psi_i^*\}_i, \{\psi_{0,i}^*\}_i, \{\psi^c\})$.

As calculated above, we can unfold the expression of c^* as

$$
c^* = \langle z_0, b/q - x^{(1)} \rangle + \langle z_1^T \mathbf{R}_{\mathsf{BP}^*}, b/q - x^{(1)} \rangle + \mathcal{D}_{\mathbb{Z}_1, \alpha}
$$

We note that $b/q - x^{(1)} = \psi_0^* - x^{(1)} - x^{(2)} - \nu \mathbf{R}_{\mathsf{BP}^*} y_1/q$, thus if we show that

$$
\langle \mathbf{R}_{\mathsf{BP}^*} z_1, \nu \mathbf{R}_{\mathsf{BP}^*} y_1/q \rangle
$$

is close to uniform distribution (modulo 1), then c^* will also be close to the uniform distribution (modulo 1), as c^* is masked by this uniformly random number. Recall in hybrids, we set $y_1 = z_1/2 + (\text{shift})$, so it is sufficient to analyze

$$
\langle \mathbf{R}_{\mathsf{BP}^*} z_1, \nu \mathbf{R}_{\mathsf{BP}^*} y_1/q \rangle = \nu \langle \mathbf{R}_{\mathsf{BP}^*} z_1, \mathbf{R}_{\mathsf{BP}^*} z_1/q \rangle = \nu \|\mathbf{R}_{\mathsf{BP}^*}^* z_1\|^2/q
$$

By applying analysis of norm inductively on matrix $\mathbf{R}_{\mathsf{BP}^*}$, we can obtain that

$$
\|\mathbf{R}_{\mathsf{BP}^*}^* z_1\|^2/q \geq \frac{(\|\mathbf{R}_{0,j} z_1\| - \Theta(m^{1.5})\ell \|z_1\|)^2}{q}
$$

Table 1. Parameter description and simple example setting

Parameters	Description	Setting
n, m	Lattice dimension	$n = \lambda, m = n^2 \log n$
ℓ	Length of input to branching program	$\ell = n$
q	Modulus (resp. bit-precision)	Smallest prime $\geq n^{1.5} m^{2.5} \omega(\log n)$
α	Sampling error terms $\boldsymbol{e}, e$	$\frac{1}{n^{2.5} \log^3 n}$
β	Sampling correlation vector $\boldsymbol{y}$	$\alpha/2$
γ	Sampling correlation coefficient μ	$\frac{1}{n \log^{1.5} n}$
s	Sampling secret key $\boldsymbol{r}$	$3\beta/2$
η	Scaling parameter for $\mathbf{R}_{0,j}$	$\Theta(m\ell)$

where $\mathbf{R}_{0,j} \in \{-1, 1\}^{m \times m}$. Since vector $\boldsymbol{z}_1$ is sampled from Gaussian with width sq, so its two-norm is at least $\sqrt{m}(sq)$ with overwhelming probability. Then by analysis of norm, the distribution $\nu \|\mathbf{R}_{\mathsf{BP}*}^* \boldsymbol{z}_1\|^2 / q$ is a Gaussian distribution with width at least

$$d = \tau \frac{(\eta sqm - \Theta(m^2\ell)sq)^2}{q} = \frac{\gamma\alpha^2(\eta sqm - \Theta(m^2\ell)sq)^2}{\beta^2 q}$$

We recall again that ν was sampled from a Gaussian with parameter $\tau = \gamma\alpha^2/\beta^2$. By our setting of parameters, we have $d/\omega(\log(n)) \geq 1$. A Gaussian with such width is statistically close to uniform in the domain $\mathbb{Z}_1$. This completes the proof of Lemma 4.4. Further, Theorem 4.2 follows from Lemmas 4.3 and 4.4. A (flexibly) bi-deniable ABE from LWE then follows. □

4.3 Parameter Setting

The parameters in Table 1 are selected in order to satisfy the following constraints:

- To ensure correctness in Lemma 4.3, we have $\alpha sqm(\eta\sqrt{m} + 2m^{1.5}\ell) \leq 1/4$.
- To ensure deniability in Hybrid H_7, we have $d/\omega(\log(n)) > \frac{\gamma\alpha^2(\eta sqm - \Theta(m^2\ell sq))^2}{\beta^2 q\omega(\log(n))} > 1$.
- To ensure large enough LWE noise, we need $\alpha \geq (\sqrt{n}\log^{1+\delta} n)/q$.
- To apply the leftover hash lemma, we need $m \geq 2n\log(q)$.
- To ensure that the matrix $\mathbf{Q}$ in FakeRCoins is positive definite, we have $\beta \geq \alpha\gamma\sqrt{\eta\sqrt{m} + 2m^{1.5}\ell}$; To ensure that the matrix $\mathbf{Q}'$ in the security proof is positive definite, we have $\alpha \geq \tau\beta\sqrt{\eta\sqrt{m} + 2m^{1.5}\ell}$. This constraint will also imply that in the security proof, both $\mathbf{Q}'$ and $\mathbf{Q}' - \beta'\mathbf{I}_{m \times m}$ are positive definite (note $\beta' = \alpha/2$).
- To ensure hybrids H_3 and H_5 are well-defined, we have $s > \beta$ and $\beta > s/2$. Let $s := (3/2)\beta$.

Regev [36] showed that for $q > \sqrt{m}/\beta'$, an efficient algorithm for $\mathsf{LWE}_{n,m,q,\chi}$ for $\chi = \mathcal{D}_{\beta'q}$ (and $\beta'q \geq \sqrt{n}\omega(\log(n))$) implies an efficient quantum algorithm for approximating the SIVP and GapSVP problems, to within $\tilde{O}(n/\beta')$ approximation factors in the worst case. Our example parameter setting yields a bi-deniable AB-BTS based on the (quantum) hardness of solving $\mathsf{SIVP}_{\tilde{O}(n^{9.5})}$, respectively $\mathsf{GapSVP}_{\tilde{O}(n^{9.5})}$. (We write this term to additionally absorb the $(1/q^2)$ loss from our LWE to eLWE^+ reduction.) We leave further optimizing the lattice problem approximation factor to future work.

4.4   From AB-BTS to Flexible Bi-Deniable ABE

We present the instantiation of a flexible bi-deniable ABE using our AB-BTS scheme described above. We let $\Sigma' = (\mathsf{Setup}', \mathsf{DenSetup}', \mathsf{Keygen}', \mathsf{SampleP}', \mathsf{SampleU}', \mathsf{TestP}', \mathsf{FakeRCoins}', \mathsf{FakeSCoins}')$ be an AB-BTS scheme. Then the flexible bi-deniable ABE $\Sigma = (\mathsf{Setup}, \mathsf{DenSetup}, \mathsf{Keygen}, \mathsf{Enc}, \mathsf{DenEnc}, \mathsf{Dec}, \mathsf{SendFake}, \mathsf{RecFake})$ is:

– $\mathsf{Setup}(1^\lambda)$: Run algorithm $(\mathsf{pp}', \mathsf{msk}') \leftarrow \mathsf{Setup}'(1^\lambda)$ in AB-BTS and set $\mathsf{pp} = \mathsf{pp}', \mathsf{msk} = \mathsf{msk}'$.
– $\mathsf{DenSetup}(1^\lambda)$: Run algorithm $(\mathsf{pp}', \mathsf{msk}', \mathsf{fk}') \leftarrow \mathsf{DenSetup}'(1^\lambda)$ in AB-BTS and set $\mathsf{pp} = \mathsf{pp}', \mathsf{msk} = \mathsf{msk}', \mathsf{fk} = (\mathsf{fk}', \mathsf{msk}')$.
– $\mathsf{Keygen}(\mathsf{msk}, f)$: Run algorithm $\mathsf{sk}'_f \leftarrow \mathsf{Keygen}'(\mathsf{msk}, f)$ in AB-BTS and set $\mathsf{sk}_f = \mathsf{sk}'_f$.
– $\mathsf{Enc}(\mathsf{pp}, \boldsymbol{x}, \mu; (r_S^{(1)}, r_S^{(2)}))$: On input the message $\mu \in \{0, 1\}$, if $\mu = 0$, then run $c_i \leftarrow \mathsf{SampleU}'(\mathsf{pp}, \boldsymbol{x}; r_S^{(i)})$ for $i = 1, 2$, otherwise, $\mu = 1$, run $c_1 \leftarrow \mathsf{SampleU}'(\mathsf{pp}, \boldsymbol{x}; r_S^{(1)})$ and $c_2 \leftarrow \mathsf{SampleP}'(\mathsf{pp}, \boldsymbol{x}; r_S^{(2)})$. Output $\mathsf{ct}_{\boldsymbol{x}} = (c_1, c_2)$.
– $\mathsf{DenEnc}(\mathsf{pp}, \boldsymbol{x}, \mu; (r_S^{(1)}, r_S^{(2)}))$: On input the message $\mu \in \{0, 1\}$, then run $c_i \leftarrow \mathsf{SampleP}'(\mathsf{pp}, \boldsymbol{x}; r_S^{(i)})$ for $i = 1, 2$, otherwise, $\mu = 1$, run $c_1 \leftarrow \mathsf{SampleU}'(\mathsf{pp}, \boldsymbol{x}; r_S^{(1)})$ and $c_2 \leftarrow \mathsf{SampleP}'(\mathsf{pp}, \boldsymbol{x}; r_S^{(2)})$. Output $\mathsf{ct}_{\boldsymbol{x}} = (c_1, c_2)$.
– $\mathsf{Dec}(\mathsf{ct}_{\boldsymbol{x}}, \mathsf{sk}_f)$: If $f(\boldsymbol{x}) \neq 0$, then output $\perp$. Otherwise, parse $\mathsf{ct}_{\boldsymbol{x}} = (c_1, c_2)$ and run $b_i \leftarrow \mathsf{TestP}'(\mathsf{sk}_f, c_i)$ for $i = 1, 2$. Output 0 if the $b_1 = b_2$ and 1 if $b_1 \neq b_2$.
– $\mathsf{SendFake}(\mathsf{pp}, r_S, \mu, \mu')$: If $\mu = \mu'$, return r_S. If $(\mu, \mu') = (0, 1)$, then run $r_S^{*(2)} \leftarrow \mathsf{FakeSCoins}'(\mathsf{pp}, r_S^{(2)})$ and return $(r_S^{(1)}, r_S^{*(2)})$. Else if $(\mu, \mu') = (1, 0)$, run $r_S^{*(1)} \leftarrow \mathsf{FakeSCoins}'(\mathsf{pp}, r_S^{(1)})$ and return $(r_S^{*(1)}, r_S^{(2)})$.
– $\mathsf{RecFake}(\mathsf{pp}, \mathsf{fk}, \mathsf{ct}_{\boldsymbol{x}}, f, \mu')$: Parse $\mathsf{ct}_{\boldsymbol{x}} = (c_1, c_2)$ and use fk to decrypt the ciphertext $\mathsf{ct}_{\boldsymbol{x}}$ then obtain the plaintext μ. If $\mu = \mu'$, then run the honest key generation of the BTS scheme, i.e. $\mathsf{sk}'_f \leftarrow \mathsf{Keygen}'(\mathsf{msk}', f)$. Otherwise, run $\mathsf{sk}'_f \leftarrow \mathsf{FakeRCoins}'(\mathsf{pp}, \mathsf{fk}, c_{\mu+1}, f)$. Return sk'_f.

Similar to the work by Canetti et al. [14] and O'Neil et al. [35], the following, desired theorem can be proven in a straightforward manner.

Theorem 4.12. *Assume that Σ' is a flexible bi-deniable AB-BTS, as in Definition 3.2. Then Σ is a flexibly bi-deniable ABE, as in Definition 3.1.*

References

1. Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (H)IBE in the standard model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg (2010)
2. Alperin-Sheriff, J., Peikert, C.: Circular and KDM security for identity-based encryption. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 334–352. Springer, Heidelberg (2012)
3. Alperin-Sheriff, J., Peikert, C.: Faster bootstrapping with polynomial error. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 297–314. Springer, Heidelberg (2014)
4. Ananth, P., Brakerski, Z., Segev, G., Vaikuntanathan, V.: From selective to adaptive security in functional encryption. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9216, pp. 657–677. Springer, Heidelberg (2015)
5. Apon, D., Fan, X., Liu, F.-H.: Bi-deniable inner product encryption from LWE. IACR Cryptology ePrint Archive, 2015:993 (2015)
6. Barrington, D.A.M.: Bounded-width polynomial-size branching programs recognize exactly those languages in nc^1. J. Comput. Syst. Sci. **38**(1), 150–164 (1989)
7. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009)
8. Bendlin, R., Nielsen, J.B., Nordholt, P.S., Orlandi, C.: Lower and upper bounds for deniable public-key encryption. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 125–142. Springer, Heidelberg (2011)
9. Boneh, D., Gentry, C., Gorbunov, S., Halevi, S., Nikolaenko, V., Segev, G., Vaikuntanathan, V., Vinayagamurthy, D.: Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 533–556. Springer, Heidelberg (2014)
10. Boneh, D., Lewi, K., David, J.W.: Constraining pseudorandom functions privately. IACR Cryptology ePrint Archive **2015**, 1167 (2015)
11. Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.): 45th ACM STOC. ACM Press, June 2013
12. Boneh, D., Sahai, A., Waters, B.: Functional encryption: definitions and challenges. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 253–273. Springer, Heidelberg (2011)
13. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Boneh et al. [11], pp. 575–584 (2013)
14. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable encryption. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 90–104. Springer, Heidelberg (1997)

15. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: 28th ACM STOC, pp. 639–648. ACM Press, May 1996

16. Canetti, R., Goldwasser, S., Poburinnaya, O.: Adaptively secure two-party computation from indistinguishability obfuscation. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 557–585. Springer, Heidelberg (2015)

17. Angelo De Caro, Vincenzo Iovino, and Adam O'Neill. Deniable functional encryption. In Public-Key Cryptography - PKC –19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6–9, Proceedings, Part I, pp. 196–222, (2016)

18. Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai trees, or how to delegate a lattice basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer, Heidelberg (2010)

19. Dachman-Soled, D., Katz, J., Rao, V.: Adaptively secure, universally composable, multiparty computation in constant rounds. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 586–613. Springer, Heidelberg (2015)

20. Dachman-Soled, D., Liu, F.-H., Zhou, H.-S.: Leakage-resilient circuits revisited - optimal number of computing components without leak-free hardware, pp. 131–158 (2015)

21. Dodis, Y., Reyzin, L., Smith, A.: Fuzzy extractors: how to generate strong keys from biometrics and other noisy data. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 523–540. Springer, Heidelberg (2004)

22. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS, pp. 40–49. IEEE Computer Society Press, October 2013

23. Garg, S., Polychroniadou, A.: Two-round adaptively secure MPC from indistinguishability obfuscation. In: Dodis, Y., Nielsen, J. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 614–637. Springer, Heidelberg (2015)

24. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: 40th ACM STOC, pp. 197–206. ACM Press, May 2008

25. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 75–92. Springer, Heidelberg (2013)

26. Gilbert, H. (ed.): EUROCRYPT 2010. LNCS, vol. 6110. Springer, Heidelberg (2010)

27. Goldwasser, S., Kalai, Y.T., Popa, R.A., Vaikuntanathan, V., Zeldovich, N.: Reusable garbled circuits and succinct functional encryption. In: Boneh et al. [11], pp. 555–564 (2013)

28. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute-based encryption for circuits. In: Boneh et al. [11], pp. 545–554 (2013)

29. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Predicate encryption for circuits from LWE. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9216, pp. 503–523. Springer, Heidelberg (2015)

30. Gorbunov, S., Vinayagamurthy, D.: Riding on asymmetry: efficient ABE for branching programs. In: Iwata, T., et al. (eds.) ASIACRYPT 2015. LNCS, vol. 9452, pp. 550–574. Springer, Heidelberg (2015). doi:10.1007/978-3-662-48797-6_23

31. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: ACM CCS 2006, pp. 89–98. ACM Press, October/November 2006. Available as Cryptology ePrint Archive Report 2006/309

32. Katz, J., Thiruvengadam, A., Zhou, H.-S.: Feasibility and infeasibility of adaptively secure fully homomorphic encryption. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 14–31. Springer, Heidelberg (2013)
33. Micciancio, D., Mol, P.: Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In: Rogaway [37], pp. 465–484
34. Micciancio, D., Peikert, C.: Trapdoors for lattices: simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer, Heidelberg (2012)
35. O'Neill, A., Peikert, C., Waters, B.: Bi-deniable public-key encryption. In: Rogaway [37], pp. 525–542 (2011)
36. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: 37th ACM STOC, pp. 84–93. ACM Press, May 2005
37. Rogaway, P. (ed.): CRYPTO 2011. LNCS, vol. 6841. Springer, Heidelberg (2011)
38. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (eds.) 46th ACM STOC, pp. 475–484. ACM Press, May/June 2014