

A General Framework for Lattice-Based ABE Using Evasive Inner-Product Functional Encryption

謝耀慶 (Yao-Ching Hsieh)  Huijia Lin 罗辑 (Ji Luo) 

Paul G. Allen School of Computer Science & Engineering,
University of Washington, Seattle
{ychsieh,rachel,luoji}@cs.washington.edu

26 May 2024

Abstract

We present a general framework for constructing attribute-based encryption (ABE) schemes for arbitrary function class based on lattices from two ingredients, *i*) a noisy linear secret sharing scheme for the class and *ii*) a new type of inner-product functional encryption (IPFE) scheme, termed *evasive* IPFE, which we introduce in this work. We propose lattice-based evasive IPFE schemes and establish their security under simple conditions based on variants of evasive learning with errors (LWE) assumption recently proposed by Wee [EUROCRYPT '22] and Tsabary [CRYPTO '22].

Our general framework is modular and conceptually simple, reducing the task of constructing ABE to that of constructing noisy linear secret sharing schemes, a more lightweight primitive. The versatility of our framework is demonstrated by three new ABE schemes based on variants of the evasive LWE assumption.

- We obtain two ciphertext-policy ABE schemes for all polynomial-size circuits with a predetermined depth bound. One of these schemes has *succinct* ciphertexts and secret keys, of size polynomial in the depth bound, rather than the circuit size. This eliminates the need for tensor LWE, another new assumption, from the previous state-of-the-art construction by Wee [EUROCRYPT '22].
- We develop ciphertext-policy and key-policy ABE schemes for deterministic finite automata (DFA) and logspace Turing machines (L). They are the first lattice-based public-key ABE schemes supporting uniform models of computation. Previous lattice-based schemes for uniform computation were limited to the secret-key setting or offered only weaker security against bounded collusion.

Lastly, the new primitive of evasive IPFE serves as the lattice-based counterpart of pairing-based IPFE, enabling the application of techniques developed in pairing-based ABE constructions to lattice-based constructions. We believe it is of independent interest and may find other applications.

Contents

1	Introduction	1
1.1	Related Works	6
1.2	Technical Overview	7
2	Preliminaries	13
2.1	Attribute-Based Encryption	15
2.2	Lattices	16
2.3	Evasive LWE Assumption	17
3	Evasive Inner-Product Functional Encryption	19
3.1	Basic Construction — Single-Identity, No Structured Noises	22
3.2	Security as an Assumption	23
3.3	Security for Restricted Samplers from Evasive LWE	27
3.4	Identity-Based Scheme	29
3.5	Scheme with Structured Noises	35
4	Noisy Linear Garbling	37
5	Noisy Linear Garbling for Circuits	39
5.1	Correctness and Shortness	42
5.2	Security with Gaussian Noise	44
6	Ciphertext-Policy ABE from Short Noisy Linear Garbling	49
6.1	Correctness	51
6.2	Security	51
6.3	Unbounded Attribute and Summary	54
7	ABE for DFA	55
7.1	Noisy Linear Garbling for DFA	55
7.2	Construction of KP-ABE for DFA	57
7.3	Security of KP-ABE for DFA	60
7.4	CP-ABE for DFA and Summary	62
8	CP-ABE with Succinct Ciphertexts	64
8.1	Succinct Noisy Linear Garbling for Circuits	64
8.2	CP-ABE from General Noisy Linear Garbling	65
8.3	Security and Summary	68
	References	72
	 Appendix	
A	Noisy Linear Garbling for Boolean Circuits	78

1 Introduction

Attribute-based encryption (ABE) [SW05,GPSW06] enhances public-key encryption with the capability of enforcing precise access control over encrypted data. ABE primarily comes in two flavors, key-policy (KP-) and ciphertext-policy ABE (CP-ABE). In the setting of CP-ABE, data owners encrypt their private message along with an access policy that dictates who can access and decrypt the data. User keys, on the other hand, contain attributes that describe their credentials. For a user to decrypt a ciphertext and access the encrypted data, their attributes must align with the policy embedded in the ciphertext. Conversely, KP-ABE inverts this structure by using attributes to describe the encrypted data and incorporating access policies into user keys.

Over the past decade, the area of ABE has witnessed substantial progress, thanks to insights on leveraging the mathematical structures present in integer lattices and the learning with errors (LWE) assumption. Notably, Gorbunov, Vaikuntanathan, and Wee [GVW13] achieved a significant milestone by constructing the first KP-ABE scheme for circuits of unbounded size, but with *a priori* bounded depth, relying on the LWE assumption. Subsequent work by Boneh *et al.* [BGG⁺14] further enhanced KP-ABE efficiency, delivering the first scheme with both short ciphertexts and short secret keys, both scaling polynomially with the depth bound and independent of the circuit size.

However, despite these remarkable advances, several challenges have persisted for over a decade in the lattice-based ABE landscape. They include *i*) the search for ABE schemes accommodating circuits of unbounded depth, *ii*) the construction of CP-ABE schemes supporting all polynomial-size circuits, even confined to predetermined depth bounds, and *iii*) extending ABE to uniform models of computation, such as deterministic finite automata (DFA), logspace Turing machines, and RAM, which have much more succinct descriptions than circuits and can process arbitrarily long inputs. Until recently, the only lattice-based solution addressing challenges *i*) and *iii*) were due to Goldwasser *et al.* [GKP⁺13], who presented a KP-ABE scheme for RAM. However, it relies on the heavy tools of SNARKs and extractable witness encryption, which can be instantiated using lattice, but require strong knowledge-type assumptions. Beyond the scope of lattice-based constructions, one could attain these objectives using indistinguishability obfuscation (*iO*) [GGH⁺13, JLL23].

Recently, the landscape has evolved with the introduction of a novel class of assumptions known as *evasive LWE* [Wee22,Tsa22]. At a high level, evasive LWE captures a “generic-model view” on lattices. When presented with LWE samples $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_1^\top$, $\mathbf{s}^\top \mathbf{A} + \mathbf{e}_0^\top$ and a short (i.e., low-norm) “trapdoor” matrix \mathbf{K} sampled conditioning on $\mathbf{BK} = \mathbf{P}$, denoted as $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{P})$, where all terms are relative to a randomly sampled matrix \mathbf{B} and jointly sampled matrices \mathbf{A} and \mathbf{P} , it postulates that if the LWE samples $\mathbf{s}^\top \mathbf{P} + \mathbf{e}_2^\top$, $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_1^\top$, $\mathbf{s}^\top \mathbf{A} + \mathbf{e}_0^\top$ are pseudorandom when \mathbf{e}_2^\top is freshly generated, then so are $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_1^\top$, $\mathbf{s}^\top \mathbf{A} + \mathbf{e}_0^\top$ given $\mathbf{B}^{-1}(\mathbf{P})$. In simpler terms, adversaries are limited to utilizing the trapdoor to acquire LWE samples ($\mathbf{s}^\top \mathbf{P} + \mathbf{e}_2^\top$) and treating them as LWE samples with fresh noises. The work by [Wee22] argued that if this type of generic attacks fail, then all known cryptanalytic techniques, including the family of zeroizing attacks (e.g., [CHL⁺15,CVW18,HJL21,JLLS23]), would also be ineffective.

The introduction of the evasive LWE assumption has catalyzed new progress on

multiple fronts, including *i*) and *ii*) mentioned above. Wee [Wee22] presented a CP-ABE scheme for bounded-depth circuits under evasive LWE and another new assumption called tensor LWE. More recently, Hsieh, Lin, and Luo [HLL23a] achieved the first KP-ABE scheme for circuits of unbounded depth. Two other works [VWW22, Tsa22] constructed witness encryption schemes, which can be viewed as CP-ABE where the secret key for an attribute is the attribute itself. Furthermore, Water, Wee, and Wu [VWW22] proposed a multi-authority ABE for subset policies without resorting to random oracles.

These accomplishments underscore the potential of evasive LWE. However, each work uses evasive LWE in a way tailored to each construction, relying on a customized variant of evasive LWE following the generic-model principle. This often requires intricate handling of algebraic details. In this work, we ask:

*“Can we formulate a general framework for
constructing ABE schemes based on evasive LWE?”*

Such a framework would provide several advantages, including *i*) deepening our understanding, specifically by elucidating the role of evasive LWE, *ii*) facilitating the development of new ABE constructions, by making them more modular and conceptually simpler, and *iii*) streamlining security proofs, by encapsulating low-level algebraic intricacies within the framework and providing natural and useful tools.

Our Results. To this end, we propose a new notion called *evasive* inner-product functional encryption (IPFE), and present a general framework for constructing CP- and KP-ABE schemes for a class of functions from evasive IPFE and *noisy* linear secret sharing for the same class. Our framework can be seen as the lattice-world counterpart of the general framework for constructing ABE from pairing-based IPFE and (noiseless) linear secret sharing, as described in the prior works of [LL20a, LL20b]. Similar ideas have been employed in various previous pairing-based ABE constructions, which have been realized through different methods, including dual system encryption and hash proof systems.

Evasive Inner-Product Functional Encryption (IPFE). Let’s start by introducing evasive IPFE. Similar to a standard IPFE [ABDP15], evasive IPFE allows encrypting a vector $\mathbf{u} \in \mathbb{Z}_q^Z$ into a ciphertext $\text{ict}(\mathbf{u})$, and generating secret keys $\{\text{isk}(\mathbf{v}_j)\}_j$ tied to vectors $\mathbf{v}_j \in \mathbb{Z}_q^Z$, except that decryption reveals *noisy* (as opposed to exact) inner products $\{\mathbf{u}^\top \mathbf{v}_j + e_j\}_j$. The more substantial difference lies in the security requirements. Standard IPFE ensures that only the inner products are revealed and no other information of the encrypted vector \mathbf{u} is leaked (the key vectors \mathbf{v}_j are public). Evasive IPFE tries to capture a “generic-model view” similar to evasive LWE, but at the higher abstraction level of IPFE. It postulates that if the noisy inner products $\{\mathbf{u}^\top \mathbf{v}_j + e_j\}_j$ with “idealized”, i.e., freshly generated (hence independent), noises are pseudorandom, then the ciphertext is pseudorandom. More formally,

$$\begin{array}{ll} \text{if} & \{\mathbf{v}_j, \mathbf{u}^\top \mathbf{v}_j + e_j\}_j \approx \{\mathbf{v}_j, \$\}_j, \\ \text{then} & \text{impk}, \text{ict}(\mathbf{u}), \{\mathbf{v}_j, \text{isk}(\mathbf{v}_j)\}_j \approx \text{impk}, \$, \{\mathbf{v}_j, \text{isk}(\mathbf{v}_j)\}_j, \end{array}$$

where e_j ’s are independent Gaussian noises. Evasive IPFE is a convenient tool for generating LWE samples. Consider a simple example where the encrypted vector

is an LWE secret $\mathbf{u} = \mathbf{s}$ and the key vectors correspond to an LWE public matrix \mathbf{A} , which may or may not be random. Security ensures that as long as the noisy outputs $(\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)$ with *fresh* noises are pseudorandom, the ciphertext encrypting \mathbf{s} is pseudorandom.

Jumping ahead, in Section 3, we present candidate lattice-based evasive IPFE schemes, and show their security for restricted distributions of plaintext vectors (observed by all our ABE constructions) based on evasive LWE.

Noisy Linear Secret Sharing Scheme (LSSS). The other ingredient of our general framework is noisy LSSS. A noisy LSSS over \mathbb{Z}_q for a function f (say, of Boolean input and output), denoted as LSSS_f , consists of a non-zero vector \mathbf{w}_{out} and matrices $\mathbf{T}, \{\mathbf{W}_{\ell,b}\}_{\ell,b}$. Given randomness \mathbf{s} , the secret and the shares for an L -bit input \mathbf{x} are created from those matrices.

$$\begin{aligned} \text{LSSS}_f : & \quad \mathbf{w}_{\text{out}}, \mathbf{T}, \{\mathbf{W}_{\ell,b}\}_{\ell \in [L], b \in \{0,1\}}; \\ \text{Secret:} & \quad \mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}}; \quad \text{Shares:} \quad \mathbf{s}^\top \mathbf{T} + \mathbf{e}_t^\top, \{\mathbf{s}^\top \mathbf{W}_{\ell, \mathbf{x}[\ell]} + \mathbf{e}_\ell^\top\}_{\ell \in [L]}. \end{aligned}$$

If $f(\mathbf{x}) = 1$, the secret $(\mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}})$ can be approximately recovered from the shares, whereas if $f(\mathbf{x}) = 0$, the secret and the shares must be jointly pseudorandom. In addition, we say that LSSS_f is *short* if its $\mathbf{w}_{\text{out}}, \mathbf{T}, \{\mathbf{W}_{\ell,b}\}_{\ell,b}$ have small norms.

Noisy LSSS is a significantly weaker primitive than ABE in that it only considers a single function f and a single input \mathbf{x} , and that the shares for \mathbf{x} can depend on f . In contrast, ABE supports publishing secret keys and ciphertexts for many functions and many inputs, and both must be generated independent of the data on the other side. Interestingly, almost every known LWE-based ABE scheme implicitly defines an LSSS scheme. In the literature, noiseless and information-theoretically secure LSSS are known for low-depth functions [KW93, Ber84, BDHM92, Mul87]. Based on LWE, prior works have constructed a similar primitive, called nearly linear secret sharing scheme, for arbitrary circuits [QWW21, AY20, AWY20, LLL22]. See related works for a comparison.

Our General Framework for Constructing ABE. Our general framework uses evasive IPFE to generically “lift” noisy LSSS into full-fledged ABE schemes. For technical reasons elaborated in Section 1.2, we need to start from a *short* noisy LSSS.

Main Result (Construction 5). *Given a short noisy LSSS for any family of functions over \mathbb{Z}_q , KP- and CP-ABE schemes for the same family can be constructed from an evasive IPFE scheme over \mathbb{Z}_q .*

When instantiated with our evasive IPFE scheme in Section 3.4, the KP-ABE scheme has

$$|\text{mpk}| = \text{poly}(\log q), \quad |\text{sk}_f| = |\text{LSSS}_f| \text{poly}(\log q), \quad |\text{ct}_\mathbf{x}| = |\mathbf{x}| \text{poly}(\log q),$$

and the CP-ABE scheme has

$$|\text{mpk}| = \text{poly}(\log q), \quad |\text{sk}_\mathbf{x}| = |\mathbf{x}| \text{poly}(\log q), \quad |\text{ct}_f| = |\text{LSSS}_f| \text{poly}(\log q).$$

Here, $|\text{LSSS}_f|$ is the length of its defining vectors and matrices, and polynomial factors in the security parameter λ are hidden.

Using our general framework, the task of building ABE for an arbitrary function class reduces to the simpler task of constructing noisy LSSS for it. This has led to the following new ABE constructions.

- In Section 5, we design new *short* noisy LSSS for all arithmetic circuits based on LWE. The schemes are adapted from the arithmetic garbling proposed by [AIK11]. Combining them with our evasive IPFE scheme yields a new CP-ABE scheme for bounded-depth bounded-arithmetic circuits from evasive LWE. Compared with Wee’s construction [Wee22], our scheme eliminates the tensor LWE assumption, but does not have succinct ciphertext.
- For deterministic finite automata (DFA) and logspace Turing machines (L), noiseless LSSS are known, which have been used in prior works to construct pairing-based ABE for DFA and L [LL20a]. In lattice-based ABE landscape, handling uniform models of computation has been a challenge. The state-of-the-art ABE schemes for DFA only tolerate bounded collusion [AS17, Wee21], are in the secret-key setting [AMY19] (this also supports NFA), or do not come with a security proof [Wee21]. In Section 7, we apply our general framework to those LSSS with simple tweaks for uniform computation, obtaining both KP- and CP-ABE for DFA and L based on evasive LWE, the first provably secure lattice-based public-key ABE schemes tolerating unbounded collusion for these uniform models.

Corollaries 21 and 29 (CP-ABE for circuits and ABE for DFA). *Assuming LWE and evasive LWE, there exists a CP-ABE scheme for arithmetic circuits with*

$$|\text{mpk}| = \text{poly}(d, \log M), \quad |\text{sk}_{\mathbf{x}}| = |\mathbf{x}| \text{poly}(d, \log M), \quad |\text{ct}_C| = |C| \text{poly}(d, \log M),$$

where $|C|$ is the circuit size, d is a bound on the depth of C , and M is a bound on the magnitude of intermediate computation values.

Under the same set of assumptions, there exist ABE schemes for DFA.

$$\begin{array}{lll} \text{KP-ABE for DFA:} & |\text{mpk}| = O(1), & |\text{sk}_{\mathbf{x}}| = O(|\mathbf{x}|), & |\text{ct}_{\Gamma}| = O(|\Gamma|); \\ \text{CP-ABE for DFA:} & |\text{mpk}| = O(1), & |\text{sk}_{\Gamma}| = O(|\Gamma|), & |\text{ct}_{\mathbf{x}}| = O(|\mathbf{x}|). \end{array}$$

Here, $|\mathbf{x}|$ is the input length and $|\Gamma|$ is the number of states of the DFA. Each $O(\cdot)$ hides a polynomial factor in λ .

Upgrading Our General Framework. An unsatisfactory aspect of the above general framework is that it requires *short* noisy LSSS. On the other hand, there are noisy LSSS that are not *short*. For instance, the noisy LSSS underlying the ABE scheme of [BGG⁺14], as described below.

$$\begin{array}{l} \text{LSSS}_f : \quad \mathbf{w}_{\text{out}} = \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{u}), \quad \mathbf{T} = \mathbf{A}_0, \quad \{\mathbf{W}_{\ell,b} = \mathbf{A}_{\ell} - b\mathbf{G}\}_{\ell \in [L], b \in \{0,1\}}; \\ \text{Secret:} \quad \mathbf{s}^{\top} \mathbf{A}_f \mathbf{G}^{-1}(\mathbf{u}); \quad \text{Shares:} \quad \mathbf{s}^{\top} \mathbf{A}_0 + \mathbf{e}_0^{\top}, \quad \{\mathbf{s}^{\top} (\mathbf{A}_{\ell} - \mathbf{x}[\ell] \mathbf{G}) + \mathbf{e}_{\ell}^{\top}\}_{\ell \in [L]}. \end{array}$$

The vector \mathbf{u} and matrices $\mathbf{A}_0, \mathbf{A}_{\ell}$ ’s are random, hence they have large entries. (So are the entries of the gadget matrix \mathbf{G} .) We would like to upgrade our general framework to accommodate the BGG-based LSSS. Beyond the consideration of generality, that LSSS is the only known LSSS for circuits of succinct size, in particular, $|\text{LSSS}_f| = \text{poly}(d)$ grows polynomially with the maximum depth of the computation, instead of size. Once combined with our framework, we immediately obtain a CP-ABE scheme supporting bounded-depth circuits with *succinct* ciphertexts.

We achieve this by introducing a stronger variant of evasive IPFE, where the noisy inner products returned from decryption additionally contains structured noises in the form of $e'\mathbf{g}$, where \mathbf{g} is the gadget vector. More precisely, in an evasive IPFE scheme with structured noises, a ciphertext $\text{ict}(\mathbf{u})$ still encrypts a vector \mathbf{u} , but a secret key $\text{isk}(\mathbf{V})$ now encodes a matrix \mathbf{V} whose number of columns is the dimension of \mathbf{g} . Decryption produces $(\mathbf{u}^\top \mathbf{V} + e'\mathbf{g}^\top + \mathbf{e}^\top)$, where \mathbf{e} is a small noise vector as before and $e'\mathbf{g}$ is the structured noise. Security is enhanced correspondingly. If the noisy outputs with idealized fresh noises e', \mathbf{e} are pseudorandom, so is the ciphertext, i.e.,

$$\begin{array}{ll} \text{if} & \{\mathbf{V}_j, \mathbf{u}^\top \mathbf{V}_j + e'_j \mathbf{g}^\top + \mathbf{e}_j^\top\} \approx \{\mathbf{V}_j, \$\}, \\ \text{then} & \text{impk}, \text{ict}(\mathbf{u}), \{\mathbf{V}_j, \text{isk}(\mathbf{V}_j)\}_j \approx \text{impk}, \$, \{\mathbf{V}_j, \text{isk}(\mathbf{V}_j)\}_j. \end{array}$$

Using evasive IPFE with structured noises, we can now “lift” arbitrary noisy LSSS, without any norm restriction, into full-fledged ABE schemes.

Construction 10 and Corollary 34. *Given a (not necessarily short) noisy LSSS for any family of functions over \mathbb{Z}_q , KP- and CP-ABE schemes for the same family can be constructed from evasive IPFE with structured noises over \mathbb{Z}_q .*

When instantiated with our evasive IPFE with structured noises in Section 3.5, the resultant ABE schemes have the same asymptotic sizes as those in Main Result (Construction 5) — compared with which, it is now possible to use LSSS with smaller sizes. Assuming LWE and a variant of evasive LWE, there exists a CP-ABE scheme for arithmetic circuits with

$$|\text{mpk}| = \text{poly}(d, \log M), \quad |\text{sk}_{\mathbf{x}}| = |\mathbf{x}| \text{poly}(d, \log M), \quad |\text{ct}_r| = \text{poly}(d, \log M).$$

Constructions of Evasive IPFE. We propose candidate lattice-based evasive IPFE schemes, both with and without structured noises in Section 3, and extensively study them. We show that when the input vectors \mathbf{u}_i have certain form, namely $\mathbf{u}_i^\top = \mathbf{s}^\top \mathbf{A}_i$ for arbitrary \mathbf{A}_i 's sampled using public coins and independent and uniformly random \mathbf{s} (e.g., $\mathbf{u}^\top = \mathbf{s}^\top = \mathbf{s}^\top \mathbf{I}$), then we can base the security of evasive IPFE on variants of evasive LWE. All ABE constructions in this work uses evasive IPFE to encrypt vectors of this form. For evasive IPFE without structured noises, we rely on the usual type of evasive LWE assumption, where all noises involved are small (Assumption 4). For the version with structured noises, we rely on a variant of evasive LWE with structured noises (Assumption 3).

In order to reduce to evasive LWE, the ABE constructions must use a single IPFE instance, instead of many instances with independently sampled master public key. However, we often want to impose some restriction on what inner products are produced, e.g., only computing $\mathbf{u}_1^\top \mathbf{v}_1$ and $\mathbf{u}_2^\top \mathbf{v}_2$, without revealing $\mathbf{u}_1^\top \mathbf{v}_2$ or $\mathbf{u}_2^\top \mathbf{v}_1$. To achieve this using a single instance of evasive IPFE, we extend the interface to be identity-based (see Section 3.4), so that a pair of ciphertext $\text{ict}_{\text{id}}(\mathbf{u})$ and secret key $\text{isk}_{\text{id}'}(\mathbf{v})$ only reveals the noisy inner product if $\text{id} = \text{id}'$.

A nice addition is that the study of evasive IPFE already employs several pairing-based techniques. Nevertheless, the constructions and security proofs are intricate and require ironing out many details. We hope our evasive IPFE schemes serve as useful tools in future works.

1.1 Related Works

Lattice-Based Linear Secret Sharing Schemes for Circuits. The works of [AY20, AWY20] constructed CP-ABE with succinct ciphertext for NC^1 relying on LWE and the generic (pairing) group model (GGM). The work of [LLL22] constructed CP-ABE for circuits with constant-size keys (i.e., of size $\text{poly}(\lambda)$) and succinct ciphertext (of size $|\mathbf{x}| \text{poly}(\lambda)$, independent of depth). Both works combine a nearly linear secret sharing scheme with a pairing-based IPFE.

Their nearly LSSS is very similar to our noisy LSSS. Both are noisy and if $f(\mathbf{x}) = 1$, evaluation reveals a random secret s perturbed by a small noise e . The major differences are *i)* nearly LSSS requires linear reconstruction and the noise e to be small, which are not required by noisy LSSS, and *ii)* noisy LSSS shares must have the specific form $\mathbf{s}^\top \mathbf{T} + \mathbf{e}_t^\top, \mathbf{s}^\top \mathbf{W}_{\ell, \mathbf{x}[\ell]} + \mathbf{e}_\ell^\top$, where the noises are small. These requirements are tailored for combination with pairing-based IPFE and evasive IPFE, respectively. We compare their secret sharing scheme and techniques with ours.

Based on the laconic function evaluation (LFE) protocol of [QWW18], which in turn is based on the circuit KP-ABE scheme of [BGG⁺14], the works of [AY20, AWY20] proposed a nearly linear secret sharing scheme for NC^1 and that of [LLL22] one for circuits. They then “lift” their nearly LSSS schemes to ABE schemes using pairing-based IPFE — their ABE schemes use pairing-based IPFE to compute rerandomized nearly LSSS shares. As a result, the computed shares reside in the exponent of the pairing group. By the fact that reconstruction is linear, the perturbed secret ($s + e$) can be recovered in the exponent if $f(\mathbf{x}) = 1$. However, to recover the secret s in the clear, the schemes rely on exhaustive search of the noise e . Therefore, e must be polynomially bounded, which is particularly difficult to achieve when evaluating high-depth circuits [LLL22]. The fact that e must be small also makes the security proof more challenging, as one cannot rely on noise smudging.

Differently, in our general framework and hence CP-ABE schemes, the secret shares are computed in the clear and the evaluation noise e just need to be sufficiently small compared to the modulus, and evaluation can be non-linear. However, we require randomization of the shares in the form of $\mathbf{s}^\top \mathbf{T} + \mathbf{e}_t^\top, \{\mathbf{s}^\top \mathbf{W}_{\ell, \mathbf{x}[\ell]} + \mathbf{e}_\ell^\top\}_{\ell \in [L]}$ with small noises, in order to match the precondition of evasive LWE, where these small noises are “idealized” to be fresh.

A technical relation is that their ABE schemes rely on GGM or knowledge-type assumption (the Knowledge of OrthogonALity Assumption, or “KOALA”) on the pairing group to argue that correlated instances of secret sharing for different computation $\{(f_i, \mathbf{x}_j)\}_{i,j}$ where $f_i(\mathbf{x}_j) = 0$ for all i, j (resulting from decrypting all pairs of secret keys and ciphertexts) are jointly secure. In our ABE scheme, we rely on evasive IPFE (and by reduction, evasive LWE) to assert the security of the many correlated instances of secret sharing. In folklore, evasive LWE is regarded as a lattice counterpart of GGM, although there has been no formal known connections between them.

The works of [AY20, AWY20, LLL22] *combine* pairing techniques with lattice ones, relying on pairing-based IPFE as well as LWE. We *translate* pairing techniques to the lattice world, creating a lattice implementation of IPFE, suitable for ABE when combined with LWE. Evasive IPFE is reminiscent to the very strong simulation security of IPFE in GGM proven in [LLL22], which can be seen as a step towards understanding the connection between GGM and evasive LWE.

Noisy IPFE. The work of [Agr19], followed by [AP20], proposed a primitive called noisy IPFE, which evaluates noisy inner products, i.e., $(\mathbf{u}^\top \mathbf{v} + e)$. It is defined with strong security property. For any two plaintext vectors \mathbf{u}, \mathbf{u}' and key vectors $\mathbf{v}_1, \dots, \mathbf{v}_J$ such that their inner products are approximately equal, i.e., $(\mathbf{u} - \mathbf{u}')^\top \mathbf{v}_j$ is small for all j , the ciphertext encrypting \mathbf{u} and that encrypting \mathbf{u}' must be indistinguishable. The rationale is that the noises resulting from decryption is sufficient to “smudge” the difference between the inner products. This primitive is strong enough to imply indistinguishability obfuscation when combined with other well-studied assumptions.

Our evasive IPFE has a different security requirement. It states that if the inner products plus fresh noises are pseudorandom, i.e., $\{\mathbf{u}_i^\top \mathbf{v}_j + e_{ij}\} \approx \{\$\}$, then the ciphertexts encrypting \mathbf{u}_i 's are pseudorandom given the keys. We achieve this under evasive LWE for a general class of distributions of inputs.

ABE for DFA from Lattices. The work of [Wee21] proposed a candidate public-key KP-ABE for DFA based on lattices. In terms of methods, it relies on the DFA garbling implicit in [Wat12] (made explicit in [LL20a]), which we also use for our scheme. While many known ABE for DFA [Wat12, GWW19, LL20a, GW20, Wee21] (and this work) are roughly based on the idea of *securely* computing a pseudorandom garbling for DFA, to our knowledge, only this work and [LL20a] explicitly use *functional encryption* (concretely, IPFE) to do so. In terms of results, in [Wee21], there was no reductionist proof of security but some preliminary cryptanalysis, in which an idea in the line of evasive LWE is used (replacing a trapdoor preimage by the intended usage result, with fresh noises). The development of evasive LWE is later to [Wee21]. While we do not immediately see a proof from the version of evasive LWE in our work or [Wee22, Tsa22], it is plausible that another variant of evasive LWE could be formulated to prove its security (or a slightly modified variant thereof).

1.2 Technical Overview

We give an overview of our techniques. Our starting point is a well-established method of constructing ABE using linear secret sharing schemes (LSSS) and pairing.

Recap of ABE from LSSS and Pairing-Based IPFE. A traditional LSSS is noiseless and information-theoretic. An LSSS for an arithmetic function $f : \mathbb{Z}^L \rightarrow \{0, 1\}$ is specified by vector $\mathbf{w}_{\text{out}} \neq \mathbf{0}$ and matrices $\mathbf{T}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,x}\}_{\ell \in [L]}$ over \mathbb{Z}_q , each of which has n_f rows. Given randomness $\mathbf{s} \in \mathbb{Z}_q^{n_f}$, the secret is $\mathbf{s}^\top \mathbf{w}_{\text{out}}$ and the shares are

$$\mathbf{s}^\top \mathbf{T} \quad \text{and} \quad \{\mathbf{s}^\top (\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,x})\}_{\ell \in [L]}.$$

For correctness, if $f(\mathbf{x}) = 1$, the shares $\mathbf{s}^\top \mathbf{T}, \{\mathbf{s}^\top \mathbf{W}_{\ell,x[\ell]}\}_{\ell \in [L]}$ must determine the secret $\mathbf{s}^\top \mathbf{w}_{\text{out}}$, where $\mathbf{W}_{\ell,x[\ell]} = \mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,x}$. For security, if $f(\mathbf{x}) = 0$ and \mathbf{s} is uniformly random, the shares must be independent of $\mathbf{s}^\top \mathbf{w}_{\text{out}}$.

To obtain (KP- or CP-) ABE from LSSS, pairing is used so that when decrypting a ciphertext using a key, tied to f, \mathbf{x} , it will first compute the LSSS $_f$ shares of \mathbf{x} encoded in the exponent of the target group, which can then be used to recover the secret (again, in the exponent) if $f(\mathbf{x}) = 1$. For the construction to be secure, the computed shares should use good randomness, which must be kept hidden. The mechanism of computing the shares vary across known constructions, and oftentimes it can be

regarded as inner-product functional encryption (IPFE) [ABDP15]. Since the shares are computed in the exponent, pseudorandomness can often be obtained from DDH-style rerandomization, e.g., in CP-ABE, keys contain random r , ciphertexts contain $\mathbf{s}^\top \mathbf{T}$, $\{\mathbf{s}^\top \mathbf{W}_{\ell, \mathbf{x}[\ell]}\}_{\ell \in [L]}$ with random \mathbf{s} , and the computed shares are $r\mathbf{s}^\top \mathbf{T}$, $\{r\mathbf{s}^\top \mathbf{W}_{\ell, \mathbf{x}[\ell]}\}_{\ell \in [L]}$ in the exponent of the target group. Intuitively, by the decisional Diffie–Hellman assumption, the computed shares, using $r\mathbf{s}$ as its randomness, is indistinguishable from being created with fresh randomness. This intuition has been implemented in multiple ways, e.g., dual system encryption, hash proof systems, or using slotted IPFE.

ABE from LSSS and Lattices. Pairing-based ABE only supports low-depth computations, because the expressive power of polynomial-size LSSS is limited in NC [KW93, Ber84, BDHM92, Mul87]. To circumvent this lower bound, we must relax the notion of LSSS to support richer classes of functions. We consider LSSS with noises, termed *noisy linear secret sharing* in this work. Such a scheme for f is again specified by \mathbf{w}_{out} , \mathbf{T} , $\{\mathbf{W}_{\ell, 0}, \mathbf{W}_{\ell, \mathbf{x}}\}_{\ell \in [L]}$, but the secret and the shares now contain noises.

$$\begin{aligned} \text{Secret: } & \mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}}; \\ \text{Shares: } & \mathbf{s}^\top \mathbf{T} + \mathbf{e}_\dagger \quad \text{and} \quad \{\mathbf{s}^\top (\mathbf{W}_{\ell, 0} + \mathbf{x}[\ell] \mathbf{W}_{\ell, \mathbf{x}}) + \mathbf{e}_\ell^\dagger\}_{\ell \in [L]}. \end{aligned}$$

We additionally allow the secret sharing to have a reusable part \mathbf{R} (always given and not randomized). If $f(\mathbf{x}) = 1$, correctness only guarantees recovering the secret approximately. For security, if $f(\mathbf{x}) = 0$, the secret and the shares should be jointly pseudorandom even given $\mathbf{R}, \mathbf{w}_{\text{out}}, \mathbf{T}, \mathbf{W}$'s. Following the previous pairing-based general framework, we now need an appropriate IPFE that can compute and rerandomize the noisy secret shares.

Insufficiency of Existing IPFE. When using pairing-based IPFE, since the shares are computed in the exponent, one can only reconstruct the secret plus noise in the exponent. To be able to recover the secret in the clear, the noise must be polynomially small to allow for exhaustive search. However, keeping the noise small is difficult when evaluating arbitrary depth circuits.

To get around this limitation, we ask whether lattice-based IPFE schemes such as [ABDP15, ALS16] can help. They yield decryption results in the clear, dispensing the limitation on the magnitude of noise. However, it is unclear how to use these scheme to generate the noises needed in each rerandomized secret shares. In fact, Agrawal [Agr19] together with Pellet-Mary [AP20] showed that a noisy IPFE scheme capable of computing inner products added with (computationally) good Gaussian noises implies $i\mathcal{O}$ when combined with other standard assumption.

Evasive IPFE. We solve the aforementioned challenges with a new primitive called *evasive IPFE*. For convenience, we directly define the identity-based variant. Each key is for a vector \mathbf{v} under some identity id' , and each ciphertext, \mathbf{u} and id . Decryption yields an approximation of $\mathbf{u}^\top \mathbf{v}$ if and only if $\text{id} = \text{id}'$.

$$\left. \begin{array}{l} \text{isk}_{\text{id}'}(\mathbf{v}) \\ \text{ict}_{\text{id}}(\mathbf{u}^\top) \end{array} \right\} \xrightarrow{\text{Dec}} \begin{cases} \mathbf{u}^\top \mathbf{v} + e, & \text{if } \text{id} = \text{id}'; \\ \perp, & \text{if } \text{id} \neq \text{id}'. \end{cases}$$

We also denote batches of keys and ciphertexts as $\text{isk}_{\text{id}}(\mathbf{V})$ and $\text{ict}_{\text{id}}(\mathbf{U}^\top)$ so that their pairwise decryptions approximate $\mathbf{U}^\top \mathbf{V}$.

For security, given arbitrarily many keys, the ciphertexts hide \mathbf{u} 's if all the inner products that can be computed are jointly pseudorandom, when they are added with fresh Gaussian noises. More precisely, given secret keys generated for \mathbf{v}_j under identity id'_j and ciphertexts for \mathbf{u}_i under id_i , the security property stipulates that

$$\begin{aligned} \text{if} & \quad \text{aux}, \{\mathbf{u}_i^\top \mathbf{v}_j + e_{i,j}\}_{\text{id}_i=\text{id}'_j} \approx \text{aux}, \{\$\}_{\text{id}_i=\text{id}'_j}, \\ \text{then} & \quad \text{aux}, \{\text{isk}_{\text{id}'_j}(\mathbf{v}_j)\}_j, \{\text{ict}_{\text{id}_i}(\mathbf{u}_i)\}_i \approx \text{aux}, \{\text{isk}_{\text{id}'_j}(\mathbf{v}_j)\}_j, \{\text{ict}_{\text{id}_i}(\$\})_i. \end{aligned}$$

Here, $e_{i,j}$'s are fresh (Gaussian) noises, and aux is auxiliary information about \mathbf{v}, \mathbf{u} 's – think of it as the public coins for sampling \mathbf{v} 's.

ABE from Noisy Linear Secret Sharing and Evasive IPFE. Given evasive IPFE and noisy LSSS, our general framework for constructing KP- and CP-ABE is modular and simple. For example, to construct CP-ABE, each ciphertext for a function f consists of a series of IPFE ciphertexts encoding $\mathbf{S}^\top \mathbf{W}$'s, where \mathbf{S} is an LWE secret matrix, and each key consists of IPFE keys encoding LWE public vector \mathbf{r} and input \mathbf{x} . Their decryption produces rerandomized secret shares of the form $(\mathbf{r}^\top \mathbf{S}^\top \mathbf{W} + \mathbf{e}^\top)$. Below is an example.

$$\begin{aligned} \text{sk}_{\mathbf{x}} : & \quad \text{isk}_0(\lfloor q/2 \rfloor, \mathbf{r}), \quad \{\text{isk}_\ell(\mathbf{r}, \mathbf{x}[\ell] \mathbf{r})\}_{\ell \in [L]}; \\ \text{ct}_f : & \quad \text{ict}_0(\mathbf{0}, \mathbf{T}^\top \mathbf{S}), \quad \{\text{ict}_\ell(\mathbf{W}_{\ell,0}^\top \mathbf{S}, \mathbf{W}_{\ell,\mathbf{x}}^\top \mathbf{S})\}_{\ell \in [L]}, \\ & \quad \text{ict}_0(\mu, \mathbf{w}_{\text{out}}^\top \mathbf{S}). \end{aligned}$$

The decryption yields a noisy version of $\mathbf{W}^\top \mathbf{r}$'s as desired.

The security proof of ABE involves invoking the security of evasive IPFE to argue that all IPFE ciphertexts are hiding, which only applies if all noisy inner products are pseudorandom. Recall that in the ABE security game, there are multiple keys, and the inner products are

$$\begin{aligned} \text{for } \mathbf{x}_j : & \quad (\mathbf{S} \mathbf{r}_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top, \quad \{(\mathbf{S} \mathbf{r}_j)^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\mathbf{x}}) + \mathbf{e}_{j,\ell}^\top\}_{\ell \in [L]}, \\ & \quad (\mathbf{S} \mathbf{r}_j)^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \mu \cdot \lfloor q/2 \rfloor. \end{aligned}$$

Those are exactly the shares and secrets generated using randomness $\mathbf{S} \mathbf{r}_j$, but the components as-is are not pseudorandom, because \mathbf{r}_j 's are public, and the same secret matrix \mathbf{S} and matrices \mathbf{T}, \mathbf{W} 's are reused in all shares for different \mathbf{x}_j 's. Hopefully, if we could somehow change the randomness from $\mathbf{S} \mathbf{r}_j$ to $(\mathbf{S} \mathbf{r}_j + \mathbf{e}_{r,j}^\top)$ with independent Gaussian $\mathbf{e}_{r,j}$'s, then by LWE (for secret \mathbf{S}), the rerandomization generates pseudorandomness, i.e., $\mathbf{S} \mathbf{r}_j + \mathbf{e}_{r,j}^\top \approx \$$, and then so are the shares. We show that how to achieve this in two ways.

Short Secret Sharing. If we require that the secret sharing be short, i.e., the matrices $\mathbf{w}_{\text{out}}, \mathbf{T}, \mathbf{W}_{\ell,\mathbf{x}[\ell]}$'s have sufficiently small norm, then thanks to flooding by $\mathbf{e}_{t,j}$'s,

$$\{(\mathbf{S} \mathbf{r}_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top\}_{j \in [J]} \approx_s \{(\mathbf{S} \mathbf{r}_j + \mathbf{e}'_{t,j})^\top \mathbf{T} + \mathbf{e}_{t,j}^\top\}_{j \in [J]}.$$

The same holds for all the shares and secrets, achieving rerandomization.

Evasive IPFE with Structured Noises. Alternatively, we can augment the notion of evasive IPFE to include structured noises to work with arbitrary noisy LSSS. Instead of having only small noises, the noisy inner products now look like $(\mathbf{u}^\top \mathbf{V} + \mathbf{e}^\top \mathbf{G} + (\mathbf{e}')^\top)$,

where \mathbf{V} is associated with the secret key. Correspondingly, security holds if the noisy inner products with fresh noises \mathbf{e}' and \mathbf{e} are pseudorandom.

To use such evasive IPFE with structured noises to rerandomize secret shares, we need an additional trick. The problem is that the structured noises are of the form $\mathbf{e}^\top \mathbf{G}$ instead of $\mathbf{e}^\top \mathbf{W}_{\ell, \mathbf{x}_j[\ell]}$, and the \mathbf{e} here is different for each decryption (e.g., for computing the shares corresponding to $\mathbf{x}_j[1]$ and $\mathbf{x}_j[2]$ for the *same* j). To get consistent $\mathbf{e}_j^\top \mathbf{W}_{\ell, \mathbf{x}_j[\ell]}$ for each \mathbf{x}_j (key) independent of ℓ (input bit index), we perform a noisy secret sharing

$$\underbrace{(\mathbf{r}_j^\top \mathbf{S} + \mathbf{e}_j^\top)} \mathbf{W}_{\ell, \mathbf{x}_j[\ell]} = \underbrace{\mathbf{r}_j^\top \mathbf{S} \mathbf{W}_{\ell, \mathbf{x}_j[\ell]}} - \underbrace{(\mathbf{r}_j')^\top \mathbf{Q} \mathbf{G}^{-1} (\mathbf{W}_{\ell, \mathbf{x}_j[\ell]})} + \underbrace{((\mathbf{r}_j')^\top \mathbf{Q} + \mathbf{e}_j^\top \mathbf{G}) \mathbf{G}^{-1} (\mathbf{W}_{\ell, \mathbf{x}_j[\ell]})},$$

where the wavy underlines indicate small noises (without \mathbf{G} structure) and each noisy term on the right-hand side can be computed using evasive IPFE, without and with structured noises respectively. Importantly, each ABE key only uses one structured noise \mathbf{e}_j , and the shares for the L input bits use the same \mathbf{e}_j in $\mathbf{e}_j^\top \mathbf{W}_{\ell, \mathbf{x}_j[\ell]}$.

Short Noisy Linear Secret Sharing for Polynomial-Size Circuits. Our construction of short noisy linear secret sharing scheme for circuits is inspired by the arithmetic garbling scheme of [AIK11] and deals with the circuit gate by gate. In this paragraph, we will use “labels” for “shares” as it is customary for circuits (noisy linear secret sharing is a partially hiding garbling).

To begin, each gate i is associated with two random low-norm matrices $\mathbf{W}_{i,0}, \mathbf{W}_{i,\times}$, which we call the label function of gate i . The label (before randomization) encoding some value x is $\mathbf{W}_{i,x} = \mathbf{W}_{i,0} + x\mathbf{W}_{i,\times}$, and its randomized version is $(\mathbf{s}^\top \mathbf{W}_{i,x} + \mathbf{e}_{i,x}^\top)$. Suppose gate i has its input wires connected to gates i_1, i_2 and their output values are x, x_1, x_2 , we can sample a random low-norm matrix \mathbf{U}_i and rewrite

$$\begin{aligned} \mathbf{W}_{i,0} + (x_1 + x_2)\mathbf{W}_{i,\times} &= (\mathbf{U}_i + x_1\mathbf{W}_{i,\times}) + (\mathbf{W}_{i,0} - \mathbf{U}_i + x_2\mathbf{W}_{i,\times}), \\ \mathbf{W}_{i,0} + x_1x_2\mathbf{W}_{i,\times} &= x_2(\mathbf{U}_i + x_1\mathbf{W}_{i,\times}) + (\mathbf{W}_{i,0} - x_2\mathbf{U}_i), \end{aligned}$$

depending on whether gate i is an addition or multiplication gate. We denote the decomposed terms that are affine in x_1 as $\widetilde{\mathbf{W}}_{i_1, x_1}$, and those affine in x_2 as $\widetilde{\mathbf{W}}_{i_2, x_2}$ — note that $\widetilde{\mathbf{W}}_{i,x}$ is contributed by all the fan-outs of gate i . The randomized, noisy $\mathbf{s}^\top \mathbf{W}_{i,x}$ can be obtained from the noisy versions of $\mathbf{s}^\top \widetilde{\mathbf{W}}_{i_1, x_1}$ and $\mathbf{s}^\top \widetilde{\mathbf{W}}_{i_2, x_2}$. We call $\widetilde{\mathbf{W}}$'s and $\mathbf{s}^\top \widetilde{\mathbf{W}}$'s the expanded labels, and consequently, \mathbf{W} 's and $\mathbf{s}^\top \mathbf{W}$'s the shrunken labels.

For each gate i , we also sample and publish a low-norm matrix \mathbf{R}_i , and publish in $\mathbf{s}^\top \mathbf{T}$ (the “garbled table entries” for gate i)

$$\mathbf{s}^\top (\mathbf{W}_{i,0} \mathbf{R}_i + \widetilde{\mathbf{W}}_{i,0}) + (\mathbf{e}'_{i,0})^\top, \quad \mathbf{s}^\top (\mathbf{W}_{i,\times} \mathbf{R}_i + \widetilde{\mathbf{W}}_{i,\times}) + (\mathbf{e}'_{i,\times})^\top.$$

The values of \mathbf{R}_i 's and noisy $\mathbf{s}^\top \mathbf{T}$ are given regardless of the input values. Note that these values enable computing the expanded label $\mathbf{s}^\top \widetilde{\mathbf{W}}_{i,x}$ from the shrunken label $\mathbf{s}^\top \mathbf{W}_{i,x}$ for any value x that gate i happens to take. We also include $\mathbf{s}^\top (\mathbf{W}_{|C|,0} \mathbf{R}_{|C|} + \mathbf{w}_{\text{out}})$ to facilitate the recovery of secret when $C(\mathbf{x}) = 0$. The shares consist of the input labels as well as the garbled table entries for all gates.

Evaluation also proceeds gate by gate, starting from the input gates. For each gate i with value x , the evaluation procedure recovers the shrunken label, then use \mathbf{R}_i and the garbled table to recover the expanded label, which is later used to recover the shrunken labels for gates taking gate i as input. Eventually, if the output is 0, we approximately recover $\mathbf{s}^\top \mathbf{w}_{\text{out}}$ from $\mathbf{s}^\top \mathbf{W}_{|C|,0}$ and $\mathbf{s}^\top (\mathbf{W}_{|C|,0} \mathbf{R}_{|C|} + \mathbf{w}_{\text{out}})$ and $\mathbf{R}_{|C|}$.

We sketch the security proof. In the first step, we replace all the noisy $\mathbf{s}^\top \mathbf{W}_{i,0}$'s and $\mathbf{s}^\top \mathbf{W}_{i,\times}$'s by random $\delta_{i,0}^\top$ and $\delta_{i,\times}^\top$, using the flipped LWE assumption. Then, for a particular input \mathbf{x} , we simulate $\delta_{i,0}$ and $\delta_{i,\times}$ using δ_{i,x_i} (call this “active”) and $\delta_{i,\times}$, where x_i is the value of gate i when evaluated on input \mathbf{x} . The key observation here is that in such hybrids, the active expanded labels and garbled table entries no longer depend on the “ \times ” ones. To argue the input labels and the garbled table are jointly pseudorandom, we deal with active labels/table and “ \times ” ones in two passes. In the first pass, we work on the “ \times ” labels/table bottom-up (output gate to input gate). This step is computational and involves flipped LWE for public matrix \mathbf{R} – for each “ \times ” table ($\delta_{i,\times}^\top \mathbf{R}_i + \tilde{\delta}_{i,\times}^\top$), the secret $\delta_{i,\times}^\top$ will be absent everywhere else when the proof comes to this table, so flipped LWE ensures that the first term is pseudorandom, hiding $\tilde{\delta}_{i,\times}$, thus removing information about the “ \times ” labels of its fan-outs. In the second pass, we work on the active labels/table top-down. This step is statistical using the information-theoretic garbling security (so technically security is already in place when we finish the first pass), and the top-down order is just a pedagogy device for understanding the proof.

Succinct CP-ABE for Circuits. A (non-short) noisy linear secret sharing scheme can be distilled from the celebrated KP-ABE for circuits due to [BGG⁺14], as observed by [AWY20]. The scheme is succinct – the sizes of \mathbf{R} , \mathbf{T} , \mathbf{W} 's only depend on the depth, not the size of the circuit. Combined with evasive IPFE with structured noises, we immediately obtain succinct CP-ABE.

Secret Sharing and ABE for DFA. The existing linear secret sharing scheme for DFA [Wat12,LL20a] can be cast as a noisy linear secret sharing and is short. The definition of noisy linear secret sharing inherently considers fixed input length, while in (say, CP-) ABE for uniform computation like DFA, we would like a ciphertext (tied to DFA) to potentially authorize a key (tied to input) of arbitrary length, and conversely a key to be accepted by a ciphertext of arbitrarily many states. In the previous scheme, ct_ℓ determines the input length it accepts, so the same construction cannot be directly used.

We follow the paradigm in [Wat12,LL20a] and exploit the locality of DFA computation. Namely, each DFA step is simply reading an input bit from a fixed location and transitioning the previous state to the next. At a very high level, the most important shares (related to state transitions) are of the form $(\mathbf{s}_\ell^\top \mathbf{\Gamma}_{\mathbf{x}[\ell]} - \mathbf{s}_{\ell-1}^\top)$, where $\mathbf{\Gamma}_0, \mathbf{\Gamma}_1$ are determined by the DFA state transition function, and $\mathbf{s}_0, \dots, \mathbf{s}_L$, each of dimension Ω (number of states of DFA), together constitutes the secret sharing randomness \mathbf{s} of length $L\Omega$. The locality of shares in \mathbf{s} is inherited from the locality of DFA computation, and the subtraction expresses the state transition. Evaluation is a telescoping sum that chains the state transitions.

In (again, say CP-) ABE, we set \mathbf{s} to be LWE samples jointly generated by the keys (size dependent on L) and the ciphertexts (size dependent on Ω), $\mathbf{s}_\ell = \mathbf{S}\mathbf{r}_\ell$ for \mathbf{r}_ℓ 's in each key and \mathbf{S} in each ciphertext. Continuing with the high level form of the most important shares, we set

$$\begin{aligned} \text{sk}_{\mathbf{x}} &: \{ \text{isk}((1 - \mathbf{x}[\ell])\mathbf{r}_\ell, \mathbf{x}[\ell]\mathbf{r}_\ell, \mathbf{r}_{\ell-1}) \}_{\ell \in [L]}, \\ \text{ct}_{\text{DFA}} &: \text{ict}(\mathbf{\Gamma}_0^\top \mathbf{S}, \mathbf{\Gamma}_1^\top \mathbf{S}, -\mathbf{S}). \end{aligned}$$

Upon decryption, they compute all the shares, which reveal the secret approximately if the computation is accepting. For security, we first argue that $\mathbf{S}\mathbf{r}$'s are pseudoran-

dom, hence the shares are indistinguishable to created using fresh randomness, then we invoke the (information-theoretic) security of DFA secret sharing, and conclude ABE security from evasive IPFE security and pseudorandomness of the shares.

Instantiation of Evasive IPFE. Our candidate evasive IPFE schemes draw inspiration heavily from known IPFE constructions, and we proceed in three steps.

Basic Scheme. The basic scheme without identities nor structured noises is reminiscent to the schemes due to [ABDP15, ALS16]. The master public key consists of two matrices \mathbf{B}, \mathbf{A} , the master secret key is a trapdoor of \mathbf{B} , and

$$\text{isk}(\mathbf{v}) = \mathbf{B}^{-1}(\mathbf{A}\mathbf{G}^{-1}(\mathbf{v})), \quad \text{ict}(\mathbf{u}) = (\mathbf{d}^\top \mathbf{B}, \mathbf{d}^\top \mathbf{A} + \mathbf{u}^\top \mathbf{G}) = (\mathbf{c}_1^\top, \mathbf{c}_2^\top),$$

where $\mathbf{B}^{-1}(\mathbf{p})$ is a low-norm vector \mathbf{k} satisfying $\mathbf{B}\mathbf{k} = \mathbf{p}$, which can be sampled using the master secret key. For correctness, observe that $\mathbf{c}_2^\top \mathbf{G}^{-1}(\mathbf{v}) - \mathbf{c}_1^\top \mathbf{k}$ approximates $\mathbf{u}^\top \mathbf{v}$. We study the security of this simple scheme extensively in Section 3. Most importantly, its security can be reduced to a variant of evasive LWE if the plaintext vectors \mathbf{u} are uniformly random over a publicly known subspace.

Identity-Based Scheme. We borrow techniques from pairing to generically achieve identity-based schemes. In pairing-based IPFE, identity binding can be achieved using a single-identity scheme (denoted with primes) by adding a blinding factor in the decryption result. Let φ, ψ be random, then

$$\begin{aligned} \text{isk}_{\text{id}'}(\mathbf{v}) : & \quad \text{isk}'(\mathbf{v}, \psi, \text{id}' \cdot \psi), \\ \text{ict}_{\text{id}}(\mathbf{u}) : & \quad \text{ict}'(\mathbf{u}, \text{id} \cdot \varphi, -\varphi). \end{aligned}$$

The decryption outcome is $(\mathbf{u}^\top \mathbf{v} + (\text{id} - \text{id}')\varphi\psi)$. When the identities match, the result is the desired. When they do not match, the result is blinded by $\varphi\psi$, which is pseudorandom in the exponent (e.g., by DDH). The similar idea can be applied to evasive IPFE, except the blinding factor is $\varphi^\top \psi$, which is pseudorandom by LWE. To see security, observe that non-matching identities, due to $\varphi^\top \psi$, yield pseudorandom inner products concerned by the underlying scheme, whereas matching inner products are already pseudorandom by premise.

Scheme with Structured Noises. The structured noise $\mathbf{g}^\top \otimes \mathbf{e}^\top$ is simply \mathbf{e} multiplied by powers of two. To obtain them, we change the format of $\mathbf{d}^\top \mathbf{B}$ in ict into

$$\mathbf{d}^\top \mathbf{B}_0 + 2^0 \mathbf{e}^\top + (\mathbf{e}'_0)^\top, \dots, \mathbf{d}^\top \mathbf{B}_{K-1} + 2^{K-1} \mathbf{e}^\top + (\mathbf{e}'_{K-1})^\top,$$

where $\mathbf{e}, \mathbf{e}'_0, \dots, \mathbf{e}'_{K-1}$ are Gaussian noises. For $\mathbf{u}^\top \mathbf{V} + \mathbf{g}^\top \otimes \mathbf{e}''$, we would like to transform the previous blocks into

$$\mathbf{d}^\top \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_0) + 2^0 \mathbf{e}'' + \mathbf{e}'''_0, \dots, \mathbf{d}^\top \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_{K-1}) + 2^{K-1} \mathbf{e}'' + \mathbf{e}'''_{K-1},$$

so that they will cancel $\mathbf{d}^\top \mathbf{A}\mathbf{G}^{-1}(\mathbf{V})$ from $(\mathbf{d}^\top \mathbf{A} + \mathbf{u}^\top \mathbf{G})\mathbf{G}^{-1}(\mathbf{V})$ while attaching $\mathbf{g}^\top \otimes \mathbf{e}''$ to it. This can be done by finding low-norm \mathbf{K} such that $\mathbf{B}_k \mathbf{K} = \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_k)$ for all k , or equivalently $\tilde{\mathbf{B}}\mathbf{K} = \tilde{\mathbf{P}}$ for

$$\tilde{\mathbf{B}} = \begin{pmatrix} \mathbf{B}_0 \\ \vdots \\ \mathbf{B}_{K-1} \end{pmatrix}, \quad \tilde{\mathbf{P}} = \begin{pmatrix} \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_1) \\ \vdots \\ \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_k) \end{pmatrix}.$$

We sample $\tilde{\mathbf{B}}$ with trapdoor to be able to sample \mathbf{K} . Some more care is required to support decryption with either structured noises or not. Relying on a modified version of evasive LWE, termed *evasive learning with structured errors* assumption, our candidate is again secure for \mathbf{u} random over public subspace, sufficient for our application of succinct ABE.

2 Preliminaries

We denote the (computational) security parameter by λ , omitted for brevity except in definitions. Efficient algorithms are probabilistic polynomial-time Turing machines. So are efficient adversaries, but in addition they might also be given $\text{poly}(\lambda)$ -bit advice dependent on λ .¹ The statistical security parameter is $\kappa \in \omega(\log \lambda) \cap \lambda^{O(1)}$.

We use boldfaced lower-case letters for vectors, and boldfaced upper-case letters for matrices. They are always indexed using brackets and never using subscripts, so \mathbf{v}_i is the i^{th} vector among a series of related vectors, and $\mathbf{U}[i, j]$ is the (i, j) -entry of \mathbf{U} . The $n \times n$ identity matrix is \mathbf{I}_n (or simply \mathbf{I} when the shape is not of importance). When the dimension is clear from context, \mathbf{e}_i denotes the i^{th} standard basis vector, e.g., $\mathbf{I} = (\mathbf{e}_1, \dots)$. We write $\mathbf{0}_{n \times m}$ for the $n \times m$ zero matrix and $\mathbf{0}_n = \mathbf{0}_{n \times 1}$, with dimensions possibly omitted. We consider the infinity norm and its operator norm:

$$\|\mathbf{v}\| = \max_i |\mathbf{v}[i]|, \quad \|\mathbf{U}\| = \max_i \sum_j |\mathbf{U}[i, j]|.$$

We strictly follow the convention of vectors being columns. If $\mathbf{v} \in \mathbb{Z}^z$ is a vector, $\|\mathbf{v}^\top\|$ is an operator norm and $\|\mathbf{v}^\top\| \leq z\|\mathbf{v}\|$. For two matrices \mathbf{A}, \mathbf{B} of shapes $n_1 \times m_1$ and $n_2 \times m_2$, their Kronecker product is an $n_1 n_2 \times m_1 m_2$ matrix,

$$\mathbf{A} \otimes \mathbf{B} = \begin{pmatrix} \mathbf{A}[1, 1]\mathbf{B} & \cdots & \mathbf{A}[1, m_1]\mathbf{B} \\ \vdots & \ddots & \vdots \\ \mathbf{A}[n_1, 1]\mathbf{B} & \cdots & \mathbf{A}[n_1, m_1]\mathbf{B} \end{pmatrix}.$$

A useful property is $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$ whenever all multiplications are compatible. The by-column flattening operator $\text{col}(\cdot)$ concatenates all the columns of its input matrix. A convenient fact is $\text{col}(\mathbf{ABC}) = (\mathbf{C}^\top \otimes \mathbf{A})\text{col}(\mathbf{B})$ whenever \mathbf{ABC} is compatible. We write \star for the Kleene star, extend its usage to rows and columns of matrices, and overload it also for unnamed suitable value. For example, $\mathbf{E} \in \mathbb{Z}^{2 \times \star}$ could mean that \mathbf{E} is a 2-row integer matrix of some suitable number of columns so that its operations with other vectors and matrices are compatible, which should be clear from the context.

For integers L, B , we write $[L, B]$ for $\{z \in \mathbb{Z} \mid L \leq z \leq B\}$. For $n \geq 0$, we write $[n]$ for $[1, n]$. For natural number $q \geq 2$, we denote by \mathbb{Z}_q the integers modulo q . Matrices over \mathbb{Z} are naturally mapped to those over \mathbb{Z}_q , and this mapping is implicit so that they can be freely mixed for various operations. When $z \in \mathbb{Z}_q$, saying $z \in [L, B]$ means that a representative of z is in $[L, B]$.

Symbols. Table 1 explains select single-letter symbols used in this work.

¹The reductions in this work do not “increase non-uniformity”. Samplers, considered in many security definitions in this work, are a part of the adversary.

Table 1. Select single-letter symbols.

context	symbol	meaning
generic	λ, κ $\beta, \delta, \delta, \Delta$	computational/statistical security parameter challenge bit, random scalar/vector/matrix
ABE	P, X, Y x, y, μ J, j \mathbf{r}, \mathbf{s}	predicate (family), key-tied set, ciphertext-tied set key-tied value, ciphertext-tied value, message key count, index rerandomizer, garbling randomness
lattices	n, m, q, ρ \mathbf{g}, \mathbf{G} $\mathbf{p}, \mathbf{P}, \mathbf{k}, \mathbf{K}$ $\mathbf{A}, \mathbf{B}, \tau$ \mathcal{D}, σ $e, \mathbf{e}, \mathbf{E}, B$ \mathbf{d}, r K, k	dimension, dimension, modulus, hardness exponent gadget vector/matrix image vector/matrix, preimage vector/matrix matrix, matrix with trapdoor, trapdoor discrete Gaussian distribution, width noise scalar/vector/matrix, noise bound (evasive) LWE secret, sampler randomness structured noise length/index
IPFE	$\mathcal{I}, q, \mathbf{Z}, z$ B, σ $\mathbf{v}, \mathbf{u}^\top$ J, I, j, i $\boldsymbol{\psi}, \boldsymbol{\varphi}^\top$	identity space (family), modulus, dimension, index error bound, noise width key/plaintext vector key/ciphertext count, key/ciphertext index key/ciphertext randomness for identity binding
garbling	$F, f, \mathbf{x}, L, \ell$ q, n, \mathbf{s} $w, \mathbf{w}, \mathbf{w}, \mathbf{W}$ \mathbf{t}, \mathbf{T} \mathbf{R} B, \mathbf{e}	function family, function, input, length, index modulus, randomness dimension, garbling randomness secret scalar/vector, label (function) vector/matrix garbled table vector/matrix reusable information noise bound, garbling noise
circuits	C, d x, M \mathbf{x}, L, ℓ	circuit, depth (bound) wire value, wire value bound input, length, index
DFA	$\Gamma, \mathcal{Q}, q, \mathbf{\Gamma}$ ι_1, ξ, \bar{L} \mathbf{x}, L, ℓ	DFA, number of states, state, transition matrix initial/rejection state vector, input length bound input, length, index

2.1 Attribute-Based Encryption

We consider *promise* ABE with respect to *partial* predicates. Correctness holds when the predicate outputs 1; security holds when it outputs 0; neither is guaranteed if the output is \perp . This feature is used to exclude arithmetic computations with out-of-bound wire values, and it enables accurate specification of bounded ABE. Promise variants of ABE are already considered in [JLL23].

Definition 1 (ABE [GPSW06]). Let $P = \{P_{\lambda, \text{param}}\}_{\lambda \in \mathbb{N}, \text{param} \in \text{Params}_\lambda}$ be a predicate family, where each $P_{\lambda, \text{param}}$ is a function $X_{\lambda, \text{param}} \times Y_{\lambda, \text{param}} \rightarrow \{0, 1, \perp\}$ and $\text{Params} = \{\text{Params}_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of predicate description sets. An *attribute-based encryption scheme* for P consists of four efficient algorithms.

- $\text{Setup}(1^\lambda, \text{param})$ takes the predicate description $\text{param} \in \text{Params}_\lambda$ as input, and outputs a pair of master public/secret keys (mpk, msk) .
- $\text{KeyGen}(1^\lambda, \text{msk}, x)$ takes as input msk and some $x \in X_{\lambda, \text{param}}$. It outputs a secret key sk for x .
- $\text{Enc}(1^\lambda, \text{mpk}, y, \mu)$ takes as input mpk , some $y \in Y_{\lambda, \text{param}}$, and a single-bit message $\mu \in \{0, 1\}$. It outputs a ciphertext ct of μ tied to y .
- $\text{Dec}(1^\lambda, \text{mpk}, x, \text{sk}, y, \text{ct})$ takes as input $\text{mpk}, x, \text{sk}, y, \text{ct}$. It is supposed to output μ if $P_{\lambda, \text{param}}(x, y) = 1$.

The scheme must be *correct*, i.e., for all $\lambda \in \mathbb{N}$, $\text{param} \in \text{Params}_\lambda$, $x \in X_{\lambda, \text{param}}$, $y \in Y_{\lambda, \text{param}}$, $\mu \in \{0, 1\}$ such that $P_{\lambda, \text{param}}(x, y) = 1$, it holds that

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, \text{param}) \\ \text{sk} \stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda, \text{msk}, x) : \text{Dec}(1^\lambda, \text{mpk}, x, \text{sk}, y, \text{ct}) = \mu \\ \text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{mpk}, y, \mu) \end{array} \right] = 1.$$

Security. We consider *very selective* [AWY20] (also known as *static* [RW15]) security in this work. Before the ABE is set up, the adversary chooses all the x 's and the y for which the secret keys and the challenge ciphertext are to be generated.

Definition 2 (ABE security [AWY20]). An ABE scheme (Definition 1) is *very selectively secure* if $\text{Exp}_{\text{ABE}}^0 \approx \text{Exp}_{\text{ABE}}^1$, where $\text{Exp}_{\text{ABE}}^\beta(1^\lambda, \mathcal{A})$ proceeds as follows.

- **Challenge.** Launch $\mathcal{A}(1^\lambda)$ and receive from it

$$\text{param} \in \text{Params}_\lambda, \quad \{x_j\}_{j \in [J]} \quad (x_j \in X_{\lambda, \text{param}} \text{ for all } j \in [J]), \quad y^* \in Y_{\lambda, \text{param}}.$$

- **Setup.** Run

$$\begin{aligned} (\text{mpk}, \text{msk}) &\stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, \text{param}), \\ \text{sk}_j &\stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda, \text{msk}, x_j) \text{ for all } j \in [J], \\ \text{ct}^* &\stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{mpk}, y^*, \beta), \end{aligned}$$

and send $\text{mpk}, \{\text{sk}_j\}_{j \in [J]}, \text{ct}^*$ to \mathcal{A} .

- **Guess.** The adversary \mathcal{A} outputs $\beta' \in \{0, 1\}$. The output of the experiment is β' if $P_{\lambda, \text{param}}(x_j, y^*) = 0$ for all $j \in [J]$. Otherwise, the output is set to 0.

2.2 Lattices

Let $n, m \geq 1$ and $q \geq 2$ be integers such that $\log_2 q \leq m/n \in \mathbb{Z}$. Let

$$\mathbf{g} = (2^0, \dots, 2^{m/n-1})^\top \quad \text{and} \quad \mathbf{G} = \mathbf{I}_n \otimes \mathbf{g}^\top$$

be the gadget vector and the gadget matrix. Given $\mathbf{p} \in \mathbb{Z}_q^n$, we denote by $\mathbf{G}^{-1}(\mathbf{p})$ the vector $\mathbf{k} \in \{0, 1\}^m$, where $\mathbf{k}[(i-1) \cdot m/n + 1], \dots, \mathbf{k}[i \cdot m/n]$ are the bits of (the non-negative representative less than q of) $\mathbf{p}[i]$, low to high. It holds that $\mathbf{G}\mathbf{G}^{-1}(\mathbf{p}) = \mathbf{p}$. Given $K \in \mathbb{N}$, we let (note the different order in the Kronecker product)

$$\tilde{\mathbf{g}} = (2^0, \dots, 2^{K-1})^\top \quad \text{and} \quad \tilde{\mathbf{G}} = \tilde{\mathbf{g}}^\top \otimes \mathbf{I}.$$

When $2^K \geq q$, the vector $\tilde{\mathbf{G}}^{-1}(\mathbf{p})$ is again the bit decomposition of $\mathbf{p} \in \mathbb{Z}_q^*$, arranged so that $\tilde{\mathbf{G}}\tilde{\mathbf{G}}^{-1}(\mathbf{p}) = \mathbf{p}$. The notations $\mathbf{G}^{-1}(\cdot), \tilde{\mathbf{G}}^{-1}(\cdot)$ extend to matrices column-wise.

Given $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{p} \in \mathbb{Z}_q^n$ such that $\mathbf{B}\mathbf{k} = \mathbf{p}$ has a solution $\mathbf{k}^* \in \mathbb{Z}^m$, we write the lattice coset

$$\Lambda_{\mathbf{p}}^\perp(\mathbf{B}) = \{\mathbf{k} \in \mathbb{Z}^m \mid \mathbf{B}\mathbf{k} = \mathbf{p}\} = \mathbf{k}^* + \Lambda_{\mathbf{0}}^\perp(\mathbf{B}).$$

Let Λ' be any lattice coset and $\sigma \geq 0$, we denote by $\mathcal{D}_{\Lambda', \sigma}$ the discrete Gaussian distribution [MP11] over Λ' with width σ . Since we insist on perfect correctness, it is necessary to truncate unbounded noises in algorithms:

Definition 3 (truncated $\mathcal{D}_{\mathbb{Z}, \sigma}$). Let $B \geq 0$, the distribution $\mathcal{D}_{\mathbb{Z}, \sigma, \leq B}$ is sampled by first sampling $x \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}, \sigma}$, then returning x if $|x| \leq B$, and 0 otherwise.

Lemma 1 (truncation of $\mathcal{D}_{\mathbb{Z}, \sigma}$). For all $\sigma \geq 1$ and $\kappa \geq 2$,

$$\Pr[x \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}, \sigma} : |x| > \sigma\sqrt{\kappa}] \leq 2^{-\kappa},$$

so $\mathcal{D}_{\mathbb{Z}, \sigma, \leq \sigma\sqrt{\kappa}}$ is $2^{-\kappa}$ -close to $\mathcal{D}_{\mathbb{Z}, \sigma}$.

We will also need noise flooding with discrete Gaussian:

Lemma 2 (noise flooding using $\mathcal{D}_{\mathbb{Z}, \sigma}$). For all $y \in \mathbb{Z}$ and $\sigma \geq 2^{\kappa+6}y$, it holds that $(y + \mathcal{D}_{\mathbb{Z}, \sigma})$ is $2^{-\kappa}$ -close to $\mathcal{D}_{\mathbb{Z}, \sigma}$.

Trapdoor Generation and Gaussian Sampling. We need the following lemma:

Lemma 3 ([MP12; Theorem 2]). There are efficient algorithms TrapGen and SampleD, and functions $m_0 \in \Theta(n \log q)$ and $\sigma_0 \in \omega(\sqrt{m \log m}) \cap \mathcal{O}(m)$ satisfying these conditions.

- TrapGen($1^n, 1^m, q$) takes as input $n \geq 1, q \geq 2$, and $m \geq m_0(n, q)$. It outputs (\mathbf{B}, τ) such that $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ and \mathbf{B} is $\text{negl}(n)$ -close to uniform over $\mathbb{Z}_q^{n \times m}$.
- SampleD($\mathbf{B}, \tau, \mathbf{p}, \sigma$) takes as input \mathbf{B}, τ from TrapGen, some $\mathbf{p} \in \mathbb{Z}_q^n$, and $\sigma \geq \sigma_0(n, m)$. It outputs $\mathbf{k} \in \mathbb{Z}^m$ such that $\mathbf{B}\mathbf{k} = \mathbf{p}$, $\|\mathbf{k}\| \leq \sigma\sqrt{m}$, and \mathbf{k} is $\text{negl}(n)$ -close to $\mathcal{D}_{\Lambda_{\mathbf{p}}^\perp(\mathbf{B}), \sigma}$.

Batch Notation. It is convenient to extend SampleD to process multiple \mathbf{p} 's in one shot. Let $\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_{m'})$ be a matrix or a batch of vectors, then $\text{SampleD}(\mathbf{B}, \tau, \mathbf{P}, \sigma)$ is

$$\mathbf{K} \stackrel{\$}{\leftarrow} (\text{SampleD}(\mathbf{B}, \tau, \mathbf{p}_1, \sigma), \dots, \text{SampleD}(\mathbf{B}, \tau, \mathbf{p}_{m'}, \sigma)),$$

with fresh randomness for each call to SampleD on the right-hand side. We also write $\mathbf{B}^{-1}(\mathbf{P})$ for an output of $\text{SampleD}(\mathbf{B}, \tau, \mathbf{P}, \sigma)$.

Learning with Errors Assumption. We rely on the LWE assumption.

Assumption 1 (LWE [Reg05]). Let $n, m \leq \text{poly}(\lambda)$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, and

$$\mathbf{A} \stackrel{\$}{\leftarrow} \mathbb{Z}_{q(\lambda)}^{n(\lambda) \times m(\lambda)}, \quad \mathbf{s} \stackrel{\$}{\leftarrow} \mathbb{Z}_{q(\lambda)}^{n(\lambda)}, \quad \mathbf{e} \stackrel{\$}{\leftarrow} \mathbb{Z}_{q(\lambda)}^{m(\lambda)}, \quad \boldsymbol{\delta} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma(\lambda)}.$$

The *LWE assumption* $\text{LWE}_{n,m,q,\sigma}$ states that

$$\{(1^\lambda, \mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)\}_{\lambda \in \mathbb{N}} \approx \{(1^\lambda, \mathbf{A}, \boldsymbol{\delta}^\top)\}_{\lambda \in \mathbb{N}}.$$

Assumption 2 (flipped LWE). Let $N, m \leq \text{poly}(\lambda)$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, and

$$\mathbf{A} \stackrel{\$}{\leftarrow} \{0, 1\}^{N(\lambda) \times m(\lambda)}, \quad \mathbf{s} \stackrel{\$}{\leftarrow} \mathbb{Z}_{q(\lambda)}^{N(\lambda)}, \quad \mathbf{e} \stackrel{\$}{\leftarrow} \mathbb{Z}_{q(\lambda)}^{m(\lambda)}, \quad \boldsymbol{\delta} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma(\lambda)}.$$

The *flipped LWE assumption* $\text{FlipLWE}_{N,m,q,\sigma}$ states that

$$\{(1^\lambda, \mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)\}_{\lambda \in \mathbb{N}} \approx \{(1^\lambda, \mathbf{A}, \boldsymbol{\delta}^\top)\}_{\lambda \in \mathbb{N}}.$$

Lemma 4 (flipped LWE). *Suppose $\text{LWE}_{n,m,q,\sigma}$ holds and $N \leq \text{poly}(\lambda)$ satisfies*

$$0 \leq N/n - \log_2 q \in \omega(\log \lambda),$$

then $\text{FlipLWE}_{N,m,q,\sigma}$ holds.

Parameters. In this work, we rely on the LWE assumption with subexponential modulus-to-noise ratio, i.e., $n = \lambda^{\Theta(1)}$, $m \leq \text{poly}(n)$, $q \leq 2^{\text{poly}(n)}$, $\frac{q}{\sigma\sqrt{n}} = \Theta(2^{n^\rho})$, where $0 < \rho < 1$ is a constant. Strictly speaking, those parameters should be chosen by the adversary subject to those constraints — most algorithms require a prime q , it is currently unknown [TCH12] how to pick $\text{poly}(\lambda)$ -bit primes in deterministic polynomial time, and the adversary's choice of scheme parameters affect the ranges of LWE parameters.

2.3 Evasive LWE Assumption

Conceptually, we consider two knowledge assumptions — evasive learning with errors (evasive LWE) and evasive learning with *structured* errors (not abbreviated). The former is a special case of the latter, and we formulate the definitions as such.

Assumption 3 (evasive learning with structured errors). Let $\mathcal{S}(1^\lambda; r)$ be an algorithm that, given randomness r , outputs

$$1^n, 1^m, 1^K, q, 1^{n'}, 1^J, 1^I, 1^{m'}, \quad \sigma_{-1}, \quad \sigma_{\text{post}} \geq \sigma_{\text{pre}} \geq 0, \\ \tilde{\mathbf{P}} \in \mathbb{Z}_q^{n(K+1) \times J}, \quad \{\mathbf{A}_{1,i}, \mathbf{A}_{0,i}\}_{i \in [I]} \quad (\mathbf{A}_{1,i} \in \mathbb{Z}_q^{n \times m'}, \mathbf{A}_{0,i} \in \mathbb{Z}_q^{n' \times m'}),$$

where $m \geq m_0(n(K+1), q)$ and $\sigma_{-1} \geq \sigma_0(n(K+1), m)$ are constrained by Lemma 3. Suppose

$$(\tilde{\mathbf{B}}, \tau) \stackrel{\$}{\leftarrow} \text{TrapGen}(1^{n(K+1)}, 1^m, q), \quad \mathbf{K} \stackrel{\$}{\leftarrow} \text{SampleD}(\tilde{\mathbf{B}}, \tau, \tilde{\mathbf{P}}, \sigma_{-1}),$$

$$\begin{aligned}
\mathbf{d}_0 &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n'}; & \mathbf{e}_{\text{pre},i,0} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{pre}}}^{m'}, & \mathbf{e}_{\text{post},i,0} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{post}}}^{m'}, & \boldsymbol{\delta}_{i,0} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{m'}, \\
\mathbf{d}_i &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^n, & \mathbf{e}_{\text{pre},i,1} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{pre}}}^{m(K+1)}, & \mathbf{e}_{\text{post},i,1} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{post}}}^{m(K+1)}, & \boldsymbol{\delta}_{i,1} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{m(K+1)}, \\
& & \mathbf{e}_{\text{pre},i,2} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{pre}}}^{J(K+1)}, & \mathbf{e}_{\text{pre},i,4} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{pre}}}^J, & \boldsymbol{\delta}_{i,2} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{J(K+1)}, \\
& & \mathbf{e}_{\text{pre},i,3} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{pre}}}^m, & \mathbf{e}_{\text{post},i,3} &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z},\sigma_{\text{post}}}^m, & & \text{for all } i \in [I].
\end{aligned}$$

Let $\widetilde{\mathbf{g}}^\top = (2^0, 2^1, \dots, 2^{K-1})^\top$ and

$$\begin{aligned}
\widetilde{\mathbf{B}} &= (\mathbf{B}_{-1}^\top, \mathbf{B}_0^\top, \dots, \mathbf{B}_{K-1}^\top)^\top, & \mathbf{B} &= (\mathbf{B}_{-1}, \mathbf{B}_0, \dots, \mathbf{B}_{K-1}), \\
\widetilde{\mathbf{P}} &= (\mathbf{P}_{-1}^\top, \mathbf{P}_0^\top, \dots, \mathbf{P}_{K-1}^\top)^\top, & \mathbf{P} &= (\mathbf{P}_{-1}, \mathbf{P}_0, \dots, \mathbf{P}_{K-1}),
\end{aligned}$$

where $\mathbf{B}_k \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{P}_k \in \mathbb{Z}_q^{n \times J}$. The precondition $\text{evLWE}_{\text{pre}}^S$ is

$$\left\{ \left(\begin{array}{c} 1^\lambda, \left\{ \mathbf{d}_i^\top \mathbf{P} + (\mathbf{0}_{1 \times J}, \widetilde{\mathbf{g}}^\top \otimes \mathbf{e}_{\text{pre},i,4}^\top + \mathbf{e}_{\text{pre},i,2}^\top) \right\} \\ r, \left\{ \mathbf{d}_i^\top \mathbf{B} + (\mathbf{0}_{1 \times m}, \widetilde{\mathbf{g}}^\top \otimes \mathbf{e}_{\text{pre},i,3}^\top + \mathbf{e}_{\text{pre},i,1}^\top) \right\} \\ \mathbf{B}, \left\{ \mathbf{d}_i^\top \mathbf{A}_{1,i} + \mathbf{d}_0^\top \mathbf{A}_{0,i} + \mathbf{e}_{\text{pre},i,0}^\top \right\}_{i \in [I]} \end{array} \right) \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \left(\begin{array}{c} 1^\lambda, \left\{ \boldsymbol{\delta}_{i,2}^\top \right\} \\ r, \left\{ \boldsymbol{\delta}_{i,1}^\top \right\} \\ \mathbf{B}, \left\{ \boldsymbol{\delta}_{i,0}^\top \right\}_{i \in [I]} \end{array} \right) \right\}_{\lambda \in \mathbb{N}}.$$

The postcondition $\text{evLWE}_{\text{post}}^S$ is

$$\left\{ \left(\begin{array}{c} 1^\lambda, r, \left\{ \mathbf{d}_i^\top \mathbf{B} + (\mathbf{0}_{1 \times m}, \widetilde{\mathbf{g}}^\top \otimes \mathbf{e}_{\text{post},i,3}^\top + \mathbf{e}_{\text{post},i,1}^\top) \right\} \\ \mathbf{B}, \mathbf{K}, \left\{ \mathbf{d}_i^\top \mathbf{A}_{1,i} + \mathbf{d}_0^\top \mathbf{A}_{0,i} + \mathbf{e}_{\text{post},i,0}^\top \right\}_{i \in [I]} \end{array} \right) \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \left(\begin{array}{c} 1^\lambda, r, \left\{ \boldsymbol{\delta}_{i,1}^\top \right\} \\ \mathbf{B}, \mathbf{K}, \left\{ \boldsymbol{\delta}_{i,0}^\top \right\}_{i \in [I]} \end{array} \right) \right\}_{\lambda \in \mathbb{N}}.$$

The *evasive learning with structured errors assumption* states that $\text{evLWE}_{\text{pre}}^S$ implies $\text{evLWE}_{\text{post}}^S$ for all efficient S .

Assumption 4 (evasive LWE). A sampler \mathcal{S} (Assumption 3) does not use structured noises if $\Pr[K=0]=1$, for which $\text{evLWE}_{\text{pre}}^S$ (resp. $\text{evLWE}_{\text{post}}^S$) is defined to be $\text{evLWE}_{\text{pre}}^S$ (resp. $\text{evLWE}_{\text{post}}^S$). The *evasive LWE assumption* states that $\text{evLWE}_{\text{pre}}^S$ implies $\text{evLWE}_{\text{post}}^S$ for all efficient S not using structured noises.

Note that for evasive LWE, $\widetilde{\mathbf{B}} = \mathbf{B} = \mathbf{B}_{-1}$ and $\widetilde{\mathbf{P}} = \mathbf{P} = \mathbf{P}_{-1}$.

Remark 1 (public-coin samplers). As in [Wee22, WWW22], our evasive assumptions are only for public-coin samplers, which we enforce by providing the sampler randomness to the distinguisher. It is weaker than the versions [Tsa22, VWW22] used for witness encryption, which considers private-coin samplers and the distinguisher, instead of the sampler randomness, gets the auxiliary information, an additional output produced by the sampler. The public-coin assumption avoids obfuscation-based counterexamples, where \mathbf{P} is sampled with a trapdoor and the auxiliary information contains an obfuscated program with the trapdoor hardwired.

Remark 2 (comparison with [Wee22, WWW22]). As suggested in [Wee22], the noise magnitude in the postcondition can be made larger than that in the precondition for a more conservative assumption. We can implement it by requiring the assumptions to hold only for samplers with a gap between σ_{post} and σ_{pre} . We additionally allow a worse Gaussian preimage \mathbf{K} . Other than the magnitudes of samples, our version of evasive LWE is stronger than the version in [Wee22], and is similar to but incomparable with that in [WWW22].

While appearing much more complex, our evasive learning with structured errors assumption follows the same rationale as all existing versions of public-coin evasive LWE. In evasive LWE (without structured noises), the basic idea is that given

$$\mathbf{K} \stackrel{\$}{\leftarrow} \mathbf{B}^{-1}(\mathbf{P}), \quad \mathbf{d}^\top \mathbf{B} + \mathbf{e}_1^\top, \quad \mathbf{d}^\top \mathbf{A} + \mathbf{e}_0^\top,$$

the trapdoor preimages can only be meaningfully used to multiply the corresponding LWE samples, i.e., $(\mathbf{d}^\top \mathbf{B} + \mathbf{e}_1^\top) \mathbf{K} = \mathbf{d}^\top \mathbf{P} + \mathbf{e}_1^\top \mathbf{K}$, and that moreover, the correlation between \mathbf{e}_1 and $\mathbf{e}_1^\top \mathbf{K}$ is not useful to the attacker if the samples have no apparent relations when $\mathbf{e}_1^\top \mathbf{K}$ is replaced by fresh \mathbf{e}_2^\top . Therefore, the precondition requires that

$$\mathbf{d}^\top \mathbf{P} + \mathbf{e}_2^\top, \quad \mathbf{d}^\top \mathbf{B} + \mathbf{e}_1^\top, \quad \mathbf{d}^\top \mathbf{A} + \mathbf{e}_0^\top$$

be pseudorandom, and evasive LWE asserts that $(\mathbf{d}^\top \mathbf{B} + \mathbf{e}_1^\top, \mathbf{d}^\top \mathbf{A} + \mathbf{e}_0^\top)$ is pseudorandom in the presence of \mathbf{K} .

For evasive learning with structured errors, suppose we are given

$$\mathbf{K}, \quad \{\mathbf{d}^\top \mathbf{B}_k + 2^k \mathbf{e}_3^\top + \mathbf{e}_{1,k}^\top\}_{0 \leq k \leq K-1}, \quad \mathbf{d}^\top \mathbf{A} + \mathbf{e}_0^\top,$$

where $\mathbf{B}_k \mathbf{K} = \mathbf{P}_k$ for all k . Again, the trapdoor preimages can only be meaningfully used to multiply the corresponding LWE samples to obtain

$$(\mathbf{d}^\top \mathbf{B}_k + 2^k \mathbf{e}_3^\top + \mathbf{e}_{1,k}^\top) \mathbf{K} = \mathbf{d}^\top \mathbf{P}_k + 2^k \mathbf{e}_3^\top \mathbf{K} + \mathbf{e}_{1,k}^\top \mathbf{K}.$$

Following the same rationale of “correlation among noises not helpful if samples with fresh noises are pseudorandom”, we replace $\{\mathbf{e}_{1,k}^\top \mathbf{K}\}_k, \mathbf{e}_3^\top \mathbf{K}$ by independent $\{\mathbf{e}_{2,k}^\top\}_k, \mathbf{e}_4^\top$. This gives rise to the conditions in our assumption. It remains to figure out how to sample \mathbf{K} simultaneously satisfying $\mathbf{B}_k \mathbf{K} = \mathbf{P}_k$ for all k . The equations can be rewritten as

$$\begin{pmatrix} \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_k \end{pmatrix} \mathbf{K} = \begin{pmatrix} \mathbf{P}_1 \\ \vdots \\ \mathbf{P}_k \end{pmatrix},$$

so it suffices to sample the vertically blocked \mathbf{B}_k 's with trapdoor. We further incorporate a block $(\mathbf{B}_{-1}$ and $\mathbf{P}_{-1})$ without structured noises and arrive at Assumption 3.

3 Evasive Inner-Product Functional Encryption

We define identity-based IPFE with structured noises and its security. A scheme might not support every combination of moduli, identities, and noise structures, which can be signaled by the algorithms aborting.

Definition 4 ((IB-)evIPFE(-w/sn)). Let $\mathcal{I} = \{\mathcal{I}_{\lambda,q,K,Z}\}_{\lambda,q,K,Z \in \mathbb{N}}$ be a family of sets. An (identity-based) evasive IPFE (with structured noises) scheme for \mathcal{I} consists of four efficient algorithms.

- $\text{Setup}(1^\lambda, q, 1^K, 1^Z)$ takes as input the modulus q , the noise parameter K , and the dimension Z . It outputs a pair $(\text{impk}, \text{imsk})$ of master public/secret keys.

- $\text{KeyGen}(1^\lambda, \text{imsk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g)$ takes as input imsk , an identity $\text{id} \in \mathcal{I}_{\lambda, q, K, Z}$, a key vector $\mathbf{v}_0 \in \mathbb{Z}_q^Z$, and a key matrix $\mathbf{V}_g \in \mathbb{Z}_q^{Z \times K}$. It outputs a secret key isk for $(\mathbf{v}_0, \mathbf{V}_g)$ under identity id .
- $\text{Enc}(1^\lambda, \text{impk}, \text{id}, \mathfrak{s}, \mathbf{u}^\top)$ takes as input impk , id , a noise structure $\mathfrak{s} \in \{\mathfrak{o}, \mathfrak{g}\}$, and a plaintext vector $\mathbf{u} \in \mathbb{Z}_q^Z$. It outputs a ciphertext ict of \mathbf{u} under id for noise structure \mathfrak{s} .
- $\text{Dec}(1^\lambda, \text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{s}, \text{ict})$ takes as input impk , id , $\mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{s}, \text{ict}$. If ict is for noise structure \mathfrak{s} and both isk and ict are under id , the decryption result is in \mathbb{Z}_q and is supposed to approximate $\mathbf{u}^\top \mathbf{v}_0$ (when $\mathfrak{s} = \mathfrak{o}$) or $\mathbf{u}^\top \mathbf{V}_g$ (when $\mathfrak{s} = \mathfrak{g}$).

Let B be a function mapping (λ, q, K, Z) to natural numbers. The scheme is B -correct for \mathfrak{o} if for all $\lambda, q, K, Z \in \mathbb{N}$, $\text{id} \in \mathcal{I}_{\lambda, q, K, Z}$, $\mathbf{V}_g \in \mathbb{Z}_q^{Z \times K}$, $\mathbf{v}_0, \mathbf{u} \in \mathbb{Z}_q^Z$, it holds that

$$\Pr \left[\begin{array}{l} (\text{impk}, \text{imsk}) \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, q, 1^K, 1^Z) \\ \text{isk} \stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda, \text{imsk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g) \\ \text{ict} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{impk}, \text{id}, \mathfrak{o}, \mathbf{u}^\top) \end{array} : \begin{array}{l} \text{Dec}(1^\lambda, \text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{g}, \text{ict}) \\ = \mathbf{u}^\top \mathbf{v}_0 + e_0 \text{ for some } \\ e_0 \in [-B(\lambda, q, K, Z), B(\lambda, q, K, Z)] \end{array} \right] = 1.$$

For B -correctness for \mathfrak{g} , the requirement is

$$\Pr \left[\begin{array}{l} \text{---} \text{ for } \text{impk}, \text{imsk}, \text{isk} \\ \text{ict} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{impk}, \text{id}, \mathfrak{g}, \mathbf{u}^\top) \end{array} : \begin{array}{l} \text{Dec}(1^\lambda, \text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{g}, \text{ict}) \\ = \mathbf{u}^\top \mathbf{V}_g + \tilde{\mathbf{g}}^\top \otimes e_{\mathfrak{g}\mathfrak{g}} + \mathbf{e}_{\mathfrak{g}\mathfrak{o}}^\top \text{ for some } \\ e_{\mathfrak{g}\mathfrak{g}}, \mathbf{e}_{\mathfrak{g}\mathfrak{o}} \in [-B(\lambda, q, K, Z), B(\lambda, q, K, Z)]^* \end{array} \right] = 1.$$

The scheme is B -correct if both are satisfied.

The correctness property is trivial when $\mathcal{I} = \emptyset$ or $B \geq q/2$. We are interested in large enough \mathcal{I} and reasonably small B , for useful choices of q, Z . Jumping ahead, Constructions 2 and 3 support exponentially large $|\mathcal{I}|$ and q/B , which suffice for our ABE schemes.

When referring to specific schemes, if $|\mathcal{I}|$ is small, we remove “identity-based” (IB), and if they only support $K = 0$, we remove “with structured noises” (w/sn).

Batch Notations. It is convenient to consider $\text{KeyGen}, \text{Enc}, \text{Dec}$ in batches. Let

$\mathbf{V}_0 = (\mathbf{v}_{1,0}, \dots, \mathbf{v}_{J,0}) \in \mathbb{Z}_q^{Z \times J}$, $\mathbf{V}_g = (\mathbf{V}_{1,g}, \dots, \mathbf{V}_{J,g}) \in \mathbb{Z}_q^{Z \times KJ}$, $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_I) \in \mathbb{Z}_q^{Z \times I}$ be matrices (batches of vectors/matrices) and $\text{id} \in \mathcal{I}$, $\mathfrak{s} \in \{\mathfrak{o}, \mathfrak{g}\}$. We write

$$\text{isk} \stackrel{\$}{\leftarrow} \text{KeyGen}(\text{imsk}, \text{id}, \mathbf{V}_0, \mathbf{V}_g) \quad \text{and} \quad \text{ict} \stackrel{\$}{\leftarrow} \text{Enc}(\text{impk}, \text{id}, \mathfrak{s}, \mathbf{U}^\top)$$

for $\text{isk} \leftarrow \{\text{isk}_j\}_{j \in [J]}$ and $\text{ict} \leftarrow \{\text{ict}_i\}_{i \in [I]}$, where

$$\begin{array}{ll} \text{isk}_j \stackrel{\$}{\leftarrow} \text{KeyGen}(\text{imsk}, \text{id}, \mathbf{v}_{j,0}, \mathbf{V}_{j,g}) & \text{for all } j \in [J], \\ \text{ict}_i \stackrel{\$}{\leftarrow} \text{Enc}(\text{impk}, \text{id}, \mathfrak{s}, \mathbf{u}_i^\top) & \text{for all } i \in [I], \end{array}$$

and $\text{Dec}(\text{impk}, \text{id}, \mathbf{V}_0, \mathbf{V}_g, \text{isk}, \mathfrak{s}, \text{ict})$ means

$$\left(\begin{array}{ccc} \text{Dec}(\text{impk}, \text{id}, \mathbf{v}_{1,0}, \mathbf{V}_{1,g}, \text{isk}_1, \mathfrak{s}, \text{ict}_1) & \cdots & \text{Dec}(\text{impk}, \text{id}, \mathbf{v}_{J,0}, \mathbf{V}_{J,g}, \text{isk}_J, \mathfrak{s}, \text{ict}_1) \\ \vdots & \ddots & \vdots \\ \text{Dec}(\text{impk}, \text{id}, \mathbf{v}_{1,0}, \mathbf{V}_{1,g}, \text{isk}_1, \mathfrak{s}, \text{ict}_I) & \cdots & \text{Dec}(\text{impk}, \text{id}, \mathbf{v}_{J,0}, \mathbf{V}_{J,g}, \text{isk}_J, \mathfrak{s}, \text{ict}_I) \end{array} \right),$$

which is approximately $\mathbf{U}^\top \mathbf{V}_\mathfrak{s}$. The mnemonic is that the decryption result is the plaintext matrix \mathbf{U}^\top multiplied by the key matrix $\mathbf{V}_\mathfrak{s}$.

Security. We consider *very selective* (also known as *static*) security with respect to a sampler producing pseudorandom structurally noisy outputs.

Definition 5 (IB-evIPFE-w/sn security). Let $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ be two algorithms with the following syntax.

- $\mathcal{S}_V(1^\lambda; r_{\text{pub}})$, given randomness r_{pub} , outputs

$$q, 1^K, 1^Z, \text{ID} \subseteq \mathcal{I}_{\lambda, q, K, Z}, \{1^{J_{\text{id}}}, 1^{I_{\text{id},0}}, 1^{I_{\text{id},g}}\}_{\text{id} \in \text{ID}}, \quad \sigma_{\text{pre}} \geq 0,$$

$$\{\mathbf{V}_{\text{id},0}, \mathbf{V}_{\text{id},g}\}_{\text{id} \in \text{ID}} \quad (\mathbf{V}_{\text{id},0} \in \mathbb{Z}_q^{Z \times J_{\text{id}}}, \mathbf{V}_{\text{id},g} \in \mathbb{Z}_q^{Z \times K J_{\text{id}}}).$$

- $\mathcal{S}_U(1^\lambda, r_{\text{pub}}; r_{\text{priv}})$, given r_{pub} and additional randomness r_{priv} , outputs

$$\{\mathbf{U}_{\text{id},s}^\top\}_{\text{id} \in \text{ID}, s \in \{0,g\}} \quad (\mathbf{U}_{\text{id},s} \in \mathbb{Z}_q^{Z \times I_{\text{id},s}}),$$

where $q, Z, \text{ID}, \{I_{\text{id},s}\}_{\text{id} \in \text{ID}, s \in \{0,g\}}$ agree with the output of $\mathcal{S}_V(1^\lambda; r_{\text{pub}})$.

Suppose $(\text{impk}, \text{imsk}) \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, q, 1^K, 1^Z)$ and for all $\text{id} \in \text{ID}, s \in \{0, g\}$,

$$\begin{aligned} \mathbf{E}_{\text{id},0} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{pre}}}^{J_{\text{id}} \times I_{\text{id},0}}, & \mathbf{E}_{\text{id},gg} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{pre}}}^{J_{\text{id}} \times I_{\text{id},g}}, & \mathbf{E}_{\text{id},g0} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{pre}}}^{K J_{\text{id}} \times I_{\text{id},g}}, \\ \Delta_{\text{pre}, \text{id},0} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{J_{\text{id}} \times I_{\text{id},0}}, & \Delta_{\text{pre}, \text{id},g} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{K J_{\text{id}} \times I_{\text{id},g}}, \\ & & \text{isk}_{\text{id}} &\stackrel{\$}{\leftarrow} \text{KeyGen}(1^\lambda, \text{imsk}, \text{id}, \mathbf{V}_{\text{id},0}, \mathbf{V}_{\text{id},g}), \\ & & \text{ict}_{\text{id},s}^* &\stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{impk}, \text{id}, s, \mathbf{U}_{\text{id},s}^\top), \\ \mathbf{U}_{\$, \text{id}, s} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{Z \times I_{\text{id},s}}, & \text{ict}_{\text{id},s}^\$ &\stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{impk}, \text{id}, s, \mathbf{U}_{\$, \text{id}, s}^\top). \end{aligned}$$

\mathcal{S} has *pseudorandom structurally noisy inner products*, denoted $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$, if

$$\left\{ \left(1^\lambda, \left\{ \mathbf{U}_{\text{id},0}^\top \mathbf{V}_{\text{id},0} + \mathbf{E}_{\text{id},0}^\top, \right\}_{\text{id} \in \text{ID}} \right) \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \left(1^\lambda, \left\{ \Delta_{\text{pre}, \text{id},0}^\top, \right\}_{\text{id} \in \text{ID}} \right) \right\}_{\lambda \in \mathbb{N}}.$$

The scheme is \mathcal{S} -secure, denoted $\text{IPFEsec}_{\text{post}}^{\mathcal{S}}$, if

$$\left\{ \left(1^\lambda, \left\{ \text{ict}_{\text{id},0}^*, \right\}_{\text{id} \in \text{ID}} \right) \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \left(1^\lambda, \left\{ \text{ict}_{\text{id},0}^\$, \right\}_{\text{id} \in \text{ID}} \right) \right\}_{\lambda \in \mathbb{N}}.$$

The scheme is $\sigma_{\text{pre},0}$ -secure if it is \mathcal{S} -secure for all efficient \mathcal{S} such that $\sigma_{\text{pre}} \leq \sigma_{\text{pre},0}$ and $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$ holds.

Remark 3 (public and private coins). It is necessary to make the sampling of \mathbf{U} private-coin. Since IPFE does not protect \mathbf{V} (the key vectors/matrices), the adversary is assumed to have full knowledge of \mathbf{V} . If \mathbf{U} were sampled with public coins, the distinguisher against the ciphertexts (in $\text{IPFEsec}_{\text{post}}^{\mathcal{S}}$) would know both \mathbf{V} and \mathbf{U} , thus $\mathbf{U}^\top \mathbf{V}$. In the first distribution (containing ciphertexts of \mathbf{U}), decryption will yield the correct output, whereas it is not the case for the second distribution (containing ciphertexts of random vectors).

We formulate \mathcal{S}_V as public-coin and \mathcal{S}_U as fully private-coin, to avoid obvious counterexamples (e.g., obfuscation-based).

3.1 Basic Construction — Single-Identity, No Structured Noises

We first present a simple candidate for one identity without structured noises. The scheme is reminiscent to many existing IPFE constructions [ABDP15, ALS16, Wee17, Agr19, AP20, LL20a, LL20b].

Ingredients of Construction 1. We rely on Lemma 3 for correctness. The security of the scheme is elaborated in Sections 3.2 and 3.3.

Construction 1 (evIPFE). Our scheme works as follows.

- $\text{Setup}(q, 1^K, 1^Z)$ takes as input q, K, Z . If $K \neq 0$, it sets $\mathcal{I}_{q,K,Z} = \emptyset$ and terminates (aborting case). Otherwise, the algorithm sets $\mathcal{I}_{q,K,Z} = \{1\}$, and *deterministically* picks n, m and $\sigma_{-1} \geq \sigma_0(n, m)$ appropriate for Lemma 3, as well as $\sigma_{\text{post}}, \kappa$. It samples and sets

$$\begin{aligned} (\mathbf{B}, \tau) &\stackrel{\$}{\leftarrow} \text{TrapGen}(1^n, 1^m, q), & \mathbf{A} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times Z \lceil \log_2 q \rceil}, \\ \text{impk} &= (1^n, 1^m, q, 1^Z, \sigma_{-1}, \sigma_{\text{post}}, 1^K, \mathbf{A}, \mathbf{B}), & \text{imsk} &= (\text{impk}, \tau). \end{aligned}$$

The algorithm outputs $(\text{impk}, \text{imsk})$.

- $\text{KeyGen}(\text{imsk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g)$ takes as input $\text{imsk}, \text{id} = 1, \mathbf{v}_0 \in \mathbb{Z}_q^Z$, and $\mathbf{V}_g = \perp$. It runs

$$\mathbf{k} \stackrel{\$}{\leftarrow} \text{SampleD}(\mathbf{B}, \tau, \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_0), \sigma_{-1}).$$

Here, \mathbf{G} is of shape $Z \times Z \lceil \log_2 q \rceil$. The algorithm outputs $\text{isk} = \mathbf{k}$.

- $\text{Enc}(\text{impk}, \text{id}, \mathfrak{s}, \mathbf{u}^\top)$ takes as input $\text{impk}, \text{id} = 1, \mathfrak{s} \in \{\mathfrak{o}, \mathfrak{g}\}$, and $\mathbf{u} \in \mathbb{Z}_q^Z$. If $\mathfrak{s} = \mathfrak{g}$, the algorithm outputs \perp and terminates (the scheme does not produce any key capable of decrypting such a ciphertext). Otherwise, it samples

$$\mathbf{d} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^n, \quad \mathbf{e}_1 \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}, \leq \sigma_{\text{post}} \sqrt{\kappa}}^m, \quad \mathbf{e}_0 \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}, \leq \sigma_{\text{post}} \sqrt{\kappa}}^{Z \lceil \log_2 q \rceil}.$$

The algorithm outputs $\text{ict} = (\mathbf{d}^\top \mathbf{B} + \mathbf{e}_1^\top, \mathbf{d}^\top \mathbf{A} + \mathbf{e}_0^\top + \mathbf{u}^\top \mathbf{G})$.

- $\text{Dec}(\text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{s}, \text{ict})$ computes and outputs

$$\text{ict} \cdot \begin{pmatrix} -\text{isk} \\ \mathbf{G}^{-1}(\mathbf{v}_0) \end{pmatrix}.$$

Correctness. By Lemma 3, we have $\mathbf{B}\mathbf{k} = \mathbf{A}\mathbf{G}^{-1}(\mathbf{v}_0)$ and $\|\mathbf{k}\| \leq \sigma_{-1} \sqrt{m}$, so

$$\begin{aligned} \text{ict} \cdot \begin{pmatrix} -\text{isk} \\ \mathbf{G}^{-1}(\mathbf{v}_0) \end{pmatrix} - \mathbf{u}^\top \mathbf{v}_0 &= (\mathbf{d}^\top \mathbf{B} + \mathbf{e}_1^\top, \mathbf{d}^\top \mathbf{A} + \mathbf{e}_0^\top + \mathbf{u}^\top \mathbf{G}) \begin{pmatrix} -\mathbf{k} \\ \mathbf{G}^{-1}(\mathbf{v}_0) \end{pmatrix} - \mathbf{u}^\top \mathbf{v}_0 \\ &= -\mathbf{e}_1^\top \mathbf{k} + \mathbf{e}_0^\top \mathbf{G}^{-1}(\mathbf{v}_0), \\ |-\mathbf{e}_1^\top \mathbf{k} + \mathbf{e}_0^\top \mathbf{G}^{-1}(\mathbf{v}_0)| &\leq m \cdot \|\mathbf{e}_1\| \cdot \|\mathbf{k}\| + Z \lceil \log_2 q \rceil \cdot \|\mathbf{e}_0\| \cdot \|\mathbf{G}^{-1}(\mathbf{v}_0)\| \\ &\leq \kappa^{1/2} \sigma_{\text{post}} (m^{3/2} \sigma_{-1} + Z \lceil \log_2 q \rceil). \end{aligned}$$

Let B be the right-hand side, then Construction 1 is B -correct.

3.2 Security as an Assumption

Inspired by the evasive LWE assumption, we propose a new assumption that is closely related and translates to the security of Construction 1.

The precondition of evasive LWE rules out a combination of two attacking strategies [Wee22], attacks against LWE and zeroizing attacks. Alternatively, the evasive LWE assumption can be seen as a belief that *i*) the only meaningful way of using $\mathbf{B}^{-1}(\mathbf{P})$ is to multiply it to a noisy version of $\mathbf{d}^\top \mathbf{B}$, and that *ii*) the skewed noise attached to $\mathbf{d}^\top \mathbf{P}$ in $(\mathbf{d}^\top \mathbf{B} + \mathbf{e}^\top) \cdot \mathbf{B}^{-1}(\mathbf{P})$ does not give adversaries more advantage as long as there is *no apparent relation* among quantities related to \mathbf{d} when the noises are not skewed. The pseudorandomness in the precondition of evasive LWE exactly captures the absence of apparent relation.

Similar to the belief captured by evasive LWE, we propose the following:

Assumption 5 (evasive IPFE). Let $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ be two algorithms with the following syntax.

- $\mathcal{S}_V(1^\lambda; r_{\text{pub}})$, given randomness r_{pub} , outputs

$$1^n, 1^m, q, 1^Z, 1^J, 1^I, \quad \sigma_{-1}, \quad \sigma_{\text{post}} \geq \sigma_{\text{pre}} \geq 0, \quad \mathbf{V} \in \mathbb{Z}_q^{Z \times J},$$

where $m \geq m_0(n, q)$ and $\sigma_{-1} \geq \sigma_0(n, m)$ are constrained by Lemma 3.

- $\mathcal{S}_U(1^\lambda, r_{\text{pub}}; r_{\text{priv}})$, given r_{pub} and additional randomness r_{priv} , outputs \mathbf{U}^\top , where $\mathbf{U} \in \mathbb{Z}_q^{Z \times I}$ and q, Z, I agree with the output of $\mathcal{S}_V(1^\lambda; r_{\text{pub}})$.

Suppose

$$\begin{aligned} (\mathbf{B}, \tau) &\stackrel{\$}{\leftarrow} \text{TrapGen}(1^n, 1^m, q), & \mathbf{A} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times Z \lceil \log_2 q \rceil}, & \mathbf{D} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times I}, \\ \Delta_0 &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{Z \lceil \log_2 q \rceil \times I}, & \Delta_1 &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{m \times I}, & \Delta_2 &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{J \times I}, \\ \mathbf{E}_{\text{pre},0} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{pre}}}^{Z \lceil \log_2 q \rceil \times I}, & \mathbf{E}_{\text{pre},1} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{pre}}}^{m \times I}, & \mathbf{E}_{\text{pre},2} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{pre}}}^{J \times I}, \\ \mathbf{E}_{\text{post},0} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{post}}}^{Z \lceil \log_2 q \rceil \times I}, & \mathbf{E}_{\text{post},1} &\stackrel{\$}{\leftarrow} \mathcal{D}_{Z, \sigma_{\text{post}}}^{m \times I}, & & \\ \mathbf{K} &\stackrel{\$}{\leftarrow} \text{SampleD}(\mathbf{B}, \tau, \mathbf{A} \mathbf{G}^{-1}(\mathbf{V}), \sigma_{-1}). \end{aligned}$$

The *inner-product* precondition $\text{evIPFE}_{\text{ipre}}^{\mathcal{S}}$ is

$$\{(1^\lambda, r_{\text{pub}}, \mathbf{U}^\top \mathbf{V} + \mathbf{E}_{\text{pre},2}^\top)\}_{\lambda \in \mathbb{N}} \approx \{(1^\lambda, r_{\text{pub}}, \Delta_2^\top)\}_{\lambda \in \mathbb{N}}.$$

The *stricter* precondition $\text{evIPFE}_{\text{spre}}^{\mathcal{S}}$ is

$$\begin{aligned} &\{(1^\lambda, r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{D}^\top \mathbf{A} \mathbf{G}^{-1}(\mathbf{V}) + \mathbf{E}_{\text{pre},2}^\top, \mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{pre},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{pre},0}^\top + \mathbf{U}^\top \mathbf{G})\}_{\lambda \in \mathbb{N}} \\ &\approx \{(1^\lambda, r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \Delta_2^\top, \Delta_1^\top, \Delta_0^\top)\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

The *postcondition* $\text{evIPFE}_{\text{post}}^{\mathcal{S}}$ is

$$\begin{aligned} &\{(1^\lambda, r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{post},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{post},0}^\top + \mathbf{U}^\top \mathbf{G})\}_{\lambda \in \mathbb{N}} \\ &\approx \{(1^\lambda, r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \Delta_1^\top, \Delta_0^\top)\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

The $(n_0, \sigma_{\text{post},0})$ -evasive IPFE assumption states that $\text{evIPFE}_{\text{ipre}}^{\mathcal{S}}$ implies $\text{evIPFE}_{\text{post}}^{\mathcal{S}}$ for all efficient \mathcal{S} such that $n \geq n_0$ and $\sigma_{\text{post}} \geq \sigma_{\text{post},0}$. The *conservative evasive IPFE assumption* states that $\text{evIPFE}_{\text{spre}}^{\mathcal{S}}$ implies $\text{evIPFE}_{\text{post}}^{\mathcal{S}}$ for all efficient \mathcal{S} .²

The rationale of public and private coins in the above formulation is similar to that (Remark 3) of evasive IPFE security definition.

Evasive IPFE Security from Evasive IPFE Assumption. There are only minor syntactical differences between the samplers of evasive IPFE (Assumption 5) and IPFEsec (Definition 5) instantiated for Construction 1. The evasive IPFE assumption implies the security of our basic scheme:

Theorem 5 (¶). *Construction 1 is σ_{post} -secure (Definition 5) under $(n, \sigma_{\text{post}})$ -evasive IPFE (Assumption 5), where n, σ_{post} are those picked by Setup.*

Proof (Theorem 5). Let $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ be an efficient sampler (Definition 5) satisfying $\sigma_{\text{pre}} \leq \sigma_{\text{post}}$ and $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$. Consider the following $\mathcal{S}' = (\mathcal{S}'_V, \mathcal{S}'_U)$ for Assumption 5.

- $\mathcal{S}'_V(r_{\text{pub}})$ runs

$$(q, 1^0, 1^Z, \{1\}, 1^{J_1}, 1^{I_{1,0}}, 1^0, \sigma_{\text{pre}}, \mathbf{V}_{1,0}, \perp) \leftarrow \mathcal{S}_V(r_{\text{pub}}),$$

picks $n, m, \sigma_{-1}, \sigma_{\text{post}}$ according to Construction 1, sets J, I, \mathbf{V} to $J_1, I_{1,0}, \mathbf{V}_{1,0}$, and outputs³

$$1^n, 1^m, q, 1^Z, 1^J, 1^I, \quad \sigma_{-1}, \quad \sigma_{\text{post}}, \sigma_{\text{pre}}, \quad \mathbf{V}.$$

- $\mathcal{S}'_U(r_{\text{pub}}; r_{\text{priv}})$ runs $(\mathbf{U}_{1,0}^\top, \perp) \leftarrow \mathcal{S}_U(r_{\text{pub}}; r_{\text{priv}})$ and outputs $\mathbf{U}^\top = \mathbf{U}_{1,0}^\top$.

Clearly, $\text{evIPFE}_{\text{ipre}}^{\mathcal{S}'}$ is simply $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$, which is assumed to hold. Therefore, by the $(n, \sigma_{\text{post}})$ -evasive IPFE assumption, $\text{evIPFE}_{\text{post}}^{\mathcal{S}'}$ holds, implying

$$\begin{aligned} & (r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{post},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{post},0}^\top + \mathbf{U}^\top \mathbf{G} \quad) \\ & \approx (r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \quad \Delta_1^\top \quad , \quad \Delta_0^\top \quad) \\ & \equiv (r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \quad \Delta_1^\top \quad , \quad \Delta_0^\top + \mathbf{U}_\$^\top \mathbf{G} \quad) \\ & \approx (r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{post},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{post},0}^\top + (\mathbf{U} + \mathbf{U}_\$)^\top \mathbf{G} \quad) \\ & \equiv (r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \mathbf{K}, \mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{post},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{post},0}^\top + \mathbf{U}_\$^\top \mathbf{G} \quad), \end{aligned}$$

where $\mathbf{U}_\$ = \mathbf{U}_{\$,1,0} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{Z \times I}$. Opening up $\text{IPFEsec}_{\text{post}}^{\mathcal{S}}$ for Construction 1, we see that

$$\begin{aligned} & \underbrace{r_{\text{pub}}}_{\mathcal{S}} = r_{\text{pub}}, \quad \underbrace{r_{\text{pub}}}_{\mathcal{S}'} = r_{\text{pub}}, \quad \text{impk} = \overbrace{(1^n, 1^m, q, 1^Z, \sigma_{-1}, \sigma_{\text{post}}, 1^\kappa, \mathbf{A}, \mathbf{B})}^{\text{efficiently computable from } r_{\text{pub}}}, \quad \text{isk}_1 = \mathbf{K}, \\ & \text{ict}_{1,0}^* = (\mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{post},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{post},0}^\top + \mathbf{U}^\top \mathbf{G}), \\ & \text{ict}_{1,0}^{\$} = (\mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{post},1}^\top, \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{post},0}^\top + \mathbf{U}_\$^\top \mathbf{G}), \end{aligned}$$

²For the first part, explicitly requiring lower bounds on n, σ_{post} is necessary, because $\text{evIPFE}_{\text{ipre}}$ could be unconditional when $J < Z$ yet $\text{evIPFE}_{\text{post}}$ is always computational. For the second part, $\text{evIPFE}_{\text{spre}}$ implies LWE, so the lower bounds on n, σ_{post} are implicit.

³Without loss of generality, we may assume that \mathcal{S}_V always outputs $\text{ID} = \{1\}$ and $I_{1,g} = 0$. The other cases can be converted to $\text{ID} = \{1\}$ and $I_{1,0} = J_1 = 0$, with $\mathbf{U}_{1,g}^\top$ discarded, or both. The quantities $n, m, \sigma_{-1}, \sigma_{\text{post}}$ are determined by q, Z , so \mathcal{S}'_V does not rely on the randomness of Setup of Construction 1.

except for noise truncation only up to a negligible statistical error by Lemma 1. Therefore, $\text{IPFSec}_{\text{post}}^S$ follows from the previous indistinguishability. \square

Full and Conservative Versions. The evasive IPFE assumption might appear custom-made for the scheme, rendering the security proof almost tautological, and it looks far-fetched from and bears little resemblance to the evasive LWE assumption. We justify it below.

Intuitively, the only meaningful way of using $\mathbf{B}^{-1}(\mathbf{A}\mathbf{G}^{-1}(\mathbf{V}))$ is to compute

$$-\underbrace{\mathbf{D}^\top \mathbf{B}} \cdot \mathbf{B}^{-1}(\mathbf{A}\mathbf{G}^{-1}(\mathbf{V})) + (\underbrace{\mathbf{D}^\top \mathbf{A}} + \mathbf{U}^\top \mathbf{G})\mathbf{G}^{-1}(\mathbf{V}) = \underbrace{\mathbf{U}^\top \mathbf{V}},$$

where wavy underlines indicate noises. In this computation, it appears that *some* cancellation has happened so that $\mathbf{U}^\top \mathbf{V}$ is approximately recovered, which is different from evasive LWE, where *no* cancellation happens. But the precondition requires that the inner products with noise be pseudorandom, i.e., the ‘‘cancellation’’ only leads to pseudorandom results, morally equivalent to that there is no *actual* cancellation. It excludes the obvious zeroizing attacks.

The *conservative* evasive IPFE assumption is more along the lines of the evasive LWE assumption, for which the only addition is the ability to combine privately sampled value \mathbf{U} with LWE sample $\mathbf{D}^\top \mathbf{A}$. It appears to be a smaller leap of faith than the full version, yet this is actually superfluous.

In fact, the full version does not demand more than the conservative version in suitable parameter regimes. We show that $\text{evIPFE}_{\text{ipre}}^S$ implies $\text{evIPFE}_{\text{spre}}^S$ under the LWE assumption with noise super-polynomially smaller than σ_{pre} :

Theorem 6 (¶). *If \mathcal{S} (Assumption 5) is efficient and $\text{LWE}_{n,m+Z\lceil \log_2 q \rceil, q, \sigma_{\text{help}}}$ (Assumption 1) holds for some $\sigma_{\text{help}} \leq \frac{\sigma_{\text{pre}}}{2^{\kappa+6}\sqrt{\kappa} \cdot Z\lceil \log_2 q \rceil}$, where $q, Z, \sigma_{\text{pre}}$ are those picked by \mathcal{S} , then $\text{evIPFE}_{\text{ipre}}^S$ implies $\text{evIPFE}_{\text{spre}}^S$.*

We obtain a conservative version of Theorem 5 as a corollary.

Corollary 7 (¶). *Construction 1 is σ_{post} -secure (Definition 5) under conservative evasive IPFE (Assumption 5) and $\text{LWE}_{n,m+Z\lceil \log_2 q \rceil, q, \sigma_{\text{help}}}$ (Assumption 1) for some $\sigma_{\text{help}} \leq \frac{\sigma_{\text{post}}}{(2^{\kappa+6}\sqrt{\kappa})^2 Z\lceil \log_2 q \rceil}$, where q, Z are those picked by \mathcal{S} and $n, m, \sigma_{\text{post}}, \kappa$ are those picked by Setup.*

The idea for the proof of Theorem 6 is

$$\begin{aligned} & \left(\underbrace{\mathbf{D}^\top \mathbf{A}\mathbf{G}^{-1}(\mathbf{V})}, \underbrace{\mathbf{D}^\top \mathbf{B}}, \underbrace{\mathbf{D}^\top \mathbf{A} + \mathbf{U}^\top \mathbf{G}} \right) \\ \text{(noise flooding)} & \approx_s \left((\underbrace{\mathbf{D}^\top \mathbf{A} + \mathbf{U}^\top \mathbf{G}})\mathbf{G}^{-1}(\mathbf{V}) - \underbrace{\mathbf{U}^\top \mathbf{V}}, \underbrace{\mathbf{D}^\top \mathbf{B}}, \underbrace{\mathbf{D}^\top \mathbf{A} + \mathbf{U}^\top \mathbf{G}} \right) \\ \text{(LWE)} & \approx \left(\underbrace{\Delta_0^\top \mathbf{G}^{-1}(\mathbf{V}) - \mathbf{U}^\top \mathbf{V}}, \underbrace{\Delta_1^\top}, \underbrace{\Delta_0^\top} \right) \\ \text{(precondition)} & \approx \left(\underbrace{\Delta_2^\top}, \underbrace{\Delta_1^\top}, \underbrace{\Delta_0^\top} \right), \end{aligned}$$

where wavy underlines indicate noises. The rewriting in the first step is essentially IPFE simulation [Wee17,ACF⁺18,ALMT20,LL20b], i.e., $\underbrace{\mathbf{D}^\top \mathbf{A}\mathbf{G}^{-1}(\mathbf{V})}$ contains exactly the information of $\underbrace{\mathbf{U}^\top \mathbf{V}}$, when given $\underbrace{\mathbf{D}^\top \mathbf{B}}$ and $(\underbrace{\mathbf{D}^\top \mathbf{A} + \mathbf{U}^\top \mathbf{G}})$.

Proof (Theorem 6). We consider the following hybrids.

- H_0 is the first distribution in $\text{evIPFE}_{\text{spre}}^S$, i.e.,

$$\begin{aligned} r_{\text{pub}}, \mathbf{A}, \mathbf{B}, & \quad \mathbf{D}^\top \mathbf{A} \mathbf{G}^{-1}(\mathbf{V}) + \mathbf{E}_{\text{pre},2}^\top, \\ \mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{pre},1}^\top, & \quad \mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{pre},0}^\top + \mathbf{U}^\top \mathbf{G}. \end{aligned}$$

- In H_1 , the matrix \mathbf{B} is sampled uniformly at random, without a known trapdoor. By Lemma 3, we have $H_0 \approx_s H_1$.
- In H_2 , helper noises are inserted and flooded by \mathbf{E}_{pre} 's, i.e.,

$$\begin{aligned} r_{\text{pub}}, \mathbf{A}, \mathbf{B}, & \quad \mathbf{D}^\top \mathbf{A} \mathbf{G}^{-1}(\mathbf{V}) + (\mathbf{E}_{\text{pre},2}^\top + \mathbf{E}_{\text{help},0}^\top \mathbf{G}^{-1}(\mathbf{V})), \\ \mathbf{D}^\top \mathbf{B} + (\mathbf{E}_{\text{pre},1} + \mathbf{E}_{\text{help},1})^\top, & \quad \mathbf{D}^\top \mathbf{A} + \mathbf{U}^\top \mathbf{G} + (\mathbf{E}_{\text{pre},0} + \mathbf{E}_{\text{help},0})^\top, \end{aligned}$$

where the entries of \mathbf{E}_{help} 's are sampled independently from $\mathcal{D}_{Z, \sigma_{\text{help}}, \leq \sigma_{\text{help}} \sqrt{\kappa}}$. Since $\|(\mathbf{G}^{-1}(\mathbf{V}))^\top\| \leq Z \lceil \log_2 q \rceil$, the entries of (\mathbf{E}_{help} 's and) $\mathbf{E}_{\text{help},0}^\top \mathbf{G}^{-1}(\mathbf{V})$ are bounded by $\sigma_{\text{help}} \sqrt{\kappa} \cdot Z \lceil \log_2 q \rceil$. Together with $\sigma_{\text{pre}} \geq 2^{\kappa+6} \sigma_{\text{help}} \sqrt{\kappa} \cdot Z \lceil \log_2 q \rceil$, it follows from Lemma 2 that $H_1 \approx_s H_2$.

- In H_3 , we change $\mathbf{E}_{\text{pre},2}$ to $-\mathbf{E}_{\text{pre},2}$, rearrange some terms, and derive the noisy $\mathbf{D}^\top \mathbf{A} \mathbf{G}^{-1}(\mathbf{V})$ from the IPFE decryption relation, i.e.,

$$\begin{aligned} r_{\text{pub}}, \mathbf{A}, \mathbf{B}, & \quad (\mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{help},0}^\top + \mathbf{U}^\top \mathbf{G}) \mathbf{G}^{-1}(\mathbf{V}) - (\mathbf{U}^\top \mathbf{V} + \mathbf{E}_{\text{pre},2}^\top), \\ (\mathbf{D}^\top \mathbf{B} + \mathbf{E}_{\text{help},1}^\top) + \mathbf{E}_{\text{pre},0}^\top, & \quad (\mathbf{D}^\top \mathbf{A} + \mathbf{E}_{\text{help},0}^\top + \mathbf{U}^\top \mathbf{G}) + \mathbf{E}_{\text{pre},0}^\top. \end{aligned}$$

Since $\mathbf{E}_{\text{pre},2}$ follows a symmetric distribution, $H_2 \equiv H_3$.

- In H_4 , the \mathbf{E}_{help} 's are no longer truncated, i.e., their entries are sampled independently from $\mathcal{D}_{Z, \sigma_{\text{help}}}$. By Lemma 1, it holds that $H_3 \approx_s H_4$.
- In H_5 , we invoke $\text{LWE}_{n, m+Z \lceil \log_2 q \rceil, q, \sigma_{\text{help}}}$ for secrets⁴ \mathbf{D} and public matrix (\mathbf{B}, \mathbf{A}) , and the distribution becomes

$$\begin{aligned} r_{\text{pub}}, \mathbf{A}, \mathbf{B}, & \quad (\mathbf{\Delta}_0^\top + \mathbf{U}^\top \mathbf{G}) \mathbf{G}^{-1}(\mathbf{V}) - (\mathbf{U}^\top \mathbf{V} + \mathbf{E}_{\text{pre},2}^\top), \\ \mathbf{\Delta}_1^\top + \mathbf{E}_{\text{pre},1}^\top, & \quad (\mathbf{\Delta}_0^\top + \mathbf{U}^\top \mathbf{G}) + \mathbf{E}_{\text{pre},0}^\top. \end{aligned}$$

By the LWE assumption, $H_4 \approx H_5$.⁵

- In H_6 , we make $\mathbf{\Delta}$'s absorb terms added to them, i.e.,

$$r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \quad (\mathbf{\Delta}_0^\top - \mathbf{E}_{\text{pre},0}^\top) \mathbf{G}^{-1}(\mathbf{V}) - (\mathbf{U}^\top \mathbf{V} + \mathbf{E}_{\text{pre},2}^\top), \quad \mathbf{\Delta}_1^\top, \quad \mathbf{\Delta}_0^\top.$$

Clearly, $H_5 \equiv H_6$.

- In H_7 , we invoke $\text{evIPFE}_{\text{ipre}}^S$ so the distribution becomes

$$r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \quad (\mathbf{\Delta}_0^\top - \mathbf{E}_{\text{pre},0}^\top) \mathbf{G}^{-1}(\mathbf{V}) - \mathbf{\Delta}_2^\top, \quad \mathbf{\Delta}_1^\top, \quad \mathbf{\Delta}_0^\top.$$

By the assumption, $H_6 \approx H_7$.

⁴Technically, there are a series of hybrids over the columns of \mathbf{D} .

⁵This step requires that \mathcal{S} (part of the reduction algorithm) be efficient.

- H_8 is the second distribution in $\text{evIPFE}_{\text{spre}}^{\mathcal{S}}$, i.e.,

$$r_{\text{pub}}, \mathbf{A}, \mathbf{B}, \quad \Delta_2^\top, \quad \Delta_1^\top, \quad \Delta_0^\top.$$

Clearly, $H_7 \equiv H_8$.

By a hybrid argument, $H_0 \approx H_8$, i.e., $\text{evIPFE}_{\text{spre}}^{\mathcal{S}}$ holds. \square

Proof (Corollary 7). Given an efficient \mathcal{S} (Definition 5) satisfying $\sigma_{\text{pre}} \leq \sigma_{\text{post}}$ and $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$, we do the following:

1. construct \mathcal{S}' (Assumption 5) *similarly* to the proof of Theorem 5 and show that $\text{evIPFE}_{\text{ipre}}^{\mathcal{S}'}$ holds,
2. apply Theorem 6 to \mathcal{S}' for proving $\text{evIPFE}_{\text{spre}}^{\mathcal{S}'}$,
3. derive $\text{evIPFE}_{\text{post}}^{\mathcal{S}'}$ from the conservative evasive IPFE assumption, and
4. conclude $\text{IPFEsec}_{\text{post}}^{\mathcal{S}}$ in the same way as the final parts of the proof of Theorem 5.

Let $\sigma_{\text{pre}}^{\mathcal{S}}, \sigma_{\text{pre}}^{\mathcal{S}'}$ be the σ_{pre} of $\mathcal{S}, \mathcal{S}'$. We set

$$\sigma_{\text{pre}}^{\mathcal{S}'} = \begin{cases} \sigma_{\text{pre}}^{\mathcal{S}}, & \text{if } \sigma_{\text{pre}}^{\mathcal{S}} \geq \frac{\sigma_{\text{post}}}{2^{\kappa+6}\sqrt{\kappa}}; \\ \sigma_{\text{post}}, & \text{otherwise.} \end{cases}$$

In the first case, $\text{evIPFE}_{\text{ipre}}^{\mathcal{S}'}$ is just $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$ and Theorem 6 applies because

$$\sigma_{\text{help}} \leq \frac{\sigma_{\text{post}}}{(2^{\kappa+6}\sqrt{\kappa})^2 Z[\log_2 q]} \leq \frac{\sigma_{\text{pre}}^{\mathcal{S}'}}{2^{\kappa+6}\sqrt{\kappa} \cdot Z[\log_2 q]}.$$

In the second case, we cannot set $\sigma_{\text{pre}}^{\mathcal{S}'}$ to $\sigma_{\text{pre}}^{\mathcal{S}}$, which might not be large enough compared to σ_{help} per the premise of Theorem 6. Instead, with our choice, $\text{evIPFE}_{\text{ipre}}^{\mathcal{S}'}$ follows from $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$ by noise truncation and flooding (Lemmas 1 and 2), and Theorem 6 again applies. \square

We remark that for security against \mathcal{S} with particular σ_{pre} 's (e.g., $\sigma_{\text{pre}} = \sigma_{\text{post}}$), the bound on σ_{help} can be tuned accordingly (e.g., without the quadratic blow-up).

3.3 Security for Restricted Samplers from Evasive LWE

We prove that Construction 1 is \mathcal{S} -secure for a restricted class of samplers under the evasive LWE assumption. These samplers are efficient, satisfy $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$, and their \mathcal{S}_U 's output a uniformly random vector in a publicly known subspace. This will suffice for our ABE applications.

Definition 6 (restricted IB-evIPFE-w/sn security). A sampler $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ per Definition 5 is *restricted* if $\mathcal{S}_U(1^\lambda, r_{\text{pub}}; r_{\text{priv}})$ works as follows.

1. It uses a fixed deterministic algorithm \mathcal{S}_{A_0} to compute

$$(1^{n'}, \{\mathbf{A}_{\text{id},s,0,i}\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}, i \in [I_{\text{id},s}]}) \leftarrow \mathcal{S}_{A_0}(1^\lambda, r_{\text{pub}}),$$

where $\mathbf{A}_{\text{id},s,0,i} \in \mathbb{Z}_q^{n' \times Z}$ for all $\text{id} \in \text{ID}$, $s \in \{\mathfrak{o}, \mathfrak{g}\}$, $i \in [I_{\text{id},s}]$, and the values of q, Z, ID , $\{I_{\text{id},s}\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}}$ agree with the output of $\mathcal{S}_V(1^\lambda; r_{\text{pub}})$.

2. It uses r_{priv} to sample $\mathbf{d}_0 \xleftarrow{\$} \mathbb{Z}_q^{n'}$. The algorithm outputs $\{\mathbf{U}_{\text{id},s}^\top\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}}$

$$\begin{aligned} \text{for } \quad & \mathbf{U}_{\text{id},s} = (\mathbf{u}_{\text{id},s,1}, \dots, \mathbf{u}_{\text{id},s,I_{\text{id},s}}) \\ \text{with } \quad & \mathbf{u}_{\text{id},s,i}^\top = \mathbf{d}_0^\top \mathbf{A}_{\text{id},s,0,i} \quad \text{for all } \text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}, i \in [I_{\text{id},s}]. \end{aligned}$$

An IB-evIPFE-w/sn scheme (Definition 4) is *restricted- $\sigma_{\text{pre},0}$ -secure* if it is \mathcal{S} -secure (Definition 5) for all efficient restricted \mathcal{S} such that $\sigma_{\text{pre}} \leq \sigma_{\text{pre},0}$ and $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$ holds.

Parallel to restricted security, we formulate the following:

Assumption 6 (restricted conservative evasive IPFE). A sampler $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ for Assumption 5 is *restricted* if $\mathcal{S}_U(1^\lambda, r_{\text{pub}}; r_{\text{priv}})$ works as follows.

1. It uses a fixed deterministic algorithm \mathcal{S}_{A_0} to compute

$$(1^{n'}, \{\mathbf{A}'_{0,i}\}_{i \in [I]}) \leftarrow \mathcal{S}_{A_0}(1^\lambda, r_{\text{pub}}) \quad \text{with } \mathbf{A}'_{0,i} \in \mathbb{Z}_q^{n' \times Z} \text{ for all } i \in [I],$$

where q, Z, I agree with the output of $\mathcal{S}_V(1^\lambda; r_{\text{pub}})$.

2. It uses r_{priv} to sample $\mathbf{d}_0 \xleftarrow{\$} \mathbb{Z}_q^{n'}$. The algorithm outputs \mathbf{U}^\top for

$$\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_I) \quad \text{and} \quad \mathbf{u}_i^\top = \mathbf{d}_0^\top \mathbf{A}'_{0,i} \text{ for all } i \in [I].$$

The *restricted conservative evasive IPFE assumption* states that $\text{evIPFE}_{\text{spre}}^{\mathcal{S}}$ implies $\text{evIPFE}_{\text{post}}^{\mathcal{S}}$ for all efficient restricted \mathcal{S} .

Lemma 8. *The sampler (Assumption 5) constructed in the proof of Corollary 7 is restricted (Assumption 6) if so is the original sampler (Definitions 5 and 6).*

Lemma 8 can be easily verified. The following implication holds:

Theorem 9 (¶). *Evasive LWE (Assumption 4) implies restricted conservative evasive IPFE (Assumption 6).*

Combining Theorem 6, the proof of Corollary 7, Lemma 8, and Theorem 9:

Corollary 10. *Construction 1 is restricted- σ_{post} -secure (Definition 6) under evasive LWE (Assumption 4) and $\text{LWE}_{n,m+Z \lceil \log_2 q \rceil, q, \sigma_{\text{help}}}$ (Assumption 1) for some $\sigma_{\text{help}} \leq \frac{\sigma_{\text{post}}}{(2^{\kappa+6} \sqrt{\kappa})^2 Z \lceil \log_2 q \rceil}$, where q, Z are those picked by \mathcal{S} and $n, m, \sigma_{\text{post}}, \kappa$ are those picked by Setup.*

Proof (Theorem 9). Let $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ be an efficient restricted sampler with \mathcal{S}_{A_0} (Assumptions 5 and 6) such that $\text{evIPFE}_{\text{spre}}^{\mathcal{S}}$ holds. Construct the following $\mathcal{S}'(r)$ for Assumption 4. It parses $r = (r_{\text{pub}}, r'_{\text{pub}})$ and runs

$$\begin{aligned} (1^n, 1^m, q, 1^Z, 1^J, 1^I, \sigma_{-1}, \sigma_{\text{post}}, \sigma_{\text{pre}}, \mathbf{V}) & \leftarrow \mathcal{S}_V(r_{\text{pub}}), \\ (1^{n'}, \{\mathbf{A}'_{0,i}\}_{i \in [I]}) & \leftarrow \mathcal{S}_{A_0}(r_{\text{pub}}). \end{aligned}$$

The algorithm samples \mathbf{A} uniformly random over $\mathbb{Z}_q^{n \times Z \lceil \log_2 q \rceil}$ using r'_{pub} in the straightforward manner,⁶ and sets $K = 0$ and $m' = Z \lceil \log_2 q \rceil$. It outputs

$$1^n, 1^m, 1^K, q, 1^{n'}, 1^J, 1^I, 1^{m'}, \quad \sigma_{-1}, \quad \sigma_{\text{post}}, \quad \sigma_{\text{pre}},$$

$$\tilde{\mathbf{P}} = \mathbf{A}\mathbf{G}^{-1}(\mathbf{V}), \quad \{\mathbf{A}_{1,i}, \mathbf{A}_{0,i}\}_{i \in [I]} \quad (\text{with } \mathbf{A}_{1,i} = \mathbf{A} \text{ and } \mathbf{A}_{0,i} = \mathbf{A}'_{0,i}\mathbf{G}),$$

By setting

$$\mathbf{D} = (\mathbf{d}_1, \dots, \mathbf{d}_I), \quad \mathbf{\Delta}_\alpha = (\boldsymbol{\delta}_{1,\alpha}, \dots, \boldsymbol{\delta}_{I,\alpha}) \text{ for all } \alpha \in \{0, 1, 2\},$$

$$\mathbf{E}_{p,\alpha} = (\mathbf{e}_{p,1,\alpha}, \dots, \mathbf{e}_{p,I,\alpha}) \text{ for all } (p, \alpha) \in \left\{ \begin{array}{l} (\text{pre}, 2), (\text{pre}, 1), (\text{pre}, 0), \\ (\text{post}, 1), (\text{post}, 0) \end{array} \right\},$$

the components in $\text{evIPFE}_{\text{spre/post}}^S, \text{evLWE}_{\text{pre/post}}^{S'}$ are perfectly lined up, except that \mathbf{A} corresponds to r'_{pub} . It is readily verified that $\text{evLWE}_{\text{pre}}^{S'}$ follows from $\text{evIPFE}_{\text{spre}}^S$. By the evasive LWE assumption, $\text{evLWE}_{\text{post}}^{S'}$ holds, implying $\text{evIPFE}_{\text{post}}^S$.⁷ \square

3.4 Identity-Based Scheme

Construction 1 only supports a single identity. We employ pairing-based techniques to generically obtain an exponentially large identity space. The core idea is to generate identity-based keys and ciphertexts using a single-identity IPFE (denoted with primes), e.g., when not considering structured noises,

$$\left. \begin{array}{l} \text{isk}_{\text{id}'}(\mathbf{v}) : \text{isk}'(\mathbf{v}, \quad \psi, \text{id}' \cdot \psi) \\ \text{ict}_{\text{id}}(\mathbf{u}) : \text{ict}'(\mathbf{u}, \text{id} \cdot \varphi, \quad -\varphi) \end{array} \right\} \xrightarrow{\text{Dec}'} \mathbf{u}^\top \mathbf{v} + (\text{id} - \text{id}')\varphi\psi$$

so that decryption yields $\mathbf{u}^\top \mathbf{v}$ if $\text{id} = \text{id}'$ and the result is masked by a blinding factor $(\text{id} - \text{id}')\varphi\psi$ if $\text{id} \neq \text{id}'$. Intuitively, the masks are pseudorandom by the DDH assumption (this idea can be materialized in various ways). In our adaptation, we use the LWE assumption to create the pseudorandom blinding factors.

We note that the transformation of [GKW16] compressing exponentially many copies of a single-identity scheme into an identity-based scheme is not as desirable as the upcoming construction. Security of $\text{evIPFE}(-\text{w/sn})$ is knowledge-type. Since Definition 5 requires fully private-coin sampling of plaintext vectors (see Remark 3), it cannot be applied simultaneously for two copies unless the two copies only encrypt independent plaintext vectors — but in our applications, plaintext vectors under different identities are correlated. Moreover, even if we consider an alternative definition with arbitrary (efficiently computable) auxiliary information about the plaintext vectors, we would have to assume evasive LWE with very tight pre/postcondition distinguishing gap for it to be applied super-constantly many times (corresponding to the number of identities; used in CP-ABE for circuits).

⁶Precisely speaking, the distribution of r'_{pub} conditioned on \mathbf{A} should be efficiently sampleable from \mathbf{A} so that it does not contain a trapdoor of \mathbf{A} or otherwise impedes security. If we insist that r'_{pub} be a binary string, we can make r'_{pub} at least $\max_{r_{\text{pub}}}((\kappa + \lceil \log_2 q \rceil) \cdot nZ \lceil \log_2 q \rceil)$ bits long and set $\mathbf{A}[i, j]$ to be its $((i-1) + (j-1)n + 1)^{\text{st}}$ chunk of $(\kappa + \lceil \log_2 q \rceil)$ bits interpreted as an integer in binary then reduced modulo q , incurring a statistical error of $2^{-\kappa} nZ \lceil \log_2 q \rceil$.

⁷The reduction algorithms reshape components between vectors and matrices, and either reversely sample r'_{pub} from \mathbf{A} or computes \mathbf{A} from r'_{pub} .

Ingredients of Construction 2. Let $\text{IPFE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ be an underlying $\text{evIPFE}(-\text{w/sn})$ such that $1 \in \mathcal{I}'_{q,K,Z}$ (e.g., Constructions 1 and 3).

Construction 2 (IB- $\text{evIPFE}(-\text{w/sn})$). Our scheme works as follows.

- $\text{Setup}(q, 1^K, 1^Z)$ takes as input q, K, Z . If q is not a prime or $1 \notin \mathcal{I}'_{q,K,Z}$, the algorithm sets $\mathcal{I}_{q,K,Z} = \emptyset$ and terminates (aborting case). Otherwise, it sets $\mathcal{I}_{q,K,Z} = [q]$ ⁸ and *deterministically* picks n .⁹ The algorithm runs

$$(\text{impk}', \text{imsk}') \stackrel{\$}{\leftarrow} \text{Setup}'(q, 1^K, 1^{Z+2n})$$

and outputs $\text{impk} = (q, 1^K, 1^Z, 1^n, \text{impk}')$ and $\text{imsk} = (q, 1^K, 1^Z, 1^n, \text{imsk}')$.

- $\text{KeyGen}(\text{imsk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g)$ takes as input imsk , $\text{id} \in [q]$, $\mathbf{v}_0 \in \mathbb{Z}_q^Z$, and $\mathbf{V}_g \in \mathbb{Z}_q^{Z \times K}$. It samples $\boldsymbol{\psi}_0 \stackrel{\$}{\leftarrow} \mathbb{Z}_q^n$ and $\boldsymbol{\Psi}_g \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times K}$, runs

$$\text{isk}' \stackrel{\$}{\leftarrow} \text{KeyGen}'\left(\text{imsk}', 1, \begin{pmatrix} \mathbf{v}_0 \\ \boldsymbol{\psi}_0 \\ \text{id} \cdot \boldsymbol{\psi}_0 \end{pmatrix}, \begin{pmatrix} \mathbf{V}_g \\ \boldsymbol{\Psi}_g \\ \text{id} \cdot \boldsymbol{\Psi}_g \end{pmatrix}\right),$$

and outputs $\text{isk} = (\boldsymbol{\psi}_0, \boldsymbol{\Psi}_g, \text{isk}')$.

- $\text{Enc}(\text{impk}, \text{id}, \mathfrak{s}, \mathbf{u}^\top)$ takes as input imsk , $\text{id} \in [q]$, $\mathfrak{s} \in \{\mathfrak{o}, \mathfrak{g}\}$, and $\mathbf{u} \in \mathbb{Z}_q^Z$. It samples $\boldsymbol{\varphi} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^n$, and runs and outputs

$$\text{ict} = \text{ict}' \stackrel{\$}{\leftarrow} \text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{u}^\top, \text{id} \cdot \boldsymbol{\varphi}^\top, -\boldsymbol{\varphi}^\top)).$$

- $\text{Dec}(\text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{s}, \text{ict})$ runs and outputs

$$\text{Dec}'\left(\text{impk}', 1, \begin{pmatrix} \mathbf{v}_0 \\ \boldsymbol{\psi}_0 \\ \text{id} \cdot \boldsymbol{\psi}_0 \end{pmatrix}, \begin{pmatrix} \mathbf{V}_g \\ \boldsymbol{\Psi}_g \\ \text{id} \cdot \boldsymbol{\Psi}_g \end{pmatrix}, \text{isk}', \mathfrak{s}, \text{ict}'\right).$$

Correctness. Suppose IPFE' is $B'(q', K', Z')$ -correct, then for noise structure \mathfrak{o} ,

$$\begin{aligned} \text{Dec}(\text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{o}, \text{ict}) &= \mathbf{u}^\top \cdot \mathbf{v}_0 + (\text{id} \cdot \boldsymbol{\varphi}^\top) \cdot \boldsymbol{\psi}_0 + (-\boldsymbol{\varphi}^\top) \cdot (\text{id} \cdot \boldsymbol{\psi}_0) + e_0 \\ &= \mathbf{u}^\top \mathbf{v}_0 + e_0 \end{aligned}$$

for some $e_0 \in [-B'(q, K, Z + 2n), B'(q, K, Z + 2n)]$, and similarly for the case of \mathfrak{g} . In summary, Construction 2 is B -correct for $B(q, K, Z) = B'(q, K, Z + 2n)$.

Security. We state and prove the security of Construction 2.

Theorem 11 (♣). *Suppose IPFE' is $\sigma_{\text{pre},0}$ -secure (Definition 5) and $\text{LWE}_{n,J'_1,q,\sigma_{\text{help}}}$ (Assumption 1) holds for some $\sigma_{\text{help}} \leq \frac{\sigma_{\text{pre},0}}{(2^{\kappa+6}\sqrt{\kappa})^2}$,¹⁰ then Construction 2 is also $\sigma_{\text{pre},0}$ -secure, where q and $J'_1 = \sum_{\text{id} \in \text{ID}} J_{\text{id}}$ are picked by \mathcal{S} and n is picked by Setup .*

Suppose instead IPFE' is restricted- $\sigma_{\text{pre},0}$ -secure (Definition 6) while keeping the other conditions intact, then Construction 2 is also restricted- $\sigma_{\text{pre},0}$ -secure.

⁸Due to homomorphic noise growth and noise flooding, q is super-polynomial in our work. When q is polynomially large, the identity space can be enlarged using full-rank difference hash [ABB10].

⁹This construction is generic. The quantity n here is unrelated to those in Construction 1 or 3, although they can be set to the same value.

¹⁰Here, κ is unrelated to that in Construction 1, although they can be set to the same value.

The proof works by arguing that when the noisy inner products for the matching identities are pseudorandom, so are the inner products concerned by IPFE', e.g., for \mathfrak{o} ,

$$\begin{cases} \text{id} \neq \text{id}' : \mathbf{U}_{\text{id},\mathfrak{o}}^\top \mathbf{V}_{\text{id}',\mathfrak{o}} + (\text{id} - \text{id}') \Phi_{\text{id},\mathfrak{o}}^\top \Psi_{\text{id}',\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id}',\mathfrak{o}})^\top & \approx \mathbf{U}_{\text{id},\mathfrak{o}}^\top \mathbf{V}_{\text{id}',\mathfrak{o}} + \$ \equiv \$, & \text{IPFEsec}_{\text{pre}}^{\mathcal{S}} \\ \text{id} = \text{id}' : \mathbf{U}_{\text{id},\mathfrak{o}}^\top \mathbf{V}_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{o}})^\top & \underset{\text{LWE}}{\approx} & \approx \$, & \downarrow \end{cases}$$

When the noise widths of IPFEsec $_{\text{pre}}^{\mathcal{S}}$ and LWE do not match, we resolve the issue similarly to the proof of Corollary 7.

Proof (Theorem 11). Let $\mathcal{S} = (\mathcal{S}_V, \mathcal{S}_U)$ be an efficient sampler (Definition 5) such that $\sigma_{\text{pre}} \leq \sigma_{\text{pre},0}$ and IPFEsec $_{\text{pre}}^{\mathcal{S}}$ holds against Construction 2. Consider the following efficient $\mathcal{S}' = (\mathcal{S}'_V, \mathcal{S}'_U)$.

- $\mathcal{S}'_V(r''_{\text{pub}})$ parses $r''_{\text{pub}} = (r_{\text{pub}}, r'_{\text{pub}})$ and runs

$$(q, 1^K, 1^Z, \text{ID}, \{1^{J_{\text{id}}}, 1^{I_{\text{id},\mathfrak{o}}}, 1^{I_{\text{id},\mathfrak{g}}}\}_{\text{id} \in \text{ID}}, \sigma_{\text{pre}}, \{\mathbf{V}_{\text{id},\mathfrak{o}}, \mathbf{V}_{\text{id},\mathfrak{g}}\}_{\text{id} \in \text{ID}}) \leftarrow \mathcal{S}_V(r_{\text{pub}}).$$

The algorithm aborts if q is not a prime or $1 \notin \mathcal{I}'_{q,K,Z}$. Otherwise, it computes n used by Setup of Construction 2, and samples $\{\Psi_{\text{id},s}\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}}$, with $\Psi_{\text{id},\mathfrak{o}}$ (resp. $\Psi_{\text{id},\mathfrak{g}}$) uniformly random over $\mathbb{Z}_q^{n \times J_{\text{id}}}$ (resp. $\mathbb{Z}_q^{n \times K_{J_{\text{id}}}}$) for all $\text{id} \in \text{ID}$, using r''_{pub} in the straight-forward manner.¹¹ The algorithm sets

$$\begin{aligned} Z' &= Z + 2n, & \text{ID}' &= \{1\}, & J'_1 &= \sum_{\text{id} \in \text{ID}} J_{\text{id}}, & I'_{1,s} &= \sum_{\text{id} \in \text{ID}} I_{\text{id},s}, \\ \sigma'_{\text{pre}} &= \begin{cases} \sigma_{\text{pre}}, & \text{if } \sigma_{\text{pre}} \geq \frac{\sigma_{\text{pre},0}}{2^{\kappa+6}\sqrt{\kappa}}; \\ \sigma_{\text{pre},0}, & \text{otherwise;} \end{cases} & \mathbf{V}'_{1,s} &= \begin{pmatrix} \cdots & \mathbf{V}_{\text{id},s} & \cdots \\ \cdots & \Psi_{\text{id},s} & \cdots \\ \cdots & \text{id} \cdot \Psi_{\text{id},s} & \cdots \end{pmatrix}_{\text{id} \in \text{ID}}, \end{aligned}$$

and outputs $(q, 1^K, 1^{Z'}, \text{ID}', 1^{J'_1}, 1^{I'_{1,\mathfrak{o}}}, 1^{I'_{1,\mathfrak{g}}}, \sigma'_{\text{pre}}, \mathbf{V}'_{1,\mathfrak{o}}, \mathbf{V}'_{1,\mathfrak{g}})$.

- $\mathcal{S}'_U(r''_{\text{pub}}; r''_{\text{priv}})$ parses $r''_{\text{pub}} = (r_{\text{pub}}, r'_{\text{pub}})$ and $r''_{\text{priv}} = (r_{\text{priv}}, r'_{\text{priv}})$. It recomputes q , ID , $\{I_{\text{id},s}\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}}$ and n as in \mathcal{S}'_V . The algorithm runs

$$\{\mathbf{U}_{\text{id},s}^\top\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}} \leftarrow \mathcal{S}_U(r_{\text{pub}}; r_{\text{priv}}),$$

and samples $\Phi_{\text{id},s}$ uniformly random over $\mathbb{Z}_q^{n \times I_{\text{id},s}}$ for all $\text{id} \in \text{ID}$ and $s \in \{\mathfrak{o}, \mathfrak{g}\}$ using r'_{priv} . It sets

$$(\mathbf{U}'_{1,s})^\top = \begin{pmatrix} \vdots & \vdots & \vdots \\ \mathbf{U}_{\text{id},s}^\top & \text{id} \cdot \Phi_{\text{id},s}^\top & -\Phi_{\text{id},s}^\top \\ \vdots & \vdots & \vdots \end{pmatrix}_{\text{id} \in \text{ID}}$$

and outputs $\{(\mathbf{U}'_{1,s})^\top\}_{s \in \{\mathfrak{o}, \mathfrak{g}\}}$.

Claim 12 (♣). IPFEsec $_{\text{pre}}^{\mathcal{S}'}$ (Definition 5) holds.

¹¹The distribution of r''_{pub} conditioned on $\{\Psi_{\text{id},s}\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}}$ must be efficiently sampleable from $\{\Psi_{\text{id},s}\}_{\text{id} \in \text{ID}, s \in \{\mathfrak{o}, \mathfrak{g}\}}$. See Footnote 6.

From the $\sigma_{\text{pre},0}$ -security of IPFE', it follows that

$$\begin{aligned} & (r''_{\text{pub}}, \text{impk}', \text{isk}'_1, \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{U}'_{1,\mathfrak{s}})^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}) \\ & \approx (r''_{\text{pub}}, \text{impk}', \text{isk}'_1, \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, \mathbf{\Delta}_{\mathfrak{s}}^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}), \end{aligned}$$

where $\mathbf{\Delta}_{\mathfrak{s}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{Z' \times I'_{1,\mathfrak{s}}}$. (This is the ‘‘initial part’’ of this proof.)

Now consider this efficient $\mathcal{S}_{\mathfrak{s}} = (\mathcal{S}_{\mathfrak{s},\mathfrak{v}}, \mathcal{S}_{\mathfrak{s},\mathfrak{u}})$:

- $\mathcal{S}_{\mathfrak{s},\mathfrak{v}}$ is just $\mathcal{S}_{\mathfrak{v}}$.
- $\mathcal{S}_{\mathfrak{s},\mathfrak{u}}(r_{\text{pub}}; r_{\mathfrak{s},\text{priv}})$ recomputes $q, Z, \text{ID}, \{I_{\text{id},\mathfrak{s}}\}_{\text{id} \in \text{ID}, \mathfrak{s} \in \{0,\mathfrak{g}\}}$ as in $\mathcal{S}_{\mathfrak{s},\mathfrak{v}}$, then uses $r_{\mathfrak{s},\text{priv}}$ to sample $\mathbf{U}_{\mathfrak{s},\text{id},\mathfrak{s}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{Z \times I_{\text{id},\mathfrak{s}}}$ for all $\text{id} \in \text{ID}$ and $\mathfrak{s} \in \{0,\mathfrak{g}\}$, and lastly outputs $\{\mathbf{U}_{\mathfrak{s},\text{id},\mathfrak{s}}^\top\}_{\text{id} \in \text{ID}, \mathfrak{s} \in \{0,\mathfrak{g}\}}$.

It can be readily verified that $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}_{\mathfrak{s}}}$ follows from $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}}$:

$$\begin{aligned} & \left(r_{\text{pub}}, \left\{ \begin{array}{l} \mathbf{U}_{\mathfrak{s},\text{id},0}^\top \mathbf{V}_{\text{id},0} + \mathbf{E}_{\text{id},0}^\top \\ \mathbf{U}_{\mathfrak{s},\text{id},\mathfrak{g}}^\top \mathbf{V}_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes \mathbf{E}_{\text{id},\mathfrak{g}\mathfrak{g}}^\top + \mathbf{E}_{\text{id},\mathfrak{g}0}^\top \end{array} \right\}_{\text{id} \in \text{ID}} \right) \\ & \equiv \left(r_{\text{pub}}, \left\{ \begin{array}{l} (\mathbf{U}_{\mathfrak{s},\text{id},0} + \mathbf{U}_{\text{id},0})^\top \mathbf{V}_{\text{id},0} + \mathbf{E}_{\text{id},0}^\top \\ (\mathbf{U}_{\mathfrak{s},\text{id},\mathfrak{g}} + \mathbf{U}_{\text{id},\mathfrak{g}})^\top \mathbf{V}_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes \mathbf{E}_{\text{id},\mathfrak{g}\mathfrak{g}}^\top + \mathbf{E}_{\text{id},\mathfrak{g}0}^\top \end{array} \right\}_{\text{id} \in \text{ID}} \right) \\ & \approx \left(r_{\text{pub}}, \left\{ \begin{array}{l} \mathbf{U}_{\mathfrak{s},\text{id},0}^\top \mathbf{V}_{\text{id},0} + \mathbf{\Delta}_{\text{pre},\text{id},0}^\top \\ \mathbf{U}_{\mathfrak{s},\text{id},\mathfrak{g}}^\top \mathbf{V}_{\text{id},\mathfrak{g}} + \mathbf{\Delta}_{\text{pre},\text{id},\mathfrak{g}}^\top \end{array} \right\}_{\text{id} \in \text{ID}} \right) \equiv \left(r_{\text{pub}}, \left\{ \begin{array}{l} \mathbf{\Delta}_{\text{pre},\text{id},0}^\top \\ \mathbf{\Delta}_{\text{pre},\text{id},\mathfrak{g}}^\top \end{array} \right\}_{\text{id} \in \text{ID}} \right). \end{aligned}$$

Applying the initial part of the proof to $\mathcal{S}_{\mathfrak{s}}$, we obtain

$$\begin{aligned} & (r''_{\text{pub}}, \text{impk}', \text{isk}'_1, \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{U}'_{\mathfrak{s},1,\mathfrak{s}})^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}) \\ & \approx (r''_{\text{pub}}, \text{impk}', \text{isk}'_1, \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, \mathbf{\Delta}_{\mathfrak{s}}^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}), \end{aligned}$$

where $\mathbf{U}'_{\mathfrak{s},1,\mathfrak{s}}$ is to $\mathbf{U}_{\mathfrak{s},\text{id},\mathfrak{s}}$'s as $\mathbf{U}'_{1,\mathfrak{s}}$ is to $\mathbf{U}_{\text{id},\mathfrak{s}}$'s for each $\mathfrak{s} \in \{0,\mathfrak{g}\}$. Therefore,

$$\begin{aligned} & (r''_{\text{pub}}, \text{impk}', \text{isk}'_1, \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{U}'_{\mathfrak{s},1,\mathfrak{s}})^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}) \\ & \approx (r''_{\text{pub}}, \text{impk}', \text{isk}'_1, \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{U}'_{1,\mathfrak{s}})^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}). \end{aligned}$$

Opening up $\text{IPFEsec}_{\text{post}}^{\mathcal{S}}$, we see that

$$\begin{aligned} & \underbrace{r_{\text{pub}}}_{\text{part of } r''_{\text{pub}}}, \quad \underbrace{\text{impk}}_{\text{from } r_{\text{pub}}} = (q, 1^K, 1^Z, 1^n, \text{impk}'), \\ & \{\text{ict}_{\text{id},\mathfrak{s}}^*\}_{\text{id} \in \text{ID}, \mathfrak{s} \in \{0,\mathfrak{g}\}} = \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{U}'_{1,\mathfrak{s}})^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}, \\ & \{\text{ict}_{\text{id},\mathfrak{s}}^\mathfrak{s}\}_{\text{id} \in \text{ID}, \mathfrak{s} \in \{0,\mathfrak{g}\}} = \{\text{Enc}'(\text{impk}', 1, \mathfrak{s}, (\mathbf{U}'_{\mathfrak{s},1,\mathfrak{s}})^\top)\}_{\mathfrak{s} \in \{0,\mathfrak{g}\}}, \\ & \{\text{isk}_{\text{id}}\}_{\text{id} \in \text{ID}} = (\underbrace{\{\mathbf{\Psi}_{\text{id},0}, \mathbf{\Psi}_{\text{id},\mathfrak{g}}\}_{\text{id} \in \text{ID}}}_{\text{from } r'_{\text{pub}} \text{ in } r''_{\text{pub}}}, \text{isk}'_1), \end{aligned}$$

so it follows from the previous indistinguishability.

For the part regarding restricted security, it remains to verify that \mathcal{S}' is restricted whenever \mathcal{S} is restricted (below), and that $\mathcal{S}_{\mathfrak{s}}$ is always restricted (clear by inspection). Let \mathcal{S}_{A0} be the algorithm in Definition 6 for \mathcal{S} . The $\mathcal{S}'_{A0}(r''_{\text{pub}})$ for \mathcal{S}' parses r''_{pub} as $(r_{\text{pub}}, r'_{\text{pub}})$, recomputes $\text{ID}, \{I_{\text{id},\mathfrak{s}}\}_{\text{id} \in \text{ID}, \mathfrak{s} \in \{0,\mathfrak{g}\}}, I'_{1,0}, I'_{1,\mathfrak{g}}, n$ as in $\mathcal{S}_{\mathfrak{v}}$, runs and sets

$$(1^{n'}, \{\mathbf{A}_{\text{id},\mathfrak{s},0,i}\}_{\text{id} \in \text{ID}, \mathfrak{s} \in \{0,\mathfrak{g}\}, i \in [I_{\text{id},\mathfrak{s}}]}) \leftarrow \mathcal{S}_{A0}(r_{\text{pub}}), \quad n'' = n' + n(I'_{1,0} + I'_{1,\mathfrak{g}}),$$

$$\mathbf{A}'_{1,\mathfrak{o},0,i'_\mathfrak{o}} = \begin{pmatrix} \mathbf{A}_{\text{id},\mathfrak{o},0,i} & & \\ & \mathbf{0}_{n(i'_\mathfrak{o}-1) \times n} & \\ & \text{id} \cdot \mathbf{I}_n & -\mathbf{I}_n \\ & & & \mathbf{0}_{n(I'_{1,\mathfrak{o}}-i'_\mathfrak{o}+I'_{1,\mathfrak{g}}) \times n} \end{pmatrix}, \quad \mathbf{A}'_{1,\mathfrak{g},0,i'_\mathfrak{g}} = \begin{pmatrix} \mathbf{A}_{\text{id},\mathfrak{g},0,i} & & \\ & \mathbf{0}_{n(I'_{1,\mathfrak{o}}+i'_\mathfrak{g}-1) \times n} & \\ & \text{id} \cdot \mathbf{I}_n & -\mathbf{I}_n \\ & & & \mathbf{0}_{n(I'_{1,\mathfrak{g}}-i'_\mathfrak{g}) \times n} \end{pmatrix},$$

with (id, i) being the $(i'_s)^{\text{th}}$ item among all such that $\text{id} \in \text{ID}$ and $i \in [I_{\text{id},s}]$. It outputs $(1^{n''}, \{\mathbf{A}'_{1,s,0,i'_s}\}_{s \in \{\mathfrak{o}, \mathfrak{g}\}, i'_s \in [I'_{1,s}]})$. Conceptually, let

$$(\mathbf{d}'_0)^\top = (\mathbf{d}'_0{}^\top, \underbrace{\dots, \boldsymbol{\varphi}_{\text{id},\mathfrak{o},i}^\top, \dots}_{\text{id} \in \text{ID}, i \in [I_{\text{id},\mathfrak{o}]}, \dots, \underbrace{\boldsymbol{\varphi}_{\text{id},\mathfrak{g},i}^\top, \dots}_{\text{id} \in \text{ID}, i \in [I_{\text{id},\mathfrak{g}]}}),$$

then

$$\begin{aligned} (\mathbf{u}'_{1,s,i'_s})^\top &= (\mathbf{u}'_{\text{id},s,i}{}^\top, \text{id} \cdot \boldsymbol{\varphi}_{\text{id},s,i}^\top, -\boldsymbol{\varphi}_{\text{id},s,i}^\top) \\ &= (\mathbf{d}'_0{}^\top \mathbf{A}_{\text{id},s,0,i}, \text{id} \cdot \boldsymbol{\varphi}_{\text{id},s,i}^\top, -\boldsymbol{\varphi}_{\text{id},s,i}^\top) = (\mathbf{d}'_0)^\top \mathbf{A}'_{1,s,0,i'_s}. \end{aligned} \quad \square$$

Proof (Claim 12). Consider the following hybrids.

- H_0 is the first distribution in $\text{IPFSec}_{\text{pre}}^{S'}$, i.e.,

$$r'_{\text{pub}}, \quad (\mathbf{U}'_{1,\mathfrak{o}})^\top \mathbf{V}'_{1,\mathfrak{o}} + (\mathbf{E}'_{1,\mathfrak{o}})^\top, \quad (\mathbf{U}'_{1,\mathfrak{g}})^\top \mathbf{V}'_{1,\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes (\mathbf{E}'_{1,\mathfrak{g}\mathfrak{g}})^\top + (\mathbf{E}'_{1,\mathfrak{g}\mathfrak{o}})^\top.$$

The components can be rearranged into

$$\begin{aligned} r'_{\text{pub}}, \quad r'_{\text{pub}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{o}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{o}} + (\text{id} - \text{id}') \boldsymbol{\Phi}'_{\text{id},\mathfrak{o}}{}^\top \boldsymbol{\Psi}'_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{o}})^\top\}_{\text{id}, \text{id}' \in \text{ID}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{g}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{g}} + (\text{id} - \text{id}') \boldsymbol{\Phi}'_{\text{id},\mathfrak{g}}{}^\top \boldsymbol{\Psi}'_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{g}})^\top + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{o}})^\top\}_{\text{id}, \text{id}' \in \text{ID}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{o}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{o}})^\top\}_{\text{id} \in \text{ID}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{g}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{g}})^\top + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{o}})^\top\}_{\text{id} \in \text{ID}}, \end{aligned}$$

where

$$\mathbf{E}'_{1,e} = \begin{pmatrix} \mathbf{E}'_{*,*,e} & \mathbf{E}'_{\text{id},*,e} & \cdots \\ \mathbf{E}'_{*,\text{id}',e} & \mathbf{E}'_{\text{id},\text{id}',e} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}_{\text{id}', \text{id} \in \text{ID}} \quad \text{for all } e \in \{\mathfrak{o}, \mathfrak{g}\mathfrak{g}, \mathfrak{g}\mathfrak{o}\}$$

with $\mathbf{E}'_{\text{id},\text{id}'\mathfrak{o}} \in \mathbb{Z}^{J_{\text{id}'} \times I_{\text{id},\mathfrak{o}}}$, $\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{g}} \in \mathbb{Z}^{J_{\text{id}'} \times I_{\text{id},\mathfrak{g}}}$, $\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{o}} \in \mathbb{Z}^{K_{J_{\text{id}'}} \times I_{\text{id},\mathfrak{g}}}$ for all $\text{id}, \text{id}' \in \text{ID}$.

- In H_1 , we insert noises \mathbf{E}^{help} into the inner products for \mathfrak{o} between non-matching identities, i.e.,

$$\begin{aligned} r'_{\text{pub}}, \quad r'_{\text{pub}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{o}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{o}} + (\text{id} - \text{id}') \boldsymbol{\Phi}'_{\text{id},\mathfrak{o}}{}^\top \boldsymbol{\Psi}'_{\text{id},\mathfrak{o}} + (\mathbf{E}^{\text{help}}_{\text{id},\text{id}'})^\top + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{o}})^\top\}_{\text{id}, \text{id}' \in \text{ID}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{g}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{g}} + (\text{id} - \text{id}') \boldsymbol{\Phi}'_{\text{id},\mathfrak{g}}{}^\top \boldsymbol{\Psi}'_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{g}})^\top + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{o}})^\top\}_{\text{id}, \text{id}' \in \text{ID}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{o}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{o}})^\top\}_{\text{id} \in \text{ID}}, \\ \{\mathbf{U}'_{\text{id},\mathfrak{g}}{}^\top \mathbf{V}'_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^\top \otimes (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{g}})^\top + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{o}})^\top\}_{\text{id} \in \text{ID}}, \end{aligned}$$

where $\mathbf{E}^{\text{help}}_{\text{id},\text{id}'} \stackrel{s}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{help}}, \leq \sigma_{\text{help}} \sqrt{\kappa}}^{J_{\text{id}'} \times I_{\text{id},\mathfrak{o}}}$, each of which is flooded by $\mathbf{E}'_{\text{id},\text{id}'\mathfrak{o}}$.

By Lemma 2, we have $H_0 \approx_s H_1$.

- In $H_{2,\alpha}$ (with $0 \leq \alpha \leq |\text{ID}|$), we no longer truncate \mathbf{E}^{help} , and change the non-matching inner products with the ciphertexts under the first α identities for \mathfrak{o} to random, i.e.,

$$\begin{aligned}
& r_{\text{pub}}, \quad r'_{\text{pub}}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{o}}^{\text{T}} \mathbf{V}_{\text{id}'\mathfrak{o}} + (\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{o}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}', \text{id} \leq \text{id}_{\alpha}}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{o}}^{\text{T}} \mathbf{V}_{\text{id}'\mathfrak{o}} + (\text{id} - \text{id}') \mathbf{\Phi}_{\text{id},\mathfrak{o}}^{\text{T}} \mathbf{\Psi}_{\text{id}'\mathfrak{o}} + (\mathbf{E}_{\text{id},\text{id}'}^{\text{help}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}', \text{id} > \text{id}_{\alpha}}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{V}_{\text{id}'\mathfrak{g}} + (\text{id} - \text{id}') \mathbf{\Phi}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{\Psi}_{\text{id}'\mathfrak{g}} + \widetilde{\mathbf{g}}^{\text{T}} \otimes (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{g}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{o}}^{\text{T}} \mathbf{V}_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{o}})^{\text{T}} \right\}_{\text{id} \in \text{ID}}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{V}_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^{\text{T}} \otimes (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{g}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{o}})^{\text{T}} \right\}_{\text{id} \in \text{ID}},
\end{aligned}$$

where $\mathbf{E}_{\text{id},\text{id}'}^{\text{help}} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{help}}}^{J_{\text{id}'} \times I_{\text{id},\mathfrak{o}}}$ and $\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{o}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{J_{\text{id}'} \times I_{\text{id},\mathfrak{o}}}$ for all $\text{id}, \text{id}' \in \text{ID}$. Here, $\text{id}_{\alpha} \in \text{ID}$ is the α^{th} -smallest among all $\text{id} \in \text{ID}$ in some fixed efficiently comparable total order for all $\alpha \in [|\text{ID}|]$, and id_0 is a sentinel value that compares less than all $\text{id} \in \text{ID}$.

By Lemma 1, we have $H_1 \approx_s H_{2,0}$. By a hybrid argument, from $\text{LWE}_{n,J_1,q,\sigma_{\text{help}}}$ it follows that $H_{2,\alpha} \approx H_{2,\alpha+1}$ for all $0 \leq \alpha < |\text{ID}|$, with the LWE public matrix being

$$\left(\dots, (\text{id}_{\alpha+1} - \text{id}') \mathbf{\Psi}_{\text{id}'\mathfrak{o}}, \dots \right)_{\text{id}' \in \text{ID}}^{\text{id}' \neq \text{id}_{\alpha+1}},$$

and the LWE secret being each column of $\mathbf{\Phi}_{\text{id},\mathfrak{o}}$. Roughly speaking, the reduction algorithm computes¹² $\mathbf{\Psi}_{\text{id}'\mathfrak{o}}$ for all $\text{id}' \in \text{ID}$ from the public matrices, and samples the other components by itself.

- In H_3 , we change all non-matching inner products for \mathfrak{o} to random, i.e.,

$$\begin{aligned}
& r_{\text{pub}}, \quad r'_{\text{pub}}, \quad \left\{ (\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{V}_{\text{id}'\mathfrak{g}} + (\text{id} - \text{id}') \mathbf{\Phi}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{\Psi}_{\text{id}'\mathfrak{g}} + \widetilde{\mathbf{g}}^{\text{T}} \otimes (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{g}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id}'\mathfrak{g}\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{o}}^{\text{T}} \mathbf{V}_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{o}})^{\text{T}} \right\}_{\text{id} \in \text{ID}}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{V}_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^{\text{T}} \otimes (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{g}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{o}})^{\text{T}} \right\}_{\text{id} \in \text{ID}}.
\end{aligned}$$

Comparing $H_{2,|\text{ID}|}$ and H_3 , we see that $(\mathbf{\Delta}')$'s absorb the terms added to them without changing the distribution, so $H_{2,|\text{ID}|} \equiv H_3$.

- In H_4 , we change all non-matching inner products for \mathfrak{g} to random, i.e.,

$$\begin{aligned}
& r_{\text{pub}}, \quad r'_{\text{pub}}, \quad \left\{ (\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'}, \quad \left\{ (\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{g}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{o}}^{\text{T}} \mathbf{V}_{\text{id},\mathfrak{o}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{o}})^{\text{T}} \right\}_{\text{id} \in \text{ID}}, \\
& \left\{ \mathbf{U}_{\text{id},\mathfrak{g}}^{\text{T}} \mathbf{V}_{\text{id},\mathfrak{g}} + \widetilde{\mathbf{g}}^{\text{T}} \otimes (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{g}})^{\text{T}} + (\mathbf{E}'_{\text{id},\text{id},\mathfrak{g}\mathfrak{o}})^{\text{T}} \right\}_{\text{id} \in \text{ID}}.
\end{aligned}$$

The transition from H_3 to H_4 is analogous to that from H_0 to H_3 , and $H_3 \approx H_4$.

- In H_5 , we alter the noises attached to the matching inner products into

$$r_{\text{pub}}, \quad r'_{\text{pub}}, \quad \left\{ (\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{o}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'}, \quad \left\{ (\mathbf{\Delta}'_{\text{id},\text{id}'\mathfrak{g}})^{\text{T}} \right\}_{\text{id},\text{id}' \in \text{ID}}^{\text{id} \neq \text{id}'},$$

¹²This step relies on q being a prime.

$$\begin{aligned} & \{\mathbf{U}_{\text{id},0}^\top \mathbf{V}_{\text{id},0} + \widetilde{\mathbf{E}}_{\text{id},0}^\top\}_{\text{id} \in \text{ID}}, \\ & \{\mathbf{U}_{\text{id},g}^\top \mathbf{V}_{\text{id},g} + \widetilde{\mathbf{g}}^\top \otimes \widetilde{\mathbf{E}}_{\text{id},gg}^\top + \widetilde{\mathbf{E}}_{\text{id},g0}^\top\}_{\text{id} \in \text{ID}}, \end{aligned}$$

where for all $\text{id} \in \text{ID}$ and $e \in \{0, gg, g0\}$,

$$\widetilde{\mathbf{E}}_{\text{id},e} = \begin{cases} \mathbf{E}_{\text{id},e}, & \text{if } \sigma_{\text{pre}} \geq \frac{\sigma_{\text{pre},0}}{2^{\kappa+6}\sqrt{\kappa}} \text{ so } \sigma'_{\text{pre}} = \sigma_{\text{pre}}; \\ \mathbf{E}_{\text{id},e} + \mathbf{E}'_{\text{id},e}, & \text{otherwise;} \end{cases}$$

$$\text{with } \mathbf{E}_{\text{id},0} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}, \leq \sigma_{\text{pre}} \sqrt{\kappa}}^{J_{\text{id}} \times I_{\text{id},0}}, \mathbf{E}_{\text{id},gg} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}, \leq \sigma_{\text{pre}} \sqrt{\kappa}}^{J_{\text{id}} \times I_{\text{id},g}}, \mathbf{E}_{\text{id},g0} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}, \leq \sigma_{\text{pre}} \sqrt{\kappa}}^{K J_{\text{id}} \times I_{\text{id},g}}.$$

In the first case, the two hybrids are statistically close by Lemma 1, and in the second case, by Lemma 2, as $\sigma'_{\text{pre}} = \sigma_{\text{pre},0} \geq 2^{\kappa+6} \sigma_{\text{pre}} \sqrt{\kappa}$. Therefore, $H_4 \approx_s H_5$.

- In H_6 , we no longer truncate $\mathbf{E}_{\text{id},e}$'s, i.e.,

$$\begin{aligned} & r_{\text{pub}}, \quad r'_{\text{pub}}, \quad \{(\Delta'_{\text{id},\text{id}'0})^\top\}_{\text{id},\text{id}' \in \text{ID}}, \quad \{(\Delta'_{\text{id},\text{id}'g})^\top\}_{\text{id},\text{id}' \in \text{ID}}, \\ & \{\mathbf{U}_{\text{id},0}^\top \mathbf{V}_{\text{id},0} + \mathbf{E}_{\text{id},0}^\top + (\mathbf{E}''_{\text{id},0})^\top\}_{\text{id} \in \text{ID}}, \\ & \{\mathbf{U}_{\text{id},g}^\top \mathbf{V}_{\text{id},g} + \widetilde{\mathbf{g}}^\top \otimes \mathbf{E}_{\text{id},gg}^\top + \mathbf{E}_{\text{id},g0}^\top + (\mathbf{E}''_{\text{id},g})^\top\}_{\text{id} \in \text{ID}}, \end{aligned}$$

where $\mathbf{E}_{\text{id},0} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}}^{J_{\text{id}} \times I_{\text{id},0}}$, $\mathbf{E}_{\text{id},gg} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}}^{J_{\text{id}} \times I_{\text{id},g}}$, $\mathbf{E}_{\text{id},g0} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}}^{K J_{\text{id}} \times I_{\text{id},g}}$, and $\mathbf{E}''_{\text{id},0}$ (resp. $\mathbf{E}''_{\text{id},g}$) is either $\mathbf{0}$ or $\mathbf{E}'_{\text{id},\text{id},0}$ (resp. $\widetilde{\mathbf{g}} \otimes \mathbf{E}'_{\text{id},\text{id},gg} + \mathbf{E}'_{\text{id},\text{id},g0}$) for all $\text{id} \in \text{ID}$.

By Lemma 1, we have $H_5 \approx_s H_6$.

- H_7 is the second distribution in $\text{IPFSec}_{\text{pre}}^{S'}$, i.e.,

$$\begin{aligned} & r_{\text{pub}}, \quad r'_{\text{pub}}, \quad \{(\Delta'_{\text{id},\text{id}'0})^\top\}_{\text{id},\text{id}' \in \text{ID}}, \quad \{(\Delta'_{\text{id},\text{id}'g})^\top\}_{\text{id},\text{id}' \in \text{ID}}, \\ & \{(\Delta'_{\text{id},\text{id},0})^\top\}_{\text{id} \in \text{ID}}, \quad \{(\Delta'_{\text{id},\text{id},g})^\top\}_{\text{id} \in \text{ID}}, \end{aligned}$$

where

$$\Delta'_{1,s} = \begin{pmatrix} \Delta'_{*,*,s} & \Delta'_{\text{id},*,s} & \cdots \\ \Delta'_{*,\text{id}',s} & \Delta'_{\text{id},\text{id}',s} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}_{\text{id}', \text{id} \in \text{ID}} \quad \text{for all } s \in \{0, g\}$$

with $\Delta'_{\text{id},\text{id},0} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{J_{\text{id}'} \times I_{\text{id},0}}$, $\Delta'_{\text{id},\text{id},g} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{K J_{\text{id}'} \times I_{\text{id},g}}$ for all $\text{id}, \text{id}' \in \text{ID}$.

$H_6 \approx H_7$ follows from $\text{IPFSec}_{\text{pre}}^S$, with $\Delta'_{\text{id},\text{id},s}$ being either $\Delta_{\text{id},s}$ or $(\Delta_{\text{id},s} - \mathbf{E}''_{\text{id},\text{id},s})$.

By hybrid argument, $H_0 \approx H_7$, which is exactly $\text{IPFSec}_{\text{pre}}^{S'}$. \square

3.5 Scheme with Structured Noises

To support structured noises, we modify Construction 1, fitting it into the shape of evasive learning with structured errors. We explain a subtlety in the adaptation. In the basic scheme, a key \mathbf{k} for $\mathbf{0}$ makes $\mathbf{d}^\top \mathbf{B} \mathbf{k}$ small (not pseudorandom). Therefore, in the presence of such a key, no ciphertext can be provably secure from evasive LWE. This is not a problem for Construction 1, because every (plaintext) vector has non-pseudorandom noisy inner product with $\mathbf{0}$. Similar properties of keys are not allowed in a scheme with structured noises, as security should hold for a ciphertext of random vector for \mathbf{g} , given a key for $\mathbf{v}_0 = \mathbf{0}$ and random \mathbf{V}_g . Preventing noise structure mix-and-match is akin to having two identities, for which we employ a mechanism similar to yet simpler than that of Construction 2.

Ingredients of Construction 3. We rely on Lemma 3 for correctness. For security, we need LWE (Assumption 1) and evasive learning with structured errors (Assumption 3).

Construction 3 (evIPFE-w/sn). Our scheme works as follows.

- $\text{Setup}(q, 1^K, 1^Z)$ takes as input q, K, Z . It sets $\mathcal{I}_{q,K,Z} = \{1\}$, and *deterministically* picks n and $m \geq m_0(n(K+1), q)$ and $\sigma_{-1} \geq \sigma_0(n, m)$ appropriate for Lemma 3, as well as $\sigma_{\text{post}}, \kappa$. The algorithm samples and sets

$$\begin{aligned} (\tilde{\mathbf{B}}, \tau) &\stackrel{\$}{\leftarrow} \text{TrapGen}(1^{n(K+1)}, 1^m, q), & \mathbf{A} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times 2(Z+n) \lceil \log_2 q \rceil}, \\ \text{impk} &= (1^n, 1^m, q, 1^K, 1^Z, \sigma_{-1}, \sigma_{\text{post}}, 1^\kappa, \mathbf{A}, \mathbf{B}), & \text{imsk} &= (\text{impk}, \tau). \end{aligned}$$

where $\mathbf{B} = (\mathbf{B}_{-1}, \dots, \mathbf{B}_{K-1})$ for $\tilde{\mathbf{B}} = (\mathbf{B}_{-1}^\top, \dots, \mathbf{B}_{K-1}^\top)^\top$. It outputs $(\text{impk}, \text{imsk})$.

- $\text{KeyGen}(\text{imsk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g)$ $\text{imsk}, \text{id} = 1, \mathbf{v}_0 \in \mathbb{Z}_q^Z$, and $\mathbf{V}_g \in \mathbb{Z}_q^{Z \times K}$. It samples and runs

$$\begin{aligned} \boldsymbol{\psi}_0 &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^n, & \boldsymbol{\Psi}_g &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times K}, \\ \mathbf{K} &\stackrel{\$}{\leftarrow} \text{SampleD}(\tilde{\mathbf{B}}, \tau, \tilde{\mathbf{P}}, \sigma_{-1}), \end{aligned}$$

where $\tilde{\mathbf{P}} = (\mathbf{P}_{-1}^\top, \dots, \mathbf{P}_{K-1}^\top)^\top$ for

$$\mathbf{P} = (\mathbf{P}_{-1}, \dots, \mathbf{P}_{K-1}) = \mathbf{A}\mathbf{G}^{-1} \begin{pmatrix} \mathbf{v}_0 \\ \boldsymbol{\psi}_0 \\ \mathbf{V}_g \\ \boldsymbol{\Psi}_g \end{pmatrix}.$$

The algorithm outputs $\text{isk} = (\boldsymbol{\psi}_0, \boldsymbol{\Psi}_g, \mathbf{K})$.

- $\text{Enc}(\text{impk}, \text{id}, \mathfrak{s}, \mathbf{u}^\top)$ takes as input $\text{impk}, \text{id} = 1, \mathfrak{s} \in \{\mathfrak{o}, \mathfrak{g}\}$, and $\mathbf{u} \in \mathbb{Z}_q^Z$. It samples and sets

$$\begin{aligned} \boldsymbol{\varphi} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^n, & \mathbf{d} &\stackrel{\$}{\leftarrow} \mathbb{Z}_q^n, & \mathbf{e}_0 &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}, \leq \sigma_{\text{post}} \sqrt{\kappa}}^{2(Z+n) \lceil \log_2 q \rceil}, \\ \mathbf{e}_1 &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}, \leq \sigma_{\text{post}} \sqrt{\kappa}}^m, & \mathbf{e}_3 &\stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}, \leq \sigma_{\text{post}} \sqrt{\kappa}}^m, \\ (\mathbf{u}')^\top &\leftarrow \begin{cases} (\mathbf{u}^\top, \mathbf{0}_{1 \times n}, \mathbf{0}_{1 \times Z}, \boldsymbol{\varphi}^\top), & \text{if } \mathfrak{s} = \mathfrak{o}; \\ (\mathbf{0}_{1 \times Z}, \boldsymbol{\varphi}^\top \mathbf{G}, \mathbf{u}^\top \mathbf{G}, \mathbf{0}_{1 \times n}), & \text{if } \mathfrak{s} = \mathfrak{g}. \end{cases} \end{aligned}$$

The algorithm outputs $\text{ict} = (\mathbf{d}^\top \mathbf{B} + (\mathbf{0}_{1 \times m}, \tilde{\mathbf{g}}^\top \otimes \mathbf{e}_3^\top) + \mathbf{e}_1^\top, \mathbf{d}^\top \mathbf{A} + (\mathbf{u}')^\top \mathbf{G} + \mathbf{e}_0^\top)$.

- $\text{Dec}(\text{impk}, \text{id}, \mathbf{v}_0, \mathbf{V}_g, \text{isk}, \mathfrak{s}, \text{ict})$ computes and outputs

$$\text{ict} \cdot \begin{pmatrix} -\mathbf{I}_{K+1} \otimes \mathbf{K} \\ \mathbf{G}^{-1} \begin{pmatrix} \mathbf{v}_0 \\ \boldsymbol{\psi}_0 \\ \mathbf{V}_g \\ \boldsymbol{\Psi}_g \end{pmatrix} \end{pmatrix} \cdot \begin{cases} (1, \mathbf{0}_{1 \times K})^\top, & \text{if } \mathfrak{s} = \mathfrak{o}; \\ (\mathbf{0}_{K \times 1}, \mathbf{I}_K)^\top, & \text{if } \mathfrak{s} = \mathfrak{g}. \end{cases}$$

Correctness. It is readily verified, similarly to Construction 1, that the scheme is B -correct for

$$B = \kappa^{1/2} \sigma_{\text{post}} (m^{3/2} \sigma_{-1} + 2(Z+n) \lceil \log_2 q \rceil).$$

Theorem 13. Construction 3 is restricted- σ_{post} -secure (Definition 6) under evasive learning with structured errors (Assumption 3) and $\text{LWE}_{n,m(K+1)+2(Z+n)\lceil \log_2 q \rceil,q,\sigma_{\text{help}}}$ (Assumption 1) for some $\sigma_{\text{help}} \leq \frac{\sigma_{\text{post}}}{(2^{\kappa+6}\sqrt{\kappa})^2 \cdot 2 \cdot (Z+n) \lceil \log_2 q \rceil}$, where q, K, Z are those picked by S and $n, m, \sigma_{\text{post}}, \kappa$ are those picked by Setup.

Proof Sketch. The proof done by combining the techniques demonstrated earlier in this section. Roughly speaking, the desired can be cast as a postcondition of evLWSE. In the precondition of evLWSE, first use IPFE simulation as in the proof of Theorem 6 so that we only consider the inner products, both for matching and non-matching noise structures. Then, use the proof of Claim 12 to argue that non-matching inner products, due to the presence of $\psi^\top \phi_0$ or $\psi^\top \Phi_0$, are pseudorandom. Lastly, use the premise that the matching inner products are pseudorandom to conclude the proof. \square

Applying Construction 2 to Construction 3, we obtain the following:

Corollary 14. Under $\text{LWE}_{n,\text{poly}(\lambda),q,\sigma_{\text{help}}}$ (Assumption 1) and evasive learning with structure errors (Assumption 3), there exists an IB-evIPFE-w/sn (Definition 4) with identity space $[q]$ and correctness bound

$$B = \kappa^{1/2} \sigma_{\text{post}} (m^{3/2} \sigma_{-1} + 2(Z+3n) \lceil \log_2 q \rceil)$$

that is restricted- σ_{post} -secure (Definition 6), where $\sigma_{\text{help}} \leq \frac{\sigma_{\text{post}}}{(2^{\kappa+6}\sqrt{\kappa})^2 \cdot 2 \cdot (Z+3n) \lceil \log_2 q \rceil}$.

4 Noisy Linear Garbling

We consider a notion of noisy linear garbling with authenticity (i.e., conditional disclosure of secrets)¹³ reusable for multiple inputs. Our formulation characterizes a *promise* primitive. Given a function $f : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}$, the garbling consists of labels decomposable and affine in the input \mathbf{x} , a garbled table, a secret, and some auxiliary information. For correctness, if $f(\mathbf{x}) = 1$, the secret can be approximately recovered from the other information; for security, if $f(\mathbf{x}) = 0$, the secret remains hidden; when $f(\mathbf{x}) = \perp$, we require neither correctness nor security. The *promise* feature of the definition is used to exclude \mathbf{x} leading to out-of-bound wire values in arithmetic computations, and $\mathbf{x} \notin \{0, 1\}^*$ in Boolean computations.

Definition 7 (noisy linear garbling). Let $F = \{F_{\lambda,\text{param}}\}_{\lambda \in \mathbb{N}, \text{param} \in \text{Params}_\lambda}$ be a sequence of function families, where $\text{Params} = \{\text{Params}_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of function family description sets and each $f \in F_{\lambda,\text{param}}$ is a function $\mathbb{Z}^L \rightarrow \{0, 1, \perp\}$ (with potentially different L for each f). A *noisy linear garbling scheme* for F consists of three efficient algorithms.

¹³The terminology is changed from “secret sharing” in introduction and technical overview, because the policy function might not be monotone.

- $\text{Setup}(1^\lambda, \text{param})$ takes the function family description $\text{param} \in \text{Params}_\lambda$ as input. It outputs some public parameter $\text{pp} = (q, \dots)$, where $q \in \mathbb{N}_{\geq 2}$ is the modulus.¹⁴
- $\text{GenF}(1^\lambda, \text{pp}, f)$ takes as input pp and $f \in F_{\lambda, \text{param}} : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}$. It outputs a deterministic dimension 1^{n_f} of garbling randomness, a secret vector \mathbf{w}_{out} , label functions $\{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}$, a garbled table \mathbf{T} , and some reusable information \mathbf{R} . Here, \mathbf{w}_{out} , \mathbf{W} 's, and \mathbf{T} are n_f -row matrices over \mathbb{Z} , and their shapes, including n_f , are fully determined by pp, f .
- $\text{Eval}(1^\lambda, \text{pp}, f, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)$ takes as input $\text{pp}, f, \mathbf{R}, \mathbf{x}$ (to f), one set of randomized labels, and the randomized garbled table \mathbf{t} . If $f(\mathbf{x}) = 1$, it is supposed to output an approximation of the randomized secret w_{out} . Here, $w_{\text{out}}, \{\mathbf{w}_\ell\}_{\ell \in [L]}, \mathbf{t}$ are over \mathbb{Z}_q .

Let B_{in} and B_{out} be functions mapping (λ, param) to natural numbers. The scheme is $(B_{\text{in}}, B_{\text{out}})$ -correct if for all $\lambda \in \mathbb{N}$, $\text{param} \in \text{Params}_\lambda$, $f \in F_{\lambda, \text{param}} : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}$, $\mathbf{x} \in \mathbb{Z}^L$, $e_{\text{out}}, \mathbf{e}_1, \dots, \mathbf{e}_L, \mathbf{e}_t \in [-B_{\text{in}}(\lambda, \text{param}), B_{\text{in}}(\lambda, \text{param})]^*$ (of suitable dimensions)¹⁵ such that $f(\mathbf{x}) = 1$, it holds that

$$\Pr \left[\begin{array}{l} \text{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, \text{param}) \\ \left(\begin{array}{l} 1^{n_f}, \mathbf{w}_{\text{out}}, \\ \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \\ \mathbf{T}, \mathbf{R} \end{array} \right) \stackrel{\$}{\leftarrow} \text{GenF}(1^\lambda, \text{pp}, f) \\ \mathbf{s} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n_f} \end{array} \right] : \begin{array}{l} \text{Eval} \left(\begin{array}{l} 1^\lambda, \text{pp}, f, \mathbf{R}, \mathbf{x}, \\ \{\mathbf{s}^\top (\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_\ell^\top\}_{\ell \in [L]}, \\ \mathbf{s}^\top \mathbf{T} + \mathbf{e}_t^\top \end{array} \right) \\ \in [-B_{\text{out}}(\lambda, \text{param}), B_{\text{out}}(\lambda, \text{param})] \\ - (\mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}}) \end{array} \right] = 1.$$

Definition 8 (shortness). A noisy linear garbling scheme (Definition 7) is B_{short} -short (for some function B_{short} mapping (λ, param) to natural numbers) if for all $\lambda \in \mathbb{N}$, $\text{param} \in \text{Params}_\lambda$, $f \in F_{\lambda, \text{param}} : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}$, $\mathbf{x} \in \mathbb{Z}^L$ such that $f(\mathbf{x}) = 0$,

$$\Pr \left[\begin{array}{l} \text{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, \text{param}) \\ \left(\begin{array}{l} 1^{n_f}, \mathbf{w}_{\text{out}}, \\ \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \\ \mathbf{T}, \mathbf{R} \end{array} \right) \stackrel{\$}{\leftarrow} \text{GenF}(1^\lambda, \text{pp}, f) \end{array} \right] : \begin{array}{l} \|\mathbf{w}_{\text{out}}^\top\|, \|\mathbf{T}^\top\| \leq B_{\text{short}}(\lambda, \text{param}), \\ \|\mathbf{W}_{\ell,0}^\top + \mathbf{x}[\ell] \mathbf{W}_{\ell,\times}^\top\| \leq B_{\text{short}}(\lambda, \text{param}) \\ \text{for all } \ell \in [L] \end{array} \right] = 1.$$

Definition 9 (fixed randomness dimension). A noisy linear garbling scheme (Definition 7) has *fixed randomness dimension* if param is of the form $(q, 1^n, \dots)$ and $n_f = n$ always holds.

Security. We require that the randomized (by noisy linear combination) secret, labels, and garbled table be pseudorandom if $f(\mathbf{x}) = 0$.

Definition 10 (noisy linear garbling security). Let $(\text{Setup}, \text{GenF}, \text{Eval})$ be a noisy linear garbling scheme for F (Definition 7) and GenNoise an efficient algorithm with suitable input/output formats. The scheme is *GenNoise-secure* if $\text{Exp}_{\text{garble}}^{\text{real}} \approx \text{Exp}_{\text{garble}}^{\text{random}}$, where $\text{Exp}_{\text{garble}}^{\text{real}}(1^\lambda, \mathcal{A})$ and $\text{Exp}_{\text{garble}}^{\text{random}}(1^\lambda, \mathcal{A})$ proceed as follows.

¹⁴In the conference version [HLL24], the definition requires q be deterministic. While that version still works (by smuggling q into param and checking its suitability in Setup), it was an unintended error – it is easier and more consistent (with the other definitions in this paper) to let Setup sample q instead. This is currently necessary since we need prime q and it is unknown [TCH12] how to generate large primes in deterministic polynomial time.

¹⁵The dimensions could depend on the randomness of Setup . Formally, they can be quantified over sufficiently large dimensions then truncated appropriately.

- **Setup.** Launch $\mathcal{A}(1^\lambda)$ and receive $\text{param} \in \text{Params}_\lambda$ from it. Sample r_{Setup} , run

$$\text{pp} \leftarrow \text{Setup}(1^\lambda, \text{param}; r_{\text{Setup}}),$$

and send r_{Setup} to \mathcal{A} .

- **Challenge.** \mathcal{A} chooses $f \in F_{\lambda, \text{param}} : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}$ and $\mathbf{x} \in \mathbb{Z}^L$. Sample r_{GenF} and run

$$(1^{nf}, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}) \leftarrow \text{GenF}(1^\lambda, \text{pp}, f; r_{\text{GenF}}).$$

Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^{nf}$, run

$$(e_{\text{out}}, \mathbf{e}_1, \dots, \mathbf{e}_L, \mathbf{e}_t) \xleftarrow{\$} \text{GenNoise}(1^\lambda, \text{param}, \text{pp}, f, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}, \mathbf{x}),$$

and compute

$$\begin{cases} w_{\text{out}} \leftarrow \mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}}, & \mathbf{w}_\ell^\top \leftarrow \mathbf{s}^\top (\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_\ell^\top, & \mathbf{t}^\top \leftarrow \mathbf{s}^\top \mathbf{T} + \mathbf{e}_t^\top, & \text{in } \text{Exp}_{\text{garble}}^{\text{real}}; \\ w_{\text{out}} \xleftarrow{\$} \mathbb{Z}_q, & \mathbf{w}_\ell \xleftarrow{\$} \mathbb{Z}_q^* \text{ (for all } \ell \in [L]), & \mathbf{t} \xleftarrow{\$} \mathbb{Z}_q^*, & \text{in } \text{Exp}_{\text{garble}}^{\text{random}}; \end{cases}$$

where $\{\mathbf{w}_\ell\}_{\ell \in [L]}$ and \mathbf{t} in $\text{Exp}_{\text{garble}}^{\text{random}}$ are of the same dimensions as their counterparts in $\text{Exp}_{\text{garble}}^{\text{real}}$. Send $(r_{\text{GenF}}, w_{\text{out}}, \{\mathbf{w}_\ell\}_{\ell \in [L]}, \mathbf{t})$ to \mathcal{A} .

- **Guess.** \mathcal{A} outputs $\beta' \in \{0, 1\}$. The output of the experiment is β' if $f(\mathbf{x}) = 0$. Otherwise, the output is set to 0.

5 Noisy Linear Garbling for Circuits

In this section, we construct a noisy linear garbling scheme for arithmetic circuits of bounded wire values. It is a variant of [AIK11, IW14, LL20a]. The garbling scheme serves as the central component in our CP-ABE for bounded-arithmetic circuits. (The notations have undergone some major changes since the conference version [HLL24] – see Remark 4.)

Construction 4 (noisy linear garbling for bounded-arithmetic circuits). The function family of bounded-arithmetic circuits is

$$\begin{aligned} \text{Params} &= \{ (M, 1^d) \mid M \in \mathbb{N}_+ \text{ and } d \in \mathbb{N} \}, \\ F_{M, 1^d} &= \{ f_{M,C} \mid C \text{ is an arithmetic circuit of depth no more than } d \}, \\ f_{M,C}(\mathbf{x}) &= \begin{cases} \perp, & \text{if some wire value } \notin [-M, M] \text{ when evaluating } C(\mathbf{x}); \\ 1, & \text{if } C(\mathbf{x}) = 0 \text{ and all wire values } \in [-M, M]; \\ 0, & \text{if } C(\mathbf{x}) \neq 0 \text{ and all wire values } \in [-M, M]; \end{cases} \\ &\text{for } C : \mathbb{Z}^L \rightarrow \mathbb{Z} \text{ and } \mathbf{x} \in \mathbb{Z}^L. \end{aligned}$$

Since M is known from param , the function $f_{M,C}$ is simply represented by C . The scheme works as follows.

- $\text{Setup}(M, 1^d)$ picks suitable q, N and outputs $\text{pp} = (q, M, 1^d, 1^N)$. The constraints of q, N are specified in Theorems 15, 16, and 17.

- GenF(pp, C) generates the garbling for $C : \mathbb{Z}^L \rightarrow \mathbb{Z}$ in the following steps.
 - It parses C into L input gates $1, \dots, L$ and $(|C| - L)$ arithmetic gates $L + 1, \dots, |C|$. They are sorted in topological order so that gate $|C|$ is the output gate of C and the inputs to every non-input gate i are connected to gates $\text{gin}[i, 1], \text{gin}[i, 2] < i$.
 - For each gate i , the algorithm samples its shrunken label function

$$\mathbf{W}_{i,0} \xleftarrow{\$} \{0, 1\}^{N \times N}, \quad \mathbf{W}_{i,\times} \xleftarrow{\$} \{0, 1\}^{N \times N}.$$

We write $\mathbf{W}_{i,x} = \mathbf{W}_{i,0} + x\mathbf{W}_{i,\times}$, where $x \in [-M, M]$ is a potential output value of gate i , for the corresponding shrunken label (before randomization).

- For each non-input gate i , the algorithm samples $\widetilde{\mathbf{W}}_{\text{gin}[i,1],0}^{(i,1)}, \widetilde{\mathbf{W}}_{\text{gin}[i,1],\times}^{(i,1)}, \widetilde{\mathbf{W}}_{\text{gin}[i,2],0}^{(i,2)}, \widetilde{\mathbf{W}}_{\text{gin}[i,2],\times}^{(i,2)}$, which contribute to the expanded label function of each input to gate i . Here, the superscript indicates the purpose of the part (due to being an input to gate i), and the subscript indicates the matrix containing that part, so

$$\widetilde{\mathbf{W}}_{i',0} = (\dots, \widetilde{\mathbf{W}}_{i',0}^{(i,\gamma)}, \dots)_{\text{gin}[i,\gamma]=i'}, \quad \widetilde{\mathbf{W}}_{i',\times} = (\dots, \widetilde{\mathbf{W}}_{i',\times}^{(i,\gamma)}, \dots)_{\text{gin}[i,\gamma]=i'}.$$

Similarly, we write $\widetilde{\mathbf{W}}_{i',x} = \widetilde{\mathbf{W}}_{i',0} + x\widetilde{\mathbf{W}}_{i',\times}$ for the expanded label (before randomization) corresponding to a potential output $x \in [-M, M]$ of gate i' . Suppose the input values to gate i are x_1, x_2 , the goal is to compute the shrunken label $\mathbf{W}_{i,x}$ from the expanded labels $\widetilde{\mathbf{W}}_{\text{gin}[i,1],x_1}, \widetilde{\mathbf{W}}_{\text{gin}[i,2],x_2}$. Following [AIK11] (adapted for authenticity only), first sample $\mathbf{U}_i \xleftarrow{\$} \{0, 1\}^{N \times N}$.

- * If gate i is addition, set

$$\begin{aligned} \widetilde{\mathbf{W}}_{\text{gin}[i,1],0}^{(i,1)} &\leftarrow \mathbf{U}_i, & \widetilde{\mathbf{W}}_{\text{gin}[i,1],\times}^{(i,1)} &\leftarrow \mathbf{W}_{i,\times}, \\ \widetilde{\mathbf{W}}_{\text{gin}[i,2],0}^{(i,2)} &\leftarrow \mathbf{W}_{i,0} - \mathbf{U}_i, & \widetilde{\mathbf{W}}_{\text{gin}[i,2],\times}^{(i,2)} &\leftarrow \mathbf{W}_{i,\times}, \end{aligned}$$

$$\text{which are subject to } \widetilde{\mathbf{W}}_{\text{gin}[i,1],x_1}^{(i,1)} + \widetilde{\mathbf{W}}_{\text{gin}[i,2],x_2}^{(i,2)} = \mathbf{W}_{i,x_1+x_2}.$$

- * If gate i is multiplication, set

$$\begin{aligned} \widetilde{\mathbf{W}}_{\text{gin}[i,1],0}^{(i,1)} &\leftarrow \mathbf{U}_i, & \widetilde{\mathbf{W}}_{\text{gin}[i,1],\times}^{(i,1)} &\leftarrow \mathbf{W}_{i,\times}, \\ \widetilde{\mathbf{W}}_{\text{gin}[i,2],0}^{(i,2)} &\leftarrow \mathbf{W}_{i,0}, & \widetilde{\mathbf{W}}_{\text{gin}[i,2],\times}^{(i,2)} &\leftarrow -\mathbf{U}_i, \end{aligned}$$

$$\text{which are subject to } x_2 \widetilde{\mathbf{W}}_{\text{gin}[i,1],x_1}^{(i,1)} + \widetilde{\mathbf{W}}_{\text{gin}[i,2],x_2}^{(i,2)} = \mathbf{W}_{i,x_1x_2}.$$

- For the output gate $|C|$, sample the expanded label function

$$\widetilde{\mathbf{W}}_{|C|,0} \xleftarrow{\$} \{0, 1\}^{N \times 1}, \quad \widetilde{\mathbf{W}}_{|C|,\times} \xleftarrow{\$} \{0, 1\}^{N \times 1}.$$

Here, $\widetilde{\mathbf{W}}_{|C|,0}, \widetilde{\mathbf{W}}_{|C|,\times}$ are vectors (single-column matrices; still in upper case for consistency with the other components).

- The algorithm computes the garbled table, which enables conversion to the expanded label $\widetilde{\mathbf{W}}_{i',x}$ from the shrunken label $\mathbf{W}_{i',x}$ for each gate i' with output value x , akin to the key-shrinking gadget in [AIK11]. For each tuple

(i', i, γ) with $\text{gin}[i, \gamma] = i'$ (i.e., gate i' is the γ^{th} input to gate i), sample and set

$$\begin{aligned} \mathbf{R}_{i'}^{(i,\gamma)} &\stackrel{\$}{\leftarrow} \{0, 1\}^{N \times N}, \\ \mathbf{T}_{i',0}^{(i,\gamma)} &\leftarrow \mathbf{W}_{i',0} \mathbf{R}_{i'}^{(i,\gamma)} + \widetilde{\mathbf{W}}_{i',0}^{(i,\gamma)}, & \mathbf{T}_{i',\times}^{(i,\gamma)} &\leftarrow \mathbf{W}_{i',\times} \mathbf{R}_{i'}^{(i,\gamma)} + \widetilde{\mathbf{W}}_{i',\times}^{(i,\gamma)}. \end{aligned}$$

Similarly define $\mathbf{T}_{i',x}^{(i,\gamma)}$, then $\mathbf{T}_{i',x}^{(i,\gamma)} = \mathbf{W}_{i',x} \mathbf{R}_{i'}^{(i,\gamma)} + \widetilde{\mathbf{W}}_{i',x}^{(i,\gamma)}$. For the output gate, sample $\mathbf{R}_{|C|} \stackrel{\$}{\leftarrow} \{0, 1\}^{N \times 1}$ and set

$$\mathbf{T}_{|C|,0} \leftarrow \mathbf{W}_{|C|,0} \mathbf{R}_{|C|} + \widetilde{\mathbf{W}}_{|C|,0}, \quad \mathbf{T}_{|C|,\times} \leftarrow \mathbf{W}_{|C|,\times} \mathbf{R}_{|C|} + \widetilde{\mathbf{W}}_{|C|,\times}.$$

Here, $\mathbf{R}_{|C|}, \mathbf{T}_{|C|,0}, \mathbf{T}_{|C|,\times}$ are vectors (single-column matrices; still in upper case for consistency with the other components).

All the samplings are done in a straight-forward way.¹⁶ The algorithm outputs

$$\begin{aligned} 1^{nC} &= \mathbf{1}^N, & \mathbf{w}_{\text{out}} &= \widetilde{\mathbf{W}}_{|C|,0}, & \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \\ \mathbf{T} &= (\dots, \mathbf{T}_{i',0}^{(i,\gamma)}, \mathbf{T}_{i',\times}^{(i,\gamma)}, \dots)_{\text{gin}[i,\gamma]=i'}, & \mathbf{R} &= (\dots, \mathbf{R}_{i'}^{(i,\gamma)}, \dots)_{\text{gin}[i,\gamma]=i'}. \end{aligned}$$

- Eval(pp, $C, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top$) parses C, \mathbf{R} as in GenF and

$$\mathbf{t}^\top = (\dots, (\mathbf{t}_{i',0}^{(i,\gamma)})^\top, (\mathbf{t}_{i',\times}^{(i,\gamma)})^\top, \dots)_{\text{gin}[i,\gamma]=i'}.$$

It evaluates $C(\mathbf{x})$ for the wire values $\{x_i\}_{i \in [|C|]}$ of each gate i , namely,

$$x_i \leftarrow \begin{cases} \mathbf{x}[\ell], & \text{if gate } i \text{ is an input gate } (i = \ell \in [L]); \\ x_{\text{gin}[i,1]} + x_{\text{gin}[i,2]}, & \text{if gate } i \text{ is addition;} \\ x_{\text{gin}[i,1]} x_{\text{gin}[i,2]}, & \text{if gate } i \text{ is multiplication.} \end{cases}$$

The algorithm has the (shrunk) input labels $\{\mathbf{w}_\ell^\top\}_{\ell \in [L]}$ as input. It computes the label of each non-input gate in increasing order by its index i .

- The algorithm decrypts for the expanded label using

$$\begin{aligned} (\widetilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1)})^\top &\leftarrow (\mathbf{t}_{\text{gin}[i,1], x_{\text{gin}[i,1]}}^{(i,1)})^\top - \mathbf{w}_{\text{gin}[i,1]}^\top \mathbf{R}_{\text{gin}[i,1]}^{(i,1)}, \\ (\widetilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2)})^\top &\leftarrow (\mathbf{t}_{\text{gin}[i,2], x_{\text{gin}[i,2]}}^{(i,2)})^\top - \mathbf{w}_{\text{gin}[i,2]}^\top \mathbf{R}_{\text{gin}[i,2]}^{(i,2)}, \end{aligned}$$

where $\mathbf{t}_{\text{gin}[i,1], x_{\text{gin}[i,1]}}^{(i,1)} = \mathbf{t}_{\text{gin}[i,1],0}^{(i,1)} + x_{\text{gin}[i,1]} \mathbf{t}_{\text{gin}[i,1],\times}^{(i,1)}$ (similarly for $\mathbf{t}_{\text{gin}[i,2], x_{\text{gin}[i,2]}}^{(i,2)}$) can be computed from \mathbf{t} (input to Eval) and x 's.

- The algorithm recovers the shrunk label as

$$\mathbf{w}_i \leftarrow \begin{cases} \widetilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1)} + \widetilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2)}, & \text{if gate } i \text{ is addition;} \\ x_{\text{gin}[i,2]} \widetilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1)} + \widetilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2)}, & \text{if gate } i \text{ is multiplication.} \end{cases}$$

¹⁶Precisely speaking, r_{GenF} conditioned on the sampled matrices must be efficiently sampleable given those matrices (with negligible statistical error), e.g., when the matrices are just the bits read sequentially from r_{GenF} . This ensures that no sampled matrix has a known trapdoor, and is important because r_{GenF} is incorporated into the sampler's randomness in a security proof.

Lastly, the algorithm computes and outputs the secret

$$w_{\text{out}} \leftarrow \mathbf{t}_{|C|,x_{|C|}}^\top - \mathbf{w}_{|C|}^\top \mathbf{R}_{|C|}.$$

Here, $\mathbf{t}_{|C|,x_{|C|}}$ is a scalar (one-dimensional vector; still boldfaced for consistency with the other components).

Remark 4 (changes since the conference version [HLL24]). We list major changes for those who have read the conference version of Construction 4.

- In Params, the wire value bound M is now encoded in binary (instead of unary) for stronger functionality since it need not be $\text{poly}(\lambda)$ -bounded.
- The dimension of secret is N (instead of n, m) since it is for flipped LWE secret. It is now encoded in unary in pp to emphasize $\text{poly}(\lambda)$ -boundedness.
- Gates are named simply by their indices (instead of g subscript index).
- The notations i, i' are more consistent. When both appear in a context, gate i' is always an input to gate i , and i' never appears as component superscript. The notations in the proof are also improved for consistency.
- The expanded label function of the output gate is now $\widetilde{\mathbf{W}}$ (instead of $\widetilde{\mathbf{w}}$) to reduce ambiguity.

5.1 Correctness and Shortness

Theorem 15 (¶). *Construction 4 is $(B_{\text{in}}, B_{\text{out}})$ -correct (Definition 7) if*

$$(6NM^2)^{d+1} B_{\text{in}}(M, d) \leq B_{\text{out}}(M, d).$$

Proof (Theorem 15). Let $C : \mathbb{Z}^L \rightarrow \mathbb{Z}$ be an arithmetic circuit of depth at most d and $\mathbf{x} \in \mathbb{Z}^L$ an input with $C(\mathbf{x}) = 0$ such that all wire values in $C(\mathbf{x})$ are bounded by M . Let $\mathbf{s} \in \mathbb{Z}_q^N$ be a secret and $e_{\text{out}}, \mathbf{e}_1, \dots, \mathbf{e}_L, \mathbf{e}_t$ any B_{in} -bounded errors. We extend the notation of \mathbf{e} 's by $\mathbf{e}_i = \mathbf{w}_i - (\mathbf{s}^\top \mathbf{W}_{i,x_i})^\top$ for all $L < i \leq |C|$, where x_i 's are the wire values. The extension symbolically agrees with $\mathbf{e}_1, \dots, \mathbf{e}_L$.

We show by induction that during evaluation, $\|\mathbf{e}_i\| \leq (6NM^2)^{d_i} B_{\text{in}}$ for every gate i , where d_i is the depth of gate i (the input gates are of depth 0). The base case (for $\mathbf{e}_1, \dots, \mathbf{e}_L$) is by assumption. For the inductive case, let $L < i \leq |C|$. Assuming the induction hypothesis for all $i' < i$, then it applies to $i' = \text{gin}[i, \gamma] < i$ for both $\gamma \in \{1, 2\}$, and together with $d_{i'} \leq d_i - 1$, we have

$$\|\mathbf{e}_{i'}\| \leq (6NM^2)^{d_{i'}} B_{\text{in}} \leq (6NM^2)^{d_i-1} B_{\text{in}}.$$

By the definition of \mathbf{t} and how Eval proceeds,

$$\begin{aligned} (\widetilde{\mathbf{w}}_{i'}^{(i,\gamma)})^\top &= (\mathbf{s}^\top \mathbf{T}_{i',0}^{(i,\gamma)} + (\mathbf{e}_{t,i',0}^{(i,\gamma)})^\top) + x_{i'} (\mathbf{s}^\top \mathbf{T}_{i',x}^{(i,\gamma)} + (\mathbf{e}_{t,i',x}^{(i,\gamma)})^\top) - (\mathbf{s}^\top \mathbf{W}_{i',x_{i'}} + \mathbf{e}_{i'}^\top) \mathbf{R}_{i'}^{(i,\gamma)} \\ &= \mathbf{s}^\top \widetilde{\mathbf{W}}_{i',x_{i'}}^{(i,\gamma)} + (\mathbf{e}_{t,i',0}^{(i,\gamma)} + x_{i'} \mathbf{e}_{t,i',x}^{(i,\gamma)})^\top - \mathbf{e}_{i'}^\top \mathbf{R}_{i'}^{(i,\gamma)} = \mathbf{s}^\top \widetilde{\mathbf{W}}_{i',x_{i'}}^{(i,\gamma)} + (\widetilde{\mathbf{e}}_{i'}^{(i,\gamma)})^\top, \end{aligned}$$

where superscripts/subscripts follow the usual meaning of taking the relevant parts. It follows that the errors on the expanded labels are

$$\begin{aligned}\|\tilde{\mathbf{e}}_{i'}^{(i,\gamma)}\| &\leq \|\mathbf{e}_{t,i',0}^{(i,\gamma)}\| + |x_{i'}| \cdot \|\mathbf{e}_{t,i',x}^{(i,\gamma)}\| + \|(\mathbf{R}_{i'}^{(i,\gamma)})^\top\| \cdot \|\mathbf{e}_{i'}\| \\ &\leq B_{\text{in}} + M \cdot B_{\text{in}} + N \cdot (6NM^2)^{d_i-1} B_{\text{in}} \\ &\leq (N + M + 1)(6NM^2)^{d_i-1} B_{\text{in}}.\end{aligned}$$

We proceed to a case analysis.

- If gate i is addition,

$$\begin{aligned}\mathbf{e}_i^\top &= \overbrace{(\mathbf{s}^\top \tilde{\mathbf{W}}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1)} + (\tilde{\mathbf{e}}_{\text{gin}[i,1]}^{(i,1)})^\top) + (\mathbf{s}^\top \tilde{\mathbf{W}}_{\text{gin}[i,2],x_{\text{gin}[i,2]}}^{(i,2)} + (\tilde{\mathbf{e}}_{\text{gin}[i,2]}^{(i,2)})^\top)}^{\mathbf{w}_i = \tilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1)} + \tilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2)}} - \mathbf{s}^\top \mathbf{W}_{i,x_i} \\ &= (\mathbf{e}_{\text{gin}[i,1]}^{(i,1)} + \mathbf{e}_{\text{gin}[i,2]}^{(i,2)})^\top,\end{aligned}$$

$$\text{so } \|\mathbf{e}_i\| \leq 2(N + M + 1)(6NM^2)^{d_i-1} B_{\text{in}} \leq (6NM^2)^{d_i} B_{\text{in}}.$$

- If gate i is multiplication,

$$\begin{aligned}\mathbf{e}_i^\top &= \overbrace{x_{\text{gin}[i,2]} (\mathbf{s}^\top \tilde{\mathbf{W}}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1)} + (\tilde{\mathbf{e}}_{\text{gin}[i,1]}^{(i,1)})^\top) + (\mathbf{s}^\top \tilde{\mathbf{W}}_{\text{gin}[i,2],x_{\text{gin}[i,2]}}^{(i,2)} + (\tilde{\mathbf{e}}_{\text{gin}[i,2]}^{(i,2)})^\top)}^{\mathbf{w}_i = x_{\text{gin}[i,2]} \tilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1)} + \tilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2)}} - \mathbf{s}^\top \mathbf{W}_{i,x_i} \\ &= (x_{\text{gin}[i,2]} \mathbf{e}_{\text{gin}[i,1]}^{(i,1)} + \mathbf{e}_{\text{gin}[i,2]}^{(i,2)})^\top,\end{aligned}$$

$$\text{so } \|\mathbf{e}_i\| \leq (M + 1)(N + M + 1)(6NM^2)^{d_i-1} B_{\text{in}} \leq (6NM^2)^{d_i} B_{\text{in}}.$$

This completes the induction.

Applying the error bound to the output gate, we have

$$\|\mathbf{e}_{|C|}\| \leq (6NM^2)^{d_{|C|}} B_{\text{in}} \leq (6NM^2)^d B_{\text{in}}.$$

Note that $x_{|C|} = C(\mathbf{x}) = 0$ by assumption, and we have

$$\begin{aligned}&w_{\text{out}} - (\mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}}) \\ &= (\mathbf{t}_{|C|,0}^\top - \mathbf{w}_{|C|}^\top \mathbf{R}_{|C|}) - (\mathbf{s}^\top \mathbf{w}_{\text{out}} + e_{\text{out}}) \\ &= (\mathbf{s}^\top \mathbf{T}_{|C|,0} + \mathbf{e}_{t,|C|,0}^\top) - (\mathbf{s}^\top \mathbf{W}_{|C|,0} + \mathbf{e}_{|C|}^\top) \mathbf{R}_{|C|} - \mathbf{s}^\top \mathbf{w}_{\text{out}} - e_{\text{out}} \\ &= \mathbf{e}_{t,|C|,0}^\top - \mathbf{e}_{|C|}^\top \mathbf{R}_{|C|} - e_{\text{out}} + \underbrace{\mathbf{s}^\top (\mathbf{W}_{|C|,0} \mathbf{R}_{|C|} + \tilde{\mathbf{W}}_{|C|,0} - \mathbf{W}_{|C|,0} \mathbf{R}_{|C|} - \tilde{\mathbf{W}}_{|C|,0})}_{\mathbf{0}}.\end{aligned}$$

Therefore, the output error is bounded by

$$\begin{aligned}\|\mathbf{e}_{t,|C|,0}\| + \|\mathbf{R}_{|C|}^\top\| \cdot \|\mathbf{e}_{|C|}\| + |e_{\text{out}}| &\leq B_{\text{in}} + N \cdot (6NM^2)^d B_{\text{in}} + B_{\text{in}} \\ &\leq (6NM^2)^{d+1} B_{\text{in}} \leq B_{\text{out}}.\end{aligned}$$

□

Theorem 16 (¶). *Construction 4 is B_{short} -short (Definition 8) if*

$$B_{\text{short}}(M, d) \geq (N + M + 2)N.$$

Proof (Theorem 16). By definition, $\mathbf{w}_{\text{out}}^\top \in \{0, 1\}^{1 \times N}$, so

$$\|\mathbf{w}_{\text{out}}^\top\| \leq N \leq (N + M + 2)N \leq B_{\text{short}}.$$

For all $\ell \in [L]$, when $|\mathbf{x}[\ell]| \leq M$,

$$\|\mathbf{W}_{\ell, \mathbf{x}[\ell]}^\top\| \leq \|\mathbf{W}_{\ell, 0}^\top\| + M\|\mathbf{W}_{\ell, \times}^\top\| \leq (M + 1)N \leq (N + M + 2)N \leq B_{\text{short}},$$

because \mathbf{W} 's are in $\{0, 1\}^{N \times N}$. For \mathbf{T} , observe that each $\tilde{\mathbf{W}}$ is a signed sum of at most two matrices in $\{0, 1\}^{N \times N}$, so each $\tilde{\mathbf{W}}^\top$ has norm bound $2N$. Therefore,

$$\|\mathbf{T}^\top\| \leq \underbrace{M}_{\mathbf{R}^\top} \cdot \underbrace{N}_{\mathbf{W}^\top} + \underbrace{2N}_{\tilde{\mathbf{W}}^\top} = (M + 2)N \leq (N + M + 2)N \leq B_{\text{short}}. \quad \square$$

5.2 Security with Gaussian Noise

Theorem 17 (¶). *Suppose*

$$\text{GenGaussian}(M, 1^d, (q, M, 1^d, 1^N), C, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell, 0}, \mathbf{W}_{\ell, \times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}, \mathbf{x})$$

samples (truncated) Gaussian noises with appropriate width σ of suitable shape

$$(e_{\text{out}}, \mathbf{e}_1, \dots, \mathbf{e}_L, \mathbf{e}_t) \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma, \leq \sigma\sqrt{k}}^*$$

*then Construction 4 is GenGaussian-secure (Definition 10) if q is always a prime and $\text{FlipLWE}_{N, \text{poly}(\lambda), q, \sigma'}$ (Assumption 2) holds for some $\sigma' \leq \frac{\sigma}{2^{k+6}(2N+M+2)\sqrt{k}}$.*¹⁷

It is known [AIK11, IW14, LL20a] that the information-theoretic version of arithmetic garbling satisfies the following simple simulation property – the labels are uniformly random conditioned on evaluation correctness. Construction 4 is similar in this regard. The garbling is a noisy linear randomization of the public matrices. Once LWE is applied:

- the expanded labels of fan-ins are uniformly random conditioned on the correct evaluation into the shrunken labels (of fan-outs),
- the garbled table is uniformly random conditioned on the correct recovery of the expanded labels from the shrunken labels (of the same gate), and
- the expanded labels of the output gate are uniformly random conditioned on the approximate recovery of $\mathbf{s}^\top(\tilde{\mathbf{W}}_{|C|, 0} + C(\mathbf{x}) \cdot \tilde{\mathbf{W}}_{|C|, \times})$.

Combining these points, the input labels and the garbled table are jointly random, conditioned on the approximate recovery of $\mathbf{s}^\top(\tilde{\mathbf{W}}_{|C|, 0} + C(\mathbf{x}) \cdot \tilde{\mathbf{W}}_{|C|, \times})$. When this latter value itself is random, the overall distribution of input labels and garbled table is simply random. Recall that security is considered only when $C(\mathbf{x}) \neq 0$. Since q is a prime, the evaluation result will be independent of the secret approximating $\mathbf{s}^\top \tilde{\mathbf{W}}_{|C|, 0}$ once LWE is applied. Therefore, the revealed components (secret, input labels, garbled table) are jointly pseudorandom. The proof below formalizes this idea.

¹⁷The constant has been changed from the conference version [HLL24]. Both versions are correct. The new version fits the proof better.

Proof (Theorem 17). We show $\text{Exp}_{\text{garble}}^{\text{real}} \approx \text{Exp}_{\text{garble}}^{\text{random}}$ with the following hybrids.

- H_0 is $\text{Exp}_{\text{garble}}^{\text{real}}$. The adversary sees $r_{\text{Setup}}, r_{\text{GenF}}$ and

$$\begin{aligned} w_{\text{out}} &= \mathbf{s}^\top \widetilde{\mathbf{W}}_{|C|,0} + e_{\text{out}}, & \mathbf{w}_\ell^\top &= \mathbf{s}^\top (\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_\ell^\top, \\ (\mathbf{t}_{i',0}^{(i,\gamma)})^\top &= \mathbf{s}^\top \mathbf{T}_{i',0}^{(i,\gamma)} + (\mathbf{e}_{i',0}^{(i,\gamma)})^\top, & (\mathbf{t}_{i',\times}^{(i,\gamma)})^\top &= \mathbf{s}^\top \mathbf{T}_{i',\times}^{(i,\gamma)} + (\mathbf{e}_{i',\times}^{(i,\gamma)})^\top, \end{aligned}$$

where the indices for \mathbf{t} are constrained by $\text{gin}[i, \gamma] = i'$ (henceforth same and omitted).

- In H_1 , we rewrite the garbling using distributivity and associativity. Recall that each block in \mathbf{T} is a signed sum of \mathbf{W} 's and \mathbf{U} 's. Instead of computing \mathbf{T} , multiplying \mathbf{s} , then adding the noise, we compute \mathbf{s} multiplied by \mathbf{W} 's and \mathbf{U} 's, sum them with appropriate signs, then add the noise. Namely, for all $i \in [|C|]$, compute

$$\mathbf{w}_{i,0}^\top = \mathbf{s}^\top \mathbf{W}_{i,0}, \quad \mathbf{w}_{i,\times}^\top = \mathbf{s}^\top \mathbf{W}_{i,\times}, \quad (\text{if } i > L) \quad \mathbf{u}_i^\top = \mathbf{s}^\top \mathbf{U}_i,$$

and $\widetilde{w}_{\text{out},0} = \mathbf{s}^\top \widetilde{\mathbf{W}}_{|C|,0}$ and $\widetilde{w}_{\text{out},\times} = \mathbf{s}^\top \widetilde{\mathbf{W}}_{|C|,\times}$. Those components are called “initial \mathbf{w} 's and \mathbf{u} 's”, which are not given directly to the adversary. Next, compute the randomized expanded label functions as follows (compare with GenF of Construction 4).

- If gate i is addition, set

$$\begin{aligned} \widetilde{\mathbf{w}}_{\text{gin}[i,1],0}^{(i,1)} &= \mathbf{u}_i, & \widetilde{\mathbf{w}}_{\text{gin}[i,1],\times}^{(i,1)} &= \mathbf{w}_{i,\times}, \\ \widetilde{\mathbf{w}}_{\text{gin}[i,2],0}^{(i,2)} &= \mathbf{w}_{i,0} - \mathbf{u}_i, & \widetilde{\mathbf{w}}_{\text{gin}[i,2],\times}^{(i,2)} &= \mathbf{w}_{i,\times}. \end{aligned}$$

- If gate i is multiplication, set

$$\begin{aligned} \widetilde{\mathbf{w}}_{\text{gin}[i,1],0}^{(i,1)} &= \mathbf{u}_i, & \widetilde{\mathbf{w}}_{\text{gin}[i,1],\times}^{(i,1)} &= \mathbf{w}_{i,\times}, \\ \widetilde{\mathbf{w}}_{\text{gin}[i,2],0}^{(i,2)} &= \mathbf{w}_{i,0}, & \widetilde{\mathbf{w}}_{\text{gin}[i,2],\times}^{(i,2)} &= -\mathbf{u}_i. \end{aligned}$$

Then, compute the \mathbf{s} -linear parts (denoted by $\widetilde{\mathbf{t}}$) of \mathbf{t} as

$$(\widetilde{\mathbf{t}}_{i',0}^{(i,\gamma)})^\top = \mathbf{w}_{i',0}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\widetilde{\mathbf{w}}_{i',0}^{(i,\gamma)})^\top, \quad (\widetilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})^\top = \mathbf{w}_{i',\times}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\widetilde{\mathbf{w}}_{i',\times}^{(i,\gamma)})^\top.$$

The randomized secret, input labels, and garble table are

$$\begin{aligned} w_{\text{out}} &= \widetilde{w}_{\text{out},0} + e_{\text{out}}, & \mathbf{w}_\ell &= \mathbf{w}_{\ell,0} + \mathbf{x}[\ell] \mathbf{w}_{\ell,\times} + \mathbf{e}_\ell, \\ \mathbf{t}_{i',0}^{(i,\gamma)} &= \widetilde{\mathbf{t}}_{i',0}^{(i,\gamma)} + \mathbf{e}_{\mathbf{t},i',0}^{(i,\gamma)}, & \mathbf{t}_{i',\times}^{(i,\gamma)} &= \widetilde{\mathbf{t}}_{i',\times}^{(i,\gamma)} + \mathbf{e}_{\mathbf{t},i',\times}^{(i,\gamma)}. \end{aligned}$$

Clearly, $H_0 \equiv H_1$.

- In H_2 , additional small noises are attached to the initial \mathbf{w} 's and \mathbf{u} 's as follows. Sample $\bar{\mathbf{e}}_{i,0}, \bar{\mathbf{e}}_{i,\times}, \bar{\mathbf{e}}_{\mathbf{U},i}$'s of appropriate dimensions and $\bar{e}_{\text{out},0}, \bar{e}_{\text{out},\times}$, each entry from $\mathcal{D}_{\mathbb{Z}, \sigma', \leq \sigma' \sqrt{\kappa}}$, and set

$$\begin{aligned} \mathbf{w}_{i,0}^\top &= \mathbf{s}^\top \mathbf{W}_{i,0} + \bar{\mathbf{e}}_{i,0}^\top, & \mathbf{w}_{i,\times}^\top &= \mathbf{s}^\top \mathbf{W}_{i,\times} + \bar{\mathbf{e}}_{i,\times}^\top, & (i > L) \quad \mathbf{u}_i^\top &= \mathbf{s}^\top \mathbf{U}_i + \bar{\mathbf{e}}_{\mathbf{U},i}^\top, \\ \widetilde{w}_{\text{out},0} &= \mathbf{s}^\top \widetilde{\mathbf{W}}_{|C|,0} + \bar{e}_{\text{out},0}, & \widetilde{w}_{\text{out},\times} &= \mathbf{s}^\top \widetilde{\mathbf{W}}_{|C|,\times} + \bar{e}_{\text{out},\times}. \end{aligned}$$

Recall that the initial \mathbf{w} 's and \mathbf{u} 's are not directly given to the adversary. In the revealed components (secret, input labels, garbled table), the magnitude of errors contributed by the additional small noises is bounded by

$$\sigma' \sqrt{\kappa} \cdot \max \{M + 1, N + 2\} \leq (N + M + 2) \sigma' \sqrt{\kappa} \leq 2^{-\kappa-6} \sigma.$$

Therefore, they are flooded (Lemma 2) by the randomizing noises (\mathbf{e} 's and e 's, without bars), i.e., $H_1 \approx_s H_2$.

- In H_3 , the small noises $\bar{\mathbf{e}}_{i,0}, \bar{\mathbf{e}}_{i,\times}, \bar{\mathbf{e}}_{\cup,i}$'s and $\bar{e}_{\text{out},0}, \bar{e}_{\text{out},\times}$ are no longer truncated, i.e., the entries are from $\mathcal{D}_{\mathbb{Z},\sigma'}$. We have $H_2 \approx_s H_3$ by Lemma 1.
- In H_4 , the initial \mathbf{w} 's and \mathbf{u} 's are replaced by random values, i.e.,

$$\mathbf{w}_{i,0} \xleftarrow{\$} \mathbb{Z}_q^N, \quad \mathbf{w}_{i,\times} \xleftarrow{\$} \mathbb{Z}_q^N, \quad (i > L) \quad \mathbf{u}_i \xleftarrow{\$} \mathbb{Z}_q^N, \quad \tilde{w}_{\text{out},0} \xleftarrow{\$} \mathbb{Z}_q, \quad \tilde{w}_{\text{out},\times} \xleftarrow{\$} \mathbb{Z}_q.$$

$H_3 \approx H_4$ follows from FlipLWE $_{N,N(L+|C|)+2,q,\sigma'}$.

For this step, we rely on the “straight-forwardness”, i.e., r_{GenF} , which must be produced by the reduction algorithm for the underlying adversary, is efficiently sampleable conditioned on and given $\mathbf{W}_{i,0}, \mathbf{W}_{i,\times}, \mathbf{U}_i$'s and $\tilde{\mathbf{W}}_{|C|,0}, \tilde{\mathbf{W}}_{|C|,\times}$, which are LWE public matrices received by the reduction algorithm.

- In H_5 , instead of sampling $\mathbf{w}_{i,0}$'s and \mathbf{u}_i 's, we sample and set

$$\begin{aligned} \mathbf{w}_{i,x_i} &\xleftarrow{\$} \mathbb{Z}_q^N, & \mathbf{w}_{i,0} &= \mathbf{w}_{i,x_i} - x_i \mathbf{w}_{i,\times}, \\ (i > L) \quad \bar{\mathbf{u}}_i &\xleftarrow{\$} \mathbb{Z}_q^N, & \mathbf{u}_i &= \bar{\mathbf{u}}_i - x_{\text{gin}[i,1]} \mathbf{w}_{i,\times}, \end{aligned}$$

where x_i 's are the wire values in the evaluation of $C(\mathbf{x})$. This is simply a change of variable, so $H_4 \equiv H_5$.

Approximations of \mathbf{w}_{i,x_i} 's would be computed during an honest evaluation — we call \mathbf{w}_{i,x_i} 's the “active” shrunken labels. Note that the input labels given to the adversary are just the active ones, i.e., $\mathbf{w}_\ell = \mathbf{w}_{i,0} + \mathbf{x}[\ell] \mathbf{w}_{i,\times} + \mathbf{e}_\ell = \mathbf{w}_{i,x_i} + \mathbf{e}_\ell$. For the “active” expanded labels, we use a case analysis. If gate i is addition,

$$\begin{aligned} \tilde{\mathbf{w}}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1)} &= \tilde{\mathbf{w}}_{\text{gin}[i,1],0}^{(i,1)} + x_{\text{gin}[i,1]} \tilde{\mathbf{w}}_{\text{gin}[i,1],\times}^{(i,1)} = \mathbf{u}_i + x_{\text{gin}[i,1]} \mathbf{w}_{i,\times} = \bar{\mathbf{u}}_i, \\ \tilde{\mathbf{w}}_{\text{gin}[i,2],x_{\text{gin}[i,2]}}^{(i,2)} &= \tilde{\mathbf{w}}_{\text{gin}[i,2],0}^{(i,2)} + x_{\text{gin}[i,2]} \tilde{\mathbf{w}}_{\text{gin}[i,2],\times}^{(i,2)} \\ &= (\mathbf{w}_{i,0} - \mathbf{u}_i) + x_{\text{gin}[i,2]} \mathbf{w}_{i,\times} \\ &= \mathbf{w}_{i,0} + (x_{\text{gin}[i,1]} + x_{\text{gin}[i,2]}) \mathbf{w}_{i,\times} - (\mathbf{u}_i + x_{\text{gin}[i,1]} \mathbf{w}_{i,\times}) \\ (x_i = x_{\text{gin}[i,1]} + x_{\text{gin}[i,2]}) &= \mathbf{w}_{i,x_i} - \bar{\mathbf{u}}_i. \end{aligned}$$

If gate i is multiplication,

$$\begin{aligned} \tilde{\mathbf{w}}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1)} &= \tilde{\mathbf{w}}_{\text{gin}[i,1],0}^{(i,1)} + x_{\text{gin}[i,1]} \tilde{\mathbf{w}}_{\text{gin}[i,1],\times}^{(i,1)} = \mathbf{u}_i + x_{\text{gin}[i,1]} \mathbf{w}_{i,\times} = \bar{\mathbf{u}}_i, \\ \tilde{\mathbf{w}}_{\text{gin}[i,2],x_{\text{gin}[i,2]}}^{(i,2)} &= \tilde{\mathbf{w}}_{\text{gin}[i,2],0}^{(i,2)} + x_{\text{gin}[i,2]} \tilde{\mathbf{w}}_{\text{gin}[i,2],\times}^{(i,2)} \\ &= \mathbf{w}_{i,0} - x_{\text{gin}[i,2]} \mathbf{u}_i \\ &= \mathbf{w}_{i,0} + x_{\text{gin}[i,1]} x_{\text{gin}[i,2]} \mathbf{w}_{i,\times} - x_{\text{gin}[i,2]} (\mathbf{u}_i + x_{\text{gin}[i,1]} \mathbf{w}_{i,\times}) \\ (x_i = x_{\text{gin}[i,1]} x_{\text{gin}[i,2]}) &= \mathbf{w}_{i,x_i} - x_{\text{gin}[i,2]} \bar{\mathbf{u}}_i. \end{aligned}$$

The point is to see that the active expanded labels are random subject to correct evaluation into active shrunken labels, and that they are independent of values with subscript “ \times ”. Moreover, $\tilde{\mathbf{t}}$ components with subscript “0” now become

$$\begin{aligned} (\tilde{\mathbf{t}}_{i',0}^{(i,\gamma)})^\top &= \mathbf{w}_{i',0}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{i',0}^{(i,\gamma)})^\top \\ &= (\mathbf{w}_{i',x_{i'}}^\top - x_{i'} \mathbf{w}_{i',\times}^\top) \mathbf{R}_{i'}^{(i,\gamma)} + ((\tilde{\mathbf{w}}_{i',x_{i'}}^{(i,\gamma)})^\top - x_{i'} (\tilde{\mathbf{w}}_{i',\times}^{(i,\gamma)})^\top) \\ &= \mathbf{w}_{i',x_{i'}}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{i',x_{i'}}^{(i,\gamma)})^\top - x_{i'} (\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})^\top, \end{aligned}$$

i.e., the active expanded label $(\tilde{\mathbf{w}}_{i',x_{i'}}^{(i,\gamma)})$ padded by the active shrunken label $(\mathbf{w}_{i',x_{i'}})$ stretched by \mathbf{R} , shifted by some known amount $(x_{i'} \tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})$.

- In H_6 , additional small noises are attached to $\tilde{\mathbf{t}}$ components with subscript “ \times ” as follows. Sample \mathbf{w}_{i,x_i} , $\mathbf{w}_{i,\times}$, $\bar{\mathbf{u}}_i$'s at random, compute $\tilde{\mathbf{w}}_{x',x_{i'}}^{(i,\gamma)}$'s as explained in H_5 , and compute $\tilde{\mathbf{w}}_{x',\times}^{(i,\gamma)}$ as usual. Sample $\bar{\mathbf{e}}_{t,i',\times}^{(i,\gamma)}$'s with entries independently from $\mathcal{D}_{\mathbb{Z},\sigma',\leq\sigma'\sqrt{\kappa}}$ and set

$$(\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})^\top = \mathbf{w}_{i',\times}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{i',\times}^{(i,\gamma)})^\top + (\bar{\mathbf{e}}_{t,i',\times}^{(i,\gamma)})^\top.$$

Compute $\tilde{\mathbf{t}}_{i',0}^{(i,\gamma)}$'s as explained in H_5 , and lastly the revealed components (secret, input labels, garbled table). The magnitude of inserted noises in the revealed components (in fact, just the randomized garbled table \mathbf{t}) is bounded by

$$\sigma' \sqrt{\kappa} \cdot \max\{|x_i|\} \leq (N + M + 2) \sigma' \sqrt{\kappa} \leq 2^{-\kappa-6} \sigma,$$

so $H_5 \approx_s H_6$ by Lemma 2.

- In H_7^0 , the small noises in $\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)}$'s are no longer truncated, i.e., entries of $\bar{\mathbf{e}}_{t,i',\times}^{(i,\gamma)}$ follow $\mathcal{D}_{\mathbb{Z},\sigma'}$. We have $H_6 \approx_s H_7^0$ by Lemma 1.
- In $H_7^{i'}$ (for i' from 1 up to $|C|$), the components $\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)}$ with $i' \leq i'$ are replaced by random. Comparing $H_7^{i'-1}$ and $H_7^{i'}$, the only change is

$$\begin{array}{ll} (\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})^\top & \text{from } \mathbf{w}_{i',\times}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{i',\times}^{(i,\gamma)})^\top + (\bar{\mathbf{e}}_{t,i',\times}^{(i,\gamma)})^\top \text{ in } H_7^{i'-1} \\ & \text{to random} \text{ in } H_7^{i'}, \text{ where } \text{gin}[i, \gamma] = i'. \end{array}$$

The indistinguishability follows from flipped LWE with secret $\mathbf{w}_{i',\times}$ and public matrix $(\dots, \mathbf{R}_{i'}^{(i,\gamma)}, \dots)_{\text{gin}[i,\gamma]=i'}$. To verify that the reduction works, we inspect where $\mathbf{w}_{i',\times}$ might appear in the revealed components.

- The secret is $w_{\text{out}} = \tilde{w}_{\text{out},0} + e_{\text{out}}$, independent of $\mathbf{w}_{i',\times}$.
- The input labels are $\mathbf{w}_\ell = \mathbf{w}_{\ell,x_\ell} + \mathbf{e}_\ell$, independent of $\mathbf{w}_{i',\times}$.
- The garbled table is $\tilde{\mathbf{t}}$ plus randomizing noises (without bars, independent of $\mathbf{w}_{i',\times}$). The $\tilde{\mathbf{t}}$ components with subscript “ \times ” require a case analysis.

* For $i' < i'$, the $\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)}$'s are already replaced by random, hence independent of $\mathbf{w}_{i',\times}$.

* For $i' \geq i'$, we have

$$(\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})^\top = \mathbf{w}_{i',\times}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{i',\times}^{(i,\gamma)})^\top + (\tilde{\mathbf{e}}_{i',\times}^{(i,\gamma)})^\top.$$

The first and third terms are either received together as input to the reduction algorithm (for $i' = i'$) or independent of $\mathbf{w}_{i',\times}$ (for $i' > i'$). The second term cannot contain $\mathbf{w}_{i',\times}$ since by construction (see H_1 and how $\bar{\mathbf{u}}$'s are set in H_5), it is computed from $\mathbf{w}_{i,\times}, \bar{\mathbf{u}}_i$ with subscript $i > i' \geq i'$ (topological sorting implies $i > i'$) as well as the wiring of C and the wire values of $C(\mathbf{x})$.

The $\tilde{\mathbf{t}}$ components with subscript “0”, as explained in H_5 , are computed from \mathbf{w}_{i,x_i} 's and $\tilde{\mathbf{w}}_{i',x_{i'}}$'s (independent of $\mathbf{w}_{i',\times}$) and $\tilde{\mathbf{t}}$ components with subscript “ \times ” (analyzed above, either received or sampled by reduction).

Therefore, $H_7^{i'-1} \approx H_7^{i'}$ follows from FlipLWE $_{N,m',q,\sigma'}$ (the minimum required m' varies by how gate i' is connected in C ; bounded by some fixed $\text{poly}(N, |C|)$). This step also relies on the “straight-forwardness”.

- In $H_8^{|C|+1}$, we replace $\tilde{\mathbf{t}}_{|C|,0}$ (actually a scalar) by random. Comparing $H_7^{|C|}$ and $H_8^{|C|+1}$, the only change is

$$\begin{array}{ccc} \tilde{\mathbf{t}}_{|C|,0}^\top & \text{from } \mathbf{w}_{|C|,x_{|C|}}^\top \mathbf{R}_{|C|} + \tilde{\mathbf{w}}_{|C|,x_{|C|}}^\top - x_{|C|} \tilde{\mathbf{t}}_{|C|,\times}^\top & \text{in } H_7^{|C|} \\ & \text{to random} & \text{in } H_8^{|C|+1}. \end{array}$$

Note that the active expanded label is $\tilde{\mathbf{w}}_{|C|,x_{|C|}} = \tilde{w}_{|C|,0} + x_{|C|} \tilde{w}_{|C|,\times}$. Since q is a prime, $x_{|C|} = C(\mathbf{x}) \neq 0$ (constraint for security), and $\tilde{w}_{|C|,\times}$ is not used elsewhere in the revealed components (secret, input labels, garbled table except $\mathbf{t}_{|C|,0}$), it follows that $x_{|C|} \tilde{w}_{|C|,\times}$ is a one-time pad for $\tilde{\mathbf{t}}_{|C|,0}$, so $H_7^{|C|} \equiv H_8^{|C|+1}$.

- In H_8^i (for i from $|C|$ down to $(L+1)$), the components $\tilde{\mathbf{t}}_{i',0}^{(i,\gamma)}$ with $i \geq i'$ (note that the indices being compared are i, i' , not i', i' as in H_7^i 's) are replaced by random. Comparing H_8^{i+1} and H_8^i , the only change is

$$\begin{array}{ccc} (\tilde{\mathbf{t}}_{\text{gin}[i,\gamma],0}^{(i,\gamma)})^\top & \text{from } \mathbf{w}_{\text{gin}[i,\gamma],x_{\text{gin}[i,\gamma]}}^\top \mathbf{R}_{\text{gin}[i,\gamma]}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{\text{gin}[i,\gamma],x_{\text{gin}[i,\gamma]}}^{(i,\gamma)})^\top - x_{\text{gin}[i,\gamma]} (\tilde{\mathbf{t}}_{\text{gin}[i,\gamma],\times}^{(i,\gamma)})^\top & \text{in } H_8^{i+1} \\ & \text{to random} & \text{in } H_8^i, \text{ where } \gamma \in \{1, 2\}. \end{array}$$

Observe that

$$(\tilde{\mathbf{w}}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1)}, \tilde{\mathbf{w}}_{\text{gin}[i,2],x_{\text{gin}[i,2]}}^{(i,2)}) = \begin{cases} (\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i} - \bar{\mathbf{u}}_i), & \text{if gate } i \text{ is addition;} \\ (\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i} - x_{\text{gin}[i,2]} \bar{\mathbf{u}}_i), & \text{if gate } i \text{ is multiplication;} \end{cases}$$

so $\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i}$ can serve as one-time pads if they do not appear elsewhere in the revealed components. We inspect their appearances.

- The secret is $w_{\text{out}} = \tilde{w}_{\text{out},0} + e_{\text{out}}$, no appearance.
- The input labels are $\mathbf{w}_\ell = \mathbf{w}_{\ell,x_\ell} + \mathbf{e}_\ell$, no appearance (note $i > L \geq \ell$).

- The garbled table is $\tilde{\mathbf{t}}$ plus randomizing noises (without bars, no $\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i}$). The $\tilde{\mathbf{t}}$ components with subscript “ \times ” are already replaced by random, so no appearance. The $\tilde{\mathbf{t}}_{i',0}^{(i,\gamma)}$ components require a case analysis.

- * For $i > i$, they are already replaced by random, so no appearance.
- * For $i \leq i$, we have

$$(\tilde{\mathbf{t}}_{i',0}^{(i,\gamma)})^\top = \mathbf{w}_{i',x_{i'}}^\top \mathbf{R}_{i'}^{(i,\gamma)} + (\tilde{\mathbf{w}}_{i',x_{i'}}^{(i,\gamma)})^\top - x_{i'} (\tilde{\mathbf{t}}_{i',\times}^{(i,\gamma)})^\top.$$

The first term has no appearance of $\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i}$, since $i' = \text{gin}[i, \gamma] < i \leq i$ by topological sorting. Neither does the third term. The second term either (if $i = i$) is among the ones being one-time-padded, or (if $i < i$) cannot contain $\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i}$, because by construction (see H_5), they are computed from only $\bar{\mathbf{u}}_i, \mathbf{w}_{i,x_i}$ (note $i < i$), the wiring of C , and the wire values of $C(\mathbf{x})$.

Therefore, $H_8^{i+1} \equiv H_8^i$.

- In H_9 , all the revealed components (secret, input labels, garbled table) are replaced by random. We inspect them in H_8^L .
 - The secret is $w_{\text{out}} = \tilde{w}_{\text{out},0} + e_{\text{out}}$ and $\tilde{w}_{\text{out},0}$ serves as a one-time pad.
 - The input labels are $\mathbf{w}_\ell = \mathbf{w}_{\ell,x_\ell} + \mathbf{e}_\ell$ and \mathbf{w}_{ℓ,x_ℓ} 's serve as one-time pads.
 - The garbled table is $\tilde{\mathbf{t}}$ plus randomizing noises. The $\tilde{\mathbf{t}}$ components with subscript “ \times ” are independent and random thanks to $H_7^{i'}$'s, and those with subscript “0”, thanks to H_8^i 's — the output component is handled in $H_8^{|C|+1}$ and the others must belong to some $\tilde{\mathbf{t}}_{i',0}^{(i,\gamma)}$ (for some $L < i \leq |C|$) hence indeed handled in H_8^i .

Therefore, $H_8^L \equiv H_9$.

Clearly, H_9 is just $\text{Exp}_{\text{garble}}^{\text{random}}$. By hybrid argument, $\text{Exp}_{\text{garble}}^{\text{real}} \equiv H_0 \approx H_9 \equiv \text{Exp}_{\text{garble}}^{\text{random}}$. \square

6 Ciphertext-Policy ABE from Short Noisy Linear Garbling

In this section, we show how to generically construct a CP-ABE scheme from a noisy linear garbling scheme and an identity-based evasive IPFE scheme.

Ingredients of Construction 5. We rely on

- a noisy linear garbling scheme NLG supporting $F = \{F_{\text{param}'}\}_{\text{param}' \in \text{Params}'}$ that is $(B_{\text{in}}, B_{\text{out}})$ -correct, B_{short} -short, and GenNoise-secure such that the output of GenNoise is in $[-B_{\text{NLG}}, B_{\text{NLG}}]^*$ (Definitions 7, 8, and 10), and
- an identity-based evasive IPFE scheme IPFE for $\mathcal{I} \supseteq \{0, 1, \dots, L\}$ that is B_{IPFE} -correct and restricted- σ_{IPFE} -secure (Definitions 4 and 6).

Construction 5 (CP-ABE). Define

$$\begin{aligned} \text{Params} &= \{ \text{param} = (\text{param}', 1^L) \mid \text{param}' \in \text{Params}', L \in \mathbb{N} \}, & X_{\text{param}', 1^L} &= \mathbb{Z}^L, \\ Y_{\text{param}', 1^L} &= \{ f : \mathbb{Z}^L \rightarrow \{0, 1, \perp\} \mid f \in F_{\text{param}'} \}, & P_{\text{param}', 1^L}(\mathbf{x}, f) &= f(\mathbf{x}). \end{aligned}$$

Our CP-ABE for P works as follows.

- Setup(param', 1^L) runs

$$\text{pp} = (q, \dots) \stackrel{\$}{\leftarrow} \text{NLG.Setup}(\text{param}'),$$

deterministically picks a suitable n_r based on q , and sets $K = 0$ and $Z = 2n_r$. It then sets up the IPFE scheme by

$$(\text{impk}, \text{imsk}) \stackrel{\$}{\leftarrow} \text{IPFE.Setup}(q, 1^K, 1^Z).$$

The algorithm outputs $\text{mpk} = (\text{pp}, 1^{n_r}, 1^K, 1^Z, \text{impk})$ and $\text{msk} = (\text{mpk}, \text{imsk})$.

- KeyGen(msk, \mathbf{x}) samples $\mathbf{r} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n_r}$ and sets

$$\mathbf{v}_0 \leftarrow \begin{pmatrix} \mathbf{r} \\ 1 \\ \mathbf{0}_{n_r-1} \end{pmatrix}, \quad \mathbf{v}_\ell \leftarrow \begin{pmatrix} \mathbf{r} \\ \mathbf{x}[\ell] \mathbf{r} \end{pmatrix} \quad \text{for all } \ell \in [L].$$

It generates IPFE secret keys id structured noises, unused

$$\text{isk}_0 \stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, \begin{matrix} \text{id} \\ 0 \end{matrix}, \mathbf{v}_0, \perp),$$

$$\text{isk}_\ell \stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, \begin{matrix} \ell \\ \mathbf{v}_\ell \end{matrix}, \perp) \quad \text{for all } \ell \in [L],$$

and outputs $\text{sk} = (\mathbf{r}, \text{isk}_0, \text{isk}_1, \dots, \text{isk}_L)$.

- Enc(mpk, f, μ) generates the noisy linear garbling of f by

$$(1^{n_f}, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}) \stackrel{\$}{\leftarrow} \text{GenF}(\text{pp}, f).$$

It samples a random secret matrix $\mathbf{S} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n_f \times n_r}$ and a sufficiently large message encoding scalar $s_{\text{msg}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa} q, 2 \cdot 2^{-\kappa} q]$. The algorithm computes

$$\mathbf{u}_{\text{msg}}^\top \leftarrow (\mathbf{w}_{\text{out}}^\top \mathbf{S}, \mu s_{\text{msg}}, \mathbf{0}_{1 \times (n_r-1)}), \quad \mathbf{U}_t^\top \leftarrow (\mathbf{T}^\top \mathbf{S}, \mathbf{0}_{\star \times n_r}),$$

$$\mathbf{U}_\ell^\top \leftarrow (\mathbf{W}_{\ell,0}^\top \mathbf{S}, \mathbf{W}_{\ell,\times}^\top \mathbf{S}) \quad \text{for all } \ell \in [L].$$

Lastly, it generates IPFE ciphertexts

$$\text{ict}_{\text{msg}} \stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 0, \mathbf{v}, \mathbf{u}_{\text{msg}}^\top), \quad \text{ict}_t \stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 0, \mathbf{v}, \mathbf{U}_t^\top),$$

$$\text{ict}_\ell \stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, \ell, \mathbf{v}, \mathbf{U}_\ell^\top) \quad \text{for all } \ell \in [L],$$

and outputs $\text{ct} = (\mathbf{R}, \text{ict}_{\text{msg}}, \text{ict}_t, \text{ict}_1, \dots, \text{ict}_L)$.

- Dec(mpk, $\mathbf{x}, \text{sk}, f, \text{ct}$) outputs \perp if $f(\mathbf{x}) \neq 1$. Otherwise, it parses sk, ct as in KeyGen, Enc, and recomputes $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_L$ as in KeyGen. The algorithm decrypts all the IPFE ciphertexts for the rerandomized garbling instance, i.e.,

$$w_{\text{msg}} \leftarrow \text{IPFE.Dec}(\text{impk}, 0, \mathbf{v}_0, \perp, \text{isk}_0, \mathbf{v}, \text{ict}_{\text{msg}}),$$

$$\mathbf{t} \leftarrow \text{IPFE.Dec}(\text{impk}, 0, \mathbf{v}_0, \perp, \text{isk}_0, \mathbf{v}, \text{ict}_t),$$

$$\mathbf{w}_\ell \leftarrow \text{IPFE.Dec}(\text{impk}, \ell, \mathbf{v}_\ell, \perp, \text{isk}_\ell, \mathbf{v}, \text{ict}_\ell) \quad \text{for all } \ell \in [L].$$

It evaluates the garbling by

$$w_{\text{out}} \leftarrow \text{NLG.Eval}(\text{pp}, f, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)$$

and outputs the decryption result

$$\mu' \leftarrow \begin{cases} 0, & \text{if } w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa} q, 2^{-\kappa} q], \\ 1, & \text{otherwise.} \end{cases}$$

6.1 Correctness

Theorem 18 (¶). *Construction 5 is correct (Definition 1) if*

- *the modulus q output by $\text{NLG.Setup}(\text{param}')$ is always a prime and always satisfies $B_{\text{out}}(\text{param}') \leq 2^{-\kappa}q$, and*
- *it holds that $B_{\text{IPFE}}(q, 0, 2n_r) \leq B_{\text{in}}(\text{param}')$.*

Proof (Theorem 18). The correctness of Construction 5 follows directly from the correctness of the underlying IPFE and garbling schemes. Since q is a prime, by the B_{IPFE} -correctness of the evasive IPFE scheme,

$$\begin{aligned} w_{\text{msg}} &= \mathbf{u}_{\text{msg}}^\top \mathbf{v}_0 + e_{\text{out}} = \mathbf{w}_{\text{out}}^\top \mathbf{S}\mathbf{r} + \mu s_{\text{msg}} + e_{\text{out}} = (\mathbf{S}\mathbf{r})^\top \mathbf{w}_{\text{out}} + e_{\text{out}} + \mu s_{\text{msg}}, \\ \mathbf{t}^\top &= \mathbf{v}_0^\top \mathbf{U}_t + \mathbf{e}_t^\top = (\mathbf{S}\mathbf{r})^\top \mathbf{T} + \mathbf{e}_t^\top, \\ \mathbf{w}_\ell^\top &= \mathbf{v}_\ell^\top \mathbf{U}_\ell + \mathbf{e}_\ell^\top = (\mathbf{S}\mathbf{r})^\top (\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_\ell^\top, \end{aligned}$$

where $\mathbf{S}\mathbf{r} \in \mathbb{Z}_q^{n_f}$ randomizes the garbling and the errors (from IPFE decryption) are bounded by $B_{\text{IPFE}}(q, 0, 2n_r) \leq B_{\text{in}}(\text{param}')$. By the $(B_{\text{in}}, B_{\text{out}})$ -correctness of garbling,

$$w_{\text{out}} - ((\mathbf{S}\mathbf{r})^\top \mathbf{w}_{\text{out}} + e_{\text{out}}) \in [-B_{\text{out}}(\text{param}'), B_{\text{out}}(\text{param}')] \subseteq [-2^{-\kappa}q, 2^{-\kappa}q].$$

If $\mu = 0$, it always holds that $w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa}q, 2^{-\kappa}q]$, which never holds when $\mu = 1$ by our choice of s_{msg} . We conclude that Construction 5 is correct. \square

6.2 Security

Theorem 19 (¶). *Construction 5 is very selectively secure (Definition 2) if¹⁸*

- $\text{LWE}_{n_r, \text{poly}(\lambda), q, \sigma_r}$ (Assumption 1) holds for some $\sigma_r \leq \frac{2^{-\kappa-6} \sigma_{\text{IPFE}}}{\sqrt{\kappa} \cdot B_{\text{short}}(\text{param}')}$, and
- *it holds that $2^{\kappa+6} B_{\text{NLG}}(\text{param}') \leq \sigma_{\text{IPFE}}$.*

Proof (Theorem 19). Let \mathcal{A} be an efficient adversary. We assume that it always chooses challenges such that $f(\mathbf{x}_j) = 0$ for all $j \in [J]$ – otherwise, we alter \mathcal{A} so that when the condition fails, it resets J to zero and aborts when it receives back the ABE components (incomplete per its expectation), as its output is irrelevant to $\text{Exp}_{\text{ABE}}^\beta$ in that case. Consider the following evasive IPFE sampler $\mathcal{S}^\beta = (\mathcal{S}_v, \mathcal{S}_u^\beta)$ per Definition 5.

- $\mathcal{S}_v(r_{\text{pub}})$ parses $r_{\text{pub}} = (r_{\mathcal{A}}, r_{\text{NLG.Setup}}, r_r, r_{\text{NLG.GenF}})$. It runs $\mathcal{A}(r_{\mathcal{A}})$ to obtain

$$\text{param} = (\text{param}', 1^L), \quad \{\mathbf{x}_j\}_{j \in [J]} \quad (\mathbf{x}_j \in \mathbb{Z}^L), \quad f : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}.$$

The algorithm next sets up the noisy linear garbling by

$$\text{pp} = (q, \dots) \leftarrow \text{NLG.Setup}(\text{param}'; r_{\text{NLG.Setup}}),$$

¹⁸The constants have been changed from the conference version [HLL24]. Both versions are correct. The new version fits the proof better.

and deterministically picks n_r, K, Z as in Setup of Construction 5. It then uses r_r to sample $\mathbf{r}_1, \dots, \mathbf{r}_J \in \mathbb{Z}_q^{n_r}$ uniformly at random in a straight-forward way.¹⁹ The algorithm sets for all $j \in [J]$,

$$\mathbf{v}_{j,0} = \begin{pmatrix} \mathbf{r}_j \\ 1 \\ \mathbf{0}_{n_r-1} \end{pmatrix}, \quad \mathbf{v}_{j,\ell} = \begin{pmatrix} \mathbf{r}_j \\ \mathbf{x}_j[\ell] \mathbf{r}_j \end{pmatrix} \quad (\text{for all } \ell \in [L]).$$

It runs

$$(1^{n_f}, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}) \leftarrow \text{NLG.GenF}(\text{pp}, f; r_{\text{NLG.GenF}}).$$

Suppose $\mathbf{W}_{\ell,0} \in \mathbb{Z}_q^{n_f \times m_\ell}$ and $\mathbf{T} \in \mathbb{Z}_q^{n_f \times m_t}$, the algorithm outputs

$$\begin{aligned} q, 1^K, 1^Z, \quad \text{ID} = \{0, 1, \dots, L\}, \quad 1^{J_{\text{id}}} = 1^J, \quad 1^{I_{0,0}} = 1^{m_t+1}, \quad 1^{I_{\ell,0}} = 1^{m_\ell}, \\ 1^{I_{\text{id},0}} = 1^0, \quad \sigma_{\text{pre}} = \sigma_{\text{PFSE}}, \quad \mathbf{V}_{\text{id},0} = (\mathbf{v}_{1,\text{id}}, \dots, \mathbf{v}_{J,\text{id}}), \quad \mathbf{V}_{\text{id},g} = \perp, \end{aligned}$$

where the indices are $\text{id} \in \text{ID}$ and $\ell \in [L]$.

- $\mathcal{S}_{\text{U}}^\beta(r_{\text{pub}}; r_{\text{priv}})$ first runs the deterministic subprocedure $\mathcal{S}_{\text{A0}}^\beta(r_{\text{pub}})$, which does the following. It first reruns $\mathcal{S}_{\text{V}}(r_{\text{pub}})$ to obtain $\mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}$. The algorithm then rearranges them into matrices $\mathbf{A}_{\text{msg}}, \{\mathbf{A}_{t,i}\}_{i \in [m_t]}, \{\mathbf{A}_{\ell,i}\}_{\ell \in [L], i \in [m_\ell]}$ so that for all $\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_{n_r} \in \mathbb{Z}_q^{n_f}$ and $s_{\text{msg}} \in \mathbb{Z}_q$,

$$\begin{aligned} \mathbf{d}_0^\top \mathbf{A}_{\text{msg}} &= \mathbf{u}_{\text{msg}}^\top = (\mathbf{w}_{\text{out}}^\top \mathbf{S}, \beta s_{\text{msg}}, \mathbf{0}_{1 \times (n_r-1)}), \\ (\mathbf{A}_{t,1}^\top \mathbf{d}_0, \dots, \mathbf{A}_{t,m_t}^\top \mathbf{d}_0)^\top &= \mathbf{U}_t^\top = (\mathbf{T}^\top \mathbf{S}, \mathbf{0}_{m_t \times n_r}), \\ (\mathbf{A}_{\ell,1}^\top \mathbf{d}_0, \dots, \mathbf{A}_{\ell,m_\ell}^\top \mathbf{d}_0)^\top &= \mathbf{U}_\ell^\top = (\mathbf{W}_{\ell,0}^\top \mathbf{S}, \mathbf{W}_{\ell,\times}^\top \mathbf{S}) \quad (\text{for all } \ell \in [L]), \end{aligned}$$

where $\mathbf{d}_0^\top = (s_{\text{msg}}, \tilde{\mathbf{s}}_1^\top, \dots, \tilde{\mathbf{s}}_{n_r}^\top)$ and $\mathbf{S} = (\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_{n_r})$. This is possible since every entry on the right-hand side is linear in \mathbf{d}_0 . The algorithm \mathcal{S}_{A0} outputs

$$\begin{aligned} 1^{n'} &= 1^{n_f n_r + 1}, \quad \mathbf{A}_{0,0,0,i} = \mathbf{A}_{t,i} \quad (\text{for all } i \in [m_t]), \\ \mathbf{A}_{\ell,0,0,m_t+1} &= \mathbf{A}_{\text{msg}}, \quad \mathbf{A}_{\ell,0,0,i} = \mathbf{A}_{\ell,i} \quad (\text{for all } \ell \in [L], i \in [m_\ell]). \end{aligned}$$

Completing \mathcal{S}_{A0} , the algorithm \mathcal{S}_{U} samples $\mathbf{d}_0 \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n'}$ using r_{priv} and outputs

$$\mathbf{U}_{\text{id},0}^\top = \begin{pmatrix} \mathbf{d}_0^\top \mathbf{A}_{\text{id},0,0,1} \\ \vdots \\ \mathbf{d}_0^\top \mathbf{A}_{\text{id},0,0,I_{\text{id},0}} \end{pmatrix} \quad \text{for all } \text{id} \in \text{ID}.$$

By definition, \mathcal{S}^β is a restricted sampler (Definition 6). We have the following:

Claim 20 (¶). \mathcal{S}^β has pseudorandom noisy inner products, i.e., $\text{IPFESec}_{\text{pre}}^{\mathcal{S}^\beta}$ (Definition 5) holds for both $\beta \in \{0, 1\}$.

The very selective security of Construction 5 follows from Claim 20. Consider the following hybrids.

¹⁹Precisely speaking, r_r must be efficiently sampleable conditioned on and given $\mathbf{r}_1, \dots, \mathbf{r}_J$.

- H_0^β is $\text{Exp}_{\text{ABE}}^\beta$ for Construction 5. Note that s_{msg} in the challenge ciphertext is uniformly random over $\mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa}q, 2 \cdot 2^{-\kappa}q]$.
- In H_1^β , the scalar s_{msg} is changed to be uniformly random over \mathbb{Z}_q . Clearly, $H_0^\beta \approx_s H_1^\beta$. Moreover, H_1^β corresponds to the left distribution of $\text{IPFEsec}_{\text{post}}^{S^\beta}$.
- In H_2^β , the vectors inside IPFE ciphertexts (components of the ABE challenge ciphertext) are replaced by random, which corresponds to the right distribution of $\text{IPFEsec}_{\text{post}}^{S^\beta}$. By Claim 20 and the restricted- σ_{IPFE} -security of IPFE, we have $H_1^\beta \approx H_2^\beta$.

Lastly, $H_2^0 \equiv H_2^1$. By hybrid argument, we conclude $\text{Exp}_{\text{ABE}}^0 \approx \text{Exp}_{\text{ABE}}^1$, completing the proof. \square

Proof (Claim 20). Fix $\beta \in \{0, 1\}$ and consider the following hybrids.

- H_0 is the left distribution of $\text{IPFEsec}_{\text{pre}}^{S^\beta}$. Recall that the public randomness is $r_{\text{pub}} = (r_{\mathcal{A}}, r_{\text{NLG.Setup}}, r_r, r_{\text{NLG.GenF}})$. Adopting the notations of the ABE decryption algorithm, the noisy inner products are

$$\begin{aligned} w_{\text{msg},j} &= \mathbf{u}_{\text{msg}}^\top \mathbf{v}_{j,0} + e_{\text{out},j} = (\mathbf{S}\mathbf{r}_j)^\top \mathbf{w}_{\text{out}} + \beta s_{\text{msg}} + e_{\text{out},j}, \\ \mathbf{t}_j^\top &= \mathbf{v}_{j,0}^\top \mathbf{U}_t + \mathbf{e}_{t,j}^\top = (\mathbf{S}\mathbf{r}_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= \mathbf{v}_{j,\ell}^\top \mathbf{U}_\ell + \mathbf{e}_{j,\ell}^\top = (\mathbf{S}\mathbf{r}_j)^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top \quad (\text{for all } \ell \in [L]), \end{aligned}$$

where $j \in [J]$ and the entries of \mathbf{e} 's are independent $\mathcal{D}_{\mathbb{Z}, \sigma_{\text{IPFE}}}$.

- In H_1 , additional small noises are attached to $\{\mathbf{S}\mathbf{r}_j\}_{j \in [J]}$. For all $j \in [J]$, sample $\mathbf{e}_{r,j} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}, \sigma_r, \leq \sigma_r \sqrt{\kappa}}^{n_f}$ and set the inner products to

$$\begin{aligned} w_{\text{msg},j} &= (\mathbf{S}\mathbf{r}_j + \mathbf{e}_{r,j})^\top \mathbf{w}_{\text{out}} + \beta s_{\text{msg}} + e_{\text{out},j}, & \mathbf{t}_j^\top &= (\mathbf{S}\mathbf{r}_j + \mathbf{e}_{r,j})^\top \mathbf{T} + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= (\mathbf{S}\mathbf{r}_j + \mathbf{e}_{r,j})^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top \quad (\text{for all } \ell \in [L]). \end{aligned}$$

The errors introduced into the inner products are bounded by

$$B_{\text{short}}(\text{param}') \cdot \sigma_r \sqrt{\kappa} \leq 2^{-\kappa-6} \sigma_{\text{IPFE}},$$

so they are flooded (Lemma 2) by $\{e_{\text{out},j}, \mathbf{e}_{t,j}, \{\mathbf{e}_{j,\ell}\}_{\ell \in [L]}\}_{j \in [J]}$, and $H_0 \approx_s H_1$.

- In H_2 , the $\mathbf{e}_{r,j}$'s are no longer truncated, i.e., $\mathbf{e}_{r,j} \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}, \sigma_r}^{n_f}$. We have $H_1 \approx_s H_2$ by Lemma 1.
- In H_3 , the garblings use uniform randomness. For all $j \in [J]$, sample $\mathbf{s}_j \xleftarrow{\$} \mathbb{Z}_q^{n_f}$ and set the inner products to

$$\begin{aligned} w_{\text{msg},j} &= \mathbf{s}_j^\top \mathbf{w}_{\text{out}} + \beta s_{\text{msg}} + e_{\text{out},j}, & \mathbf{t}_j^\top &= \mathbf{s}_j^\top \mathbf{T} + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= \mathbf{s}_j^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top \quad (\text{for all } \ell \in [L]). \end{aligned}$$

By n_f hybrids of applying $\text{LWE}_{n_r, J, q, \sigma_r}$ over the rows of \mathbf{S} with public matrix $(\mathbf{r}_1, \dots, \mathbf{r}_J)$, we have

$$\mathbf{S}(\mathbf{r}_1, \dots, \mathbf{r}_J) + (\mathbf{e}_{r,1}, \dots, \mathbf{e}_{r,J}) \approx (\mathbf{s}_1, \dots, \mathbf{s}_J),$$

which implies $H_2 \approx H_3$. In this step, we rely on the ‘‘straight-forwardness’’ of sampling \mathbf{r}_j 's from r_r in S^β .

- In H_4 , additional noises generated by GenNoise are attached to the garblings. For all $j \in [J]$, run

$$(\bar{e}_{\text{out},j}, \{\bar{\mathbf{e}}_{j,\ell}\}_{\ell \in [L]}, \bar{\mathbf{e}}_{t,j}) \stackrel{\$}{\leftarrow} \text{GenNoise}(\text{param}', \text{pp}, f, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]} \mathbf{T}, \mathbf{R}, \mathbf{x}_j),$$

and set the inner products to

$$\begin{aligned} w_{\text{msg},j} &= \mathbf{s}_j^\top \mathbf{w}_{\text{out}} + \bar{e}_{\text{out},j} + \beta s_{\text{msg}} + e_{\text{out},j}, & \mathbf{t}_j^\top &= \mathbf{s}_j^\top \mathbf{T} + \bar{\mathbf{e}}_{t,j}^\top + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= \mathbf{s}_j^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \bar{\mathbf{e}}_{j,\ell}^\top + \mathbf{e}_{j,\ell}^\top & (\text{for all } \ell \in [L]). \end{aligned}$$

The $\bar{e}, \bar{\mathbf{e}}$'s introduced into the inner products are bounded by $B_{\text{NLG}} \leq 2^{-\kappa-6} \sigma_{\text{IPFE}}$, so they are flooded (Lemma 2) by e, \mathbf{e} 's. We have $H_3 \approx_s H_4$.

- In H_5 , all the inner products are replaced by random, i.e., the right distribution of $\text{IPFEsec}_{\text{pre}}^{S^\beta}$. Note that the inner products in H_4 are well-randomized garblings (by $\mathbf{s}, \bar{e}, \bar{\mathbf{e}}$'s) plus independent noises (e, \mathbf{e} 's). Therefore, by J hybrids of applying GenNoise-security of NLG, we conclude $H_4 \approx H_5$.

By hybrid argument, $H_0 \approx H_5$, i.e., $\text{IPFEsec}_{\text{pre}}^{S^\beta}$ holds. \square

6.3 Unbounded Attribute and Summary

Unbounded Attribute. Construction 5 can be modified for attributes of unbounded length. Before presenting the scheme, we shall clarify the predicate family of such a scheme, in particular, what $P(\mathbf{x}, f)$ is when \mathbf{x} is shorter than an input to f . There are multiple approaches in the literature. The following is the discussion in [HLL23b; Section 5.4]. For monotone functions, the missing attribute bits can be assumed to be zero. For general functions: the specification is silent in many works, which should be interpreted as $P(\mathbf{x}, f) = \perp$ (neither correct nor secure); it might require exact length matching, or only reject \mathbf{x} shorter than input to f (prefix-matching), or attach index sets to \mathbf{x}, f and require equal-match or subset-match.

The “silent” specification is straight-forward to achieve. We explain how to achieve prefix-matching. For simplicity, we require that $f(\mathbf{x}) \neq \perp$ for all $\mathbf{x} \in \{0, 1\}^L$. The predicate family is

$$\begin{aligned} \text{Params} &= \{ \text{param} = \text{param}' \mid \text{param}' \in \text{Params}', \}, \\ X_{\text{param}'} &= \mathbb{Z}^{<2^\lambda}, & Y_{\text{param}'} &= \{ f : \mathbb{Z}^L \rightarrow \{0, 1, \perp\} \mid f \in F_{\text{param}'} \}, \\ P_{\text{param}'}(\mathbf{x}, f) &= \begin{cases} f(\mathbf{x}'), & \text{if the input length of } f \text{ is } L \text{ and } \mathbf{x}' \text{ is the } L\text{-prefix of } \mathbf{x}; \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Construction 5 with L (textually) removed is already quite close to an unbounded ABE – it becomes one with the following modifications similar to [HLL23b].

- The identity space satisfies $\mathcal{I} \supseteq \{0, 1, \dots, \lambda, \lambda + 1, \dots, \lambda + 2^\lambda - 1\}$.
- When generating a key for \mathbf{x} , instead generate a key for $(L_{\mathbf{x}}, \mathbf{x})$, where $L_{\mathbf{x}}$ is the λ -bit encoding of \mathbf{x} .

- When generating a ciphertext for f , instead generate a ciphertext for

$$f'(L_x, \mathbf{x}') = \begin{cases} 0, & \text{if } L_x \text{ is less than the input length of } f; \\ f(\mathbf{x}'), & \text{otherwise.} \end{cases}$$

- The security reduces to that of Construction 5 by appropriately appending zeros to every key attribute.

Summary. By instantiating our CP-ABE with the garbling scheme of Construction 4 and the IPFE scheme of Construction 2, we obtain CP-ABE for bounded-arithmetic circuits.

Corollary 21. *Under LWE (Assumption 1) and evasive LWE (Assumption 4), there exist CP-ABE schemes (bounded/unbounded in attribute length) for bounded-arithmetic circuits with $|\text{mpk}| = \text{poly}(\lambda, \log M, d)$ and*

$$|\text{sk}_x| = (|\mathbf{x}| + 1) \text{poly}(\lambda, \log M, d), \quad |\text{ct}_C| = (|C| + 1) \text{poly}(\lambda, \log M, d),$$

where M is the maximum wire value and d is the maximum depth.

7 ABE for DFA

We define DFA in a minimal format convenient for garbling.

Definition 11 (DFA). A *deterministic finite automaton* Γ over the binary alphabet is a tuple $(1^\Omega, \mathbf{\Gamma}_0, \mathbf{\Gamma}_1, \boldsymbol{\xi})$, where $\Omega \in \mathbb{N}$ is its *number of states*, $\mathbf{\Gamma}_0, \mathbf{\Gamma}_1 \in \{0, 1\}^{\Omega \times \Omega}$ are its *transition matrices*, $\boldsymbol{\xi} \in \{0, 1\}^\Omega$ is its *rejection state vector*, and each column of $\mathbf{\Gamma}_0, \mathbf{\Gamma}_1$ contains exactly one 1. The DFA *accepts* an input $\mathbf{x} \in \{0, 1\}^L$ if and only if

$$\boldsymbol{\xi}^\top \cdot \prod_{\ell=L}^1 \mathbf{\Gamma}_{\mathbf{x}[\ell]} \cdot \boldsymbol{\iota}_1 = 0.$$

Implicitly, for Γ , the set of states is $[\Omega]$, the initial state is 1, and the set of accept states is $\{q \in [\Omega] \mid \boldsymbol{\xi}[q] = 0\}$. After reading $x \in \{0, 1\}$, the machine transitions from q to q' if and only if $\mathbf{\Gamma}_x[q', q] = 1$, which is equivalent to $\mathbf{\Gamma}_x \boldsymbol{\iota}_q = \boldsymbol{\iota}_{q'}$ (recall that $\boldsymbol{\iota}_q \in \{0, 1\}^\Omega$ is the q^{th} standard basis vector). The following lemma, readily verified, will be handy:

Lemma 22. *For all DFA $(1^\Omega, \mathbf{\Gamma}_0, \mathbf{\Gamma}_1, \boldsymbol{\xi})$ and input $\mathbf{x} \in \{0, 1\}^L$, it holds that*

$$\left\| \prod_{\ell=1}^L \mathbf{\Gamma}_{\mathbf{x}[\ell]}^\top \right\| \leq 1, \quad \boldsymbol{\xi}^\top \cdot \prod_{\ell=L}^1 \mathbf{\Gamma}_{\mathbf{x}[\ell]} \cdot \boldsymbol{\iota}_1 \in \{0, 1\}.$$

7.1 Noisy Linear Garbling for DFA

DFA is defined for inputs of arbitrary length, yet our notion of garbling is only for functions with fixed-length input. Following the paradigm in [LL20a], we consider the garbling scheme for each possible input length. The existing [Wat12, LL20a] linear secret sharing scheme for DFA can be cast into a noisy linear garbling (Definition 7).

Construction 6 (DFA garbling). Let²⁰

$$\begin{aligned} \text{Params} &= \{ (q, \bar{L}) \mid q, \bar{L} \in \mathbb{N}, q \geq 2 \}, \quad F = \{ F_{\text{param}} \}_{\text{param} \in \text{Params}}, \\ F_{\text{param}} &= \{ f_{\Gamma, L} : \mathbb{Z}^L \rightarrow \{0, 1, \perp\} \mid \Gamma \text{ is a DFA and } L \in [0, \bar{L}] \}, \\ f_{\Gamma, L}(\mathbf{x}) &= \begin{cases} 1, & \text{if } \mathbf{x} \in \{0, 1\}^L \text{ and } \Gamma \text{ accepts } \mathbf{x}; \\ 0, & \text{if } \mathbf{x} \in \{0, 1\}^L \text{ and } \Gamma \text{ rejects } \mathbf{x}; \\ \perp, & \text{if } \mathbf{x} \notin \{0, 1\}^L. \end{cases} \end{aligned}$$

The function $f_{\Gamma, L}$ is represented by $(\Gamma, 1^L)$. The noisy linear garbling scheme for F works as follows.

- Setup(param) outputs $\text{pp} = \text{param} = (q, \bar{L})$.
- GenF(pp, $f_{\Gamma, L}$) parses $\Gamma = (1^{\mathfrak{Q}}, \mathbf{\Gamma}_0, \mathbf{\Gamma}_1, \boldsymbol{\xi})$, sets

$$\begin{aligned} n_f &= (L+1)\mathfrak{Q} + 1, & \mathbf{w}_{\text{out}}^{\top} &= (-\boldsymbol{\iota}_1^{\top}, \mathbf{0}_{1 \times L\mathfrak{Q}}, 0), & \mathbf{R} &= \perp, \\ \mathbf{W}_{\ell, 0} &= \begin{pmatrix} \mathbf{0}_{(\ell-1)\mathfrak{Q} \times \mathfrak{Q}} \\ -\mathbf{I}_{\mathfrak{Q}} \\ \mathbf{\Gamma}_0 \\ \mathbf{0}_{(L-\ell)\mathfrak{Q} \times \mathfrak{Q}} \\ \mathbf{0}_{1 \times \mathfrak{Q}} \end{pmatrix}, & \mathbf{W}_{\ell, \times} &= \begin{pmatrix} \mathbf{0}_{(\ell-1)\mathfrak{Q} \times \mathfrak{Q}} \\ \mathbf{0}_{\mathfrak{Q} \times \mathfrak{Q}} \\ \mathbf{\Gamma}_1 - \mathbf{\Gamma}_0 \\ \mathbf{0}_{(L-\ell)\mathfrak{Q} \times \mathfrak{Q}} \\ \mathbf{0}_{1 \times \mathfrak{Q}} \end{pmatrix}, & \mathbf{T} &= \begin{pmatrix} \mathbf{0}_{L\mathfrak{Q} \times \mathfrak{Q}} \\ -\mathbf{I}_{\mathfrak{Q}} \\ \boldsymbol{\xi}^{\top} \end{pmatrix}, \end{aligned}$$

and outputs $(1^{n_f}, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell, 0}, \mathbf{W}_{\ell, \times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R})$. Here, $\boldsymbol{\iota}_1 \in \{0, 1\}^{\mathfrak{Q}}$.

- Eval(pp, $f_{\Gamma, L}$, $\mathbf{R}, \mathbf{x}, \{\mathbf{w}_{\ell}^{\top}\}_{\ell \in [L]}, \mathbf{t}^{\top}$) computes and outputs

$$\mathbf{t}^{\top} \cdot \prod_{\ell'=L}^1 \mathbf{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 + \sum_{\ell=L}^1 \left(\mathbf{w}_{\ell}^{\top} \cdot \prod_{\ell'=\ell-1}^1 \mathbf{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 \right).$$

Theorem 23 (¶). *Construction 6 is $(B_{\text{in}}, B_{\text{out}})$ -correct (Definition 7) if $B_{\text{out}} \geq (\bar{L} + 2)B_{\text{in}}$, 2-short (Definition 8), and GenNoise-secure for all GenNoise (Definition 10).*

Proof (Theorem 23). The proof is basically that in [LL20a] plus handling the noise. It suffices to consider $\mathbf{x} \in \{0, 1\}^L$. Note that (the second by Lemma 22)

$$\begin{aligned} \mathbf{w}_{\text{out}}^{\top} &= (-\boldsymbol{\iota}_1^{\top}, \mathbf{0}) & \implies & & \|\mathbf{w}_{\text{out}}^{\top}\| &\leq 1 &\leq 2, \\ \mathbf{W}_{\ell, 0}^{\top} + \mathbf{x}[\ell]\mathbf{W}_{\ell, \times}^{\top} &= (\mathbf{0}, -\mathbf{I}_{\mathfrak{Q}}, \mathbf{\Gamma}_{\mathbf{x}[\ell]}^{\top}, \mathbf{0}) & \implies & & \|\mathbf{W}_{\ell, 0}^{\top} + \mathbf{x}[\ell]\mathbf{W}_{\ell, \times}^{\top}\| &\leq 1 + 1 \leq 2, \\ \mathbf{T}^{\top} &= (\mathbf{0}, -\mathbf{I}_{\mathfrak{Q}}, \boldsymbol{\xi}) & \implies & & \|\mathbf{T}^{\top}\| &\leq 1 + 1 \leq 2, \end{aligned}$$

so the scheme is 2-short.

Let the garbling randomness/noise be

$$\mathbf{s} = (\mathbf{s}_0^{\top}, \mathbf{s}_1^{\top}, \dots, \mathbf{s}_L^{\top}, \mathbf{s}_{L+1}^{\top})^{\top}, \quad \mathbf{e} = (e_0, \mathbf{e}_1^{\top}, \dots, \mathbf{e}_L^{\top}, \mathbf{e}_{L+1}^{\top})^{\top},$$

where each $\mathbf{s}_{\ell}, \mathbf{e}_{\ell}$ is \mathfrak{Q} -dimensional, then

$$w_{\text{out}} = -\mathbf{s}_0^{\top} \boldsymbol{\iota}_1 + e_0, \quad \mathbf{w}_{\ell}^{\top} = \mathbf{s}_{\ell}^{\top} \mathbf{\Gamma}_{\mathbf{x}[\ell]} - \mathbf{s}_{\ell-1}^{\top} + \mathbf{e}_{\ell}^{\top}, \quad \mathbf{t}^{\top} = \mathbf{s}_{L+1}^{\top} \boldsymbol{\xi}^{\top} - \mathbf{s}_L^{\top} + \mathbf{e}_{L+1}^{\top}.$$

Since evaluation is linear in \mathbf{w} 's and \mathbf{t} , we first consider the non-noisy part,

$$(\mathbf{t} - \mathbf{e}_{L+1})^{\top} \cdot \prod_{\ell'=L}^1 \mathbf{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 + \sum_{\ell=L}^1 \left((\mathbf{w}_{\ell} - \mathbf{e}_{\ell})^{\top} \cdot \prod_{\ell'=\ell-1}^1 \mathbf{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 \right)$$

²⁰Here, the DFA input length upper bound \bar{L} is needed due to perfect correctness requirement in the presence of noise accumulation. For our garbling and ABE constructions, we can simply set $\bar{L} = 2^{\lambda}$ to support arbitrary polynomial-size computations.

$$\begin{aligned}
&= (s_{L+1} \boldsymbol{\xi}^\top - \mathbf{s}_L^\top) \cdot \prod_{\ell'=L}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 + \sum_{\ell=L}^1 \left((\mathbf{s}_\ell^\top \boldsymbol{\Gamma}_{\mathbf{x}[\ell]} - \mathbf{s}_{\ell-1}^\top) \cdot \prod_{\ell'=\ell-1}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 \right) \\
&= s_{L+1} \boldsymbol{\xi}^\top \cdot \prod_{\ell'=L}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 - \mathbf{s}_L^\top \cdot \prod_{\ell'=L}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 \\
&\quad (\text{telescoping}) \quad + \sum_{\ell=L}^1 \left(\mathbf{s}_\ell^\top \cdot \prod_{\ell'=\ell}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 - \mathbf{s}_{\ell-1}^\top \cdot \prod_{\ell'=\ell-1}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 \right) \\
&= s_{L+1} \boldsymbol{\xi}^\top \cdot \prod_{\ell'=L}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 - \mathbf{s}_0^\top \boldsymbol{\iota}_1 = (w_{\text{out}} - e_0) + \begin{cases} 0, & \text{if } \Gamma \text{ accepts } \mathbf{x}; \\ s_{L+1}, & \text{otherwise.} \end{cases}
\end{aligned}$$

Writing

$$e_{\text{Eval}} = \mathbf{e}_{L+1}^\top \cdot \prod_{\ell'=L}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 + \sum_{\ell=L}^1 \left(\mathbf{e}_\ell^\top \cdot \prod_{\ell'=\ell-1}^1 \boldsymbol{\Gamma}_{\mathbf{x}[\ell']} \cdot \boldsymbol{\iota}_1 \right) - e_0,$$

we have

$$\text{Eval}(\text{pp}, f_{\Gamma, L}, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top) = w_{\text{out}} + e_{\text{Eval}} + \begin{cases} 0, & \text{if } \Gamma \text{ accepts } \mathbf{x}; \\ s_{L+1}, & \text{otherwise.} \end{cases}$$

Transposing e_{Eval} and applying Lemma 22, we find

$$\begin{aligned}
|e_{\text{Eval}}| &\leq \underbrace{\|\boldsymbol{\iota}_1^\top\| \cdot \prod_{\ell'=1}^L \|\boldsymbol{\Gamma}_{\mathbf{x}[\ell']}\|}_{\leq 1} \cdot \|\mathbf{e}_{L+1}\| + \sum_{\ell=L}^1 \left(\|\boldsymbol{\iota}_1^\top\| \cdot \underbrace{\prod_{\ell'=1}^{\ell-1} \|\boldsymbol{\Gamma}_{\mathbf{x}[\ell']}\|}_{\leq 1} \cdot \|\mathbf{e}_\ell\| \right) + |e_0| \\
&\leq (L+2)B_{\text{in}} \leq (\bar{L}+2)B_{\text{in}} \leq B_{\text{out}}.
\end{aligned}$$

This proves $(B_{\text{in}}, B_{\text{out}})$ -correctness.

Perfect security without noise is proven in [LL20a], which implies security with any noise. To recap, observe that with param, pp, $f_{\Gamma, L}$, output of GenF, and $\mathbf{x} \in \{0, 1\}^L$ fixed, the values $(\mathbf{w}_1, \dots, \mathbf{w}_L, s_{L+1})$ are jointly uniformly random — thanks to the subtraction of \mathbf{s}_ℓ in each \mathbf{w}_ℓ . When \mathbf{x} is rejected by Γ ,

$$w_{\text{out}} = \underbrace{\text{Eval}(\text{pp}, f_{\Gamma, L}, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)}_{\text{independent of } s_{L+1}} - e_{\text{Eval}} - s_{L+1},$$

so $(w_{\text{out}}, \mathbf{w}_1, \dots, \mathbf{w}_L)$ is again jointly uniformly random. \square

7.2 Construction of KP-ABE for DFA

We present our KP-ABE for DFA, utilizing the telescoping-sum structure of the DFA garbling. The idea is similar to [LL20a] — an ABE key contains roughly Ω IPFE keys and an ABE ciphertext, L IPFE ciphertexts, so that decryption yields $\Theta(L\Omega)$ inner products corresponding to a garbling generated using pseudorandomness. The proof, yet, is much simpler than that of [LL20a] (as well as those of all previous pairing-based ABE for DFA), thanks to the easy-to-use security of evasive IPFE (Definition 5).

Ingredients of Construction 7. We rely on

- Construction 6, a specific noisy linear garbling NLG for DFA, and
- an identity-based evasive IPFE scheme IPFE for $\mathcal{I} \supseteq \{0, 1, 2\}$ that is B_{IPFE} -correct and restricted- σ_{IPFE} -secure (Definitions 4 and 6).

Construction 7 (KP-ABE for DFA). Let

$$\text{Params} = \{ \bar{L} \mid \bar{L} \in \mathbb{N} \}, \quad X_{\bar{L}} = \{ \text{DFA } \Gamma \}, \quad Y_{\bar{L}} = \{ \mathbf{x} \in \{0, 1\}^L \mid L \leq \bar{L} \},$$

$$P_{\bar{L}}(\Gamma, \mathbf{x}) = \begin{cases} 1, & \text{if } \Gamma \text{ accepts } \mathbf{x}; \\ 0, & \text{if } \Gamma \text{ rejects } \mathbf{x}. \end{cases}$$

Our KP-ABE for DFA works as follows.

- $\text{Setup}(\bar{L})$ picks suitable q, n in a straight-forward way.²¹ The algorithm sets $\text{pp} = (q, \bar{L})$ and $K = 0, Z = 3n$. It runs $(\text{impk}, \text{imsk}) \xleftarrow{\$} \text{IPFE.Setup}(q, 1^K, 1^Z)$ and outputs

$$\text{mpk} = (\text{pp}, 1^n, 1^K, 1^Z, \text{impk}), \quad \text{msk} = (\text{mpk}, \text{imsk}).$$

- $\text{KeyGen}(\text{msk}, \Gamma)$ parses $\Gamma = (1^\Omega, \mathbf{\Gamma}_0, \mathbf{\Gamma}_1, \xi)$. It samples as sets

$$\mathbf{R} \xleftarrow{\$} \mathbb{Z}_q^{n \times \Omega}, \quad \mathbf{r} \xleftarrow{\$} \mathbb{Z}_q^n,$$

$$\mathbf{v}_{\text{msg}} = \begin{pmatrix} -\mathbf{R}\boldsymbol{\mu}_1 \\ \mathbf{0}_n \\ 1 \\ \mathbf{0}_{n-1} \end{pmatrix}, \quad \mathbf{V}_t = \begin{pmatrix} -\mathbf{R} \\ \mathbf{r}\xi^\top \\ \mathbf{0}_{n \times \Omega} \end{pmatrix}, \quad \mathbf{V}_{\text{in}} = \begin{pmatrix} \mathbf{R}\mathbf{\Gamma}_0 \\ \mathbf{R}\mathbf{\Gamma}_1 \\ -\mathbf{R} \end{pmatrix},$$

where $\boldsymbol{\mu}_1 = (1, 0, \dots, 0)^\top \in \{0, 1\}^\Omega$ is the first standard basis vector. The algorithm generates IPFE secret keys

$$\text{isk}_{\text{msg}} \xleftarrow{\$} \text{IPFE.KeyGen}(\text{imsk}, 0, \mathbf{v}_{\text{msg}}, \perp),$$

$$\text{isk}_t \xleftarrow{\$} \text{IPFE.KeyGen}(\text{imsk}, 1, \mathbf{V}_t, \perp),$$

$$\text{isk}_{\text{in}} \xleftarrow{\$} \text{IPFE.KeyGen}(\text{imsk}, 2, \mathbf{V}_{\text{in}}, \perp).$$

It outputs $\text{sk} = (\mathbf{R}, \mathbf{r}, \text{isk}_{\text{msg}}, \text{isk}_t, \text{isk}_{\text{in}})$.

- $\text{Enc}(\text{mpk}, \mathbf{x}, \mu)$ samples

$$\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_L, \tilde{\mathbf{s}}_{L+1} \xleftarrow{\$} \mathbb{Z}_q^n, \quad s_{\text{msg}} \xleftarrow{\$} \mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa} q, 2 \cdot 2^{-\kappa} q].$$

It sets

$$\mathbf{u}_{\text{msg}}^\top = \left(\tilde{\mathbf{s}}_0^\top, \mathbf{0}_{1 \times n}, \mu s_{\text{msg}}, \mathbf{0}_{1 \times (n-1)} \right),$$

$$\mathbf{u}_t^\top = \left(\tilde{\mathbf{s}}_L^\top, \tilde{\mathbf{s}}_{L+1}^\top, \mathbf{0}_{1 \times n} \right),$$

$$\mathbf{u}_\ell^\top = \left((1 - \mathbf{x}[\ell])\tilde{\mathbf{s}}_\ell^\top, \mathbf{x}[\ell]\tilde{\mathbf{s}}_\ell^\top, \tilde{\mathbf{s}}_{\ell-1}^\top \right) \quad \text{for all } \ell \in [L],$$

²¹Precisely speaking, the randomness used to sample q, n must be separate from that to run IPFE.Setup and efficiently sampleable conditioned on and given q, n .

generates IPFE ciphertexts

$$\begin{aligned} \text{ict}_{\text{msg}} &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 0, \mathbf{o}, \mathbf{u}_{\text{msg}}^\top), \\ \text{ict}_t &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 1, \mathbf{o}, \mathbf{u}_t^\top), \\ \text{ict}_\ell &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 2, \mathbf{o}, \mathbf{u}_\ell^\top) \quad \text{for all } \ell \in [L], \end{aligned}$$

and outputs $\text{ct} = (\text{ict}_{\text{msg}}, \text{ict}_t, \text{ict}_1, \dots, \text{ict}_L)$.

- $\text{Dec}(\text{mpk}, \Gamma, \text{sk}, \mathbf{x}, \text{ct})$ outputs \perp and terminates if Γ rejects \mathbf{x} . Otherwise, it parses sk, ct as in $\text{KeyGen}, \text{Enc}$. Suppose $\mathbf{x} \in \{0, 1\}^L$, the algorithm recomputes $\mathbf{v}_{\text{msg}}, \mathbf{V}_t, \mathbf{V}_{\text{in}}$ as in KeyGen and performs IPFE decryptions

$$\begin{aligned} w_{\text{msg}} &\leftarrow \text{IPFE.Dec}(\text{impk}, 0, \mathbf{v}_{\text{msg}}, \perp, \text{isk}_{\text{msg}}, \mathbf{o}, \text{ict}_{\text{msg}}), \\ \mathbf{t}^\top &\leftarrow \text{IPFE.Dec}(\text{impk}, 1, \mathbf{V}_t, \perp, \text{isk}_t, \mathbf{o}, \text{ict}_t), \\ \mathbf{w}_\ell^\top &\leftarrow \text{IPFE.Dec}(\text{impk}, 2, \mathbf{V}_{\text{in}}, \perp, \text{isk}_{\text{in}}, \mathbf{o}, \text{ict}_\ell) \quad \text{for all } \ell \in [L]. \end{aligned}$$

It evaluates the DFA garbling

$$w_{\text{out}} \leftarrow \text{NLG.Eval}(\text{pp}, (\Gamma, 1^L), \perp, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)$$

and outputs

$$\mu' \leftarrow \begin{cases} 0, & \text{if } w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa}q, 2^{-\kappa}q]; \\ 1, & \text{otherwise.} \end{cases}$$

Theorem 24 (¶). *Construction 7 is correct (Definition 1) if $(\bar{L} + 2)B_{\text{IPFE}}(q, 0, 3n) \leq 2^{-\kappa}q$.*

Proof (Theorem 24). The IPFE decryption results in Dec are

$$\begin{aligned} w_{\text{msg}} &= -\tilde{\mathbf{s}}_0^\top \mathbf{R} \boldsymbol{\mu}_1 + \mu s_{\text{msg}} + e_{\text{out}} = \tilde{\mathbf{s}}_0^\top \mathbf{R} \cdot (-\boldsymbol{\mu}_1) + \mu s_{\text{msg}} + e_{\text{out}}, \\ \mathbf{t}^\top &= \mathbf{u}_t^\top \mathbf{V}_t + \mathbf{e}_t^\top = -\tilde{\mathbf{s}}_L^\top \mathbf{R} + \tilde{\mathbf{s}}_{L+1}^\top \mathbf{r} \cdot \boldsymbol{\xi}^\top + \mathbf{e}_t^\top, \\ \mathbf{w}_\ell^\top &= \mathbf{u}_\ell^\top \mathbf{V}_{\text{in}} + \mathbf{e}_\ell^\top = (1 - \mathbf{x}[\ell])\tilde{\mathbf{s}}_\ell^\top \mathbf{R} \boldsymbol{\Gamma}_0 + \mathbf{x}[\ell]\tilde{\mathbf{s}}_\ell^\top \mathbf{R} \boldsymbol{\Gamma}_1 - \tilde{\mathbf{s}}_{\ell-1}^\top \mathbf{R} + \mathbf{e}_\ell^\top \\ &= \tilde{\mathbf{s}}_\ell^\top \mathbf{R} \cdot \boldsymbol{\Gamma}_{\mathbf{x}[\ell]} - \tilde{\mathbf{s}}_{\ell-1}^\top \mathbf{R} + \mathbf{e}_\ell^\top \quad \text{for all } \ell \in [L], \end{aligned}$$

where the errors $e_{\text{out}}, \mathbf{e}_t, \{\mathbf{e}_\ell\}_{\ell \in [L]}$ are B_{IPFE} -bounded by the correctness of IPFE. They form a DFA garbling with randomness

$$\mathbf{s} = (\tilde{\mathbf{s}}_0^\top \mathbf{R}, \tilde{\mathbf{s}}_1^\top \mathbf{R}, \dots, \tilde{\mathbf{s}}_L^\top \mathbf{R}, \tilde{\mathbf{s}}_{L+1}^\top \mathbf{r})^\top.$$

Therefore, by the $(B_{\text{IPFE}}, (\bar{L} + 2)B_{\text{IPFE}})$ -correctness of NLG (Theorem 23), we have

$$\begin{aligned} w_{\text{out}} - (\tilde{\mathbf{s}}_0^\top \mathbf{R} \cdot (-\boldsymbol{\mu}_1) + e_{\text{out}}) &\in [-(\bar{L} + 2)B_{\text{IPFE}}(q, 0, 3n), (\bar{L} + 2)B_{\text{IPFE}}(q, 0, 3n)] \\ &\subseteq [-2^{-\kappa}q, 2^\kappa q]. \end{aligned}$$

If $\mu = 0$, it always holds that $w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa}q, 2^{-\kappa}q]$, which never holds when $\mu = 1$ by our choice of s_{msg} . We conclude that Construction 7 is correct. \square

7.3 Security of KP-ABE for DFA

Theorem 25 (¶). *Construction 7 is very selectively secure (Definition 2) if $\text{LWE}_{n, \text{poly}(\lambda), q, \sigma_r}$ (Assumption 1) holds for some $\sigma_r \leq \frac{2^{-\kappa-6} \sigma_{\text{IPFE}}}{2\sqrt{\kappa}}$.*²²

Proof (Theorem 25). The proof is similar to that of Theorem 19. Let \mathcal{A} be an efficient adversary and assume that it always chooses challenges such that Γ_j rejects \mathbf{x} for all $j \in [J]$. Consider the following evasive IPFE sampler $\mathcal{S}^\beta = (\mathcal{S}_V, \mathcal{S}_U^\beta)$ per Definition 5.

- $\mathcal{S}_V(r_{\text{pub}})$ parses $r_{\text{pub}} = (r_{\mathcal{A}}, r_{\text{qn}}, r_r)$. It runs $\mathcal{A}(r_{\mathcal{A}})$ to obtain

$$\text{param} = \bar{L}, \quad \{\Gamma_j\}_{j \in [J]} \quad (\Gamma_j = (1^{\mathfrak{Q}_j}, \mathbf{\Gamma}_{j,0}, \mathbf{\Gamma}_{j,1}, \xi_j)), \quad \mathbf{x} \in \{0, 1\}^{<\bar{L}}.$$

The algorithm uses r_{qn} to sample q, n as in Setup of Construction 7. It uses r_r to sample $\mathbf{R}_j \in \mathbb{Z}_q^{n \times \mathfrak{Q}_j}$ and $\mathbf{r}_j \in \mathbb{Z}_q^n$ uniformly at random in a straight-forward way²³ for all $j \in [J]$. It computes

$$\mathbf{v}_{\text{msg},j} = \begin{pmatrix} -\mathbf{R}_j \mathbf{t}_1 \\ \mathbf{0}_n \\ 1 \\ \mathbf{0}_{n-1} \end{pmatrix}, \quad \mathbf{v}_{t,j} = \begin{pmatrix} -\mathbf{R}_j \\ \mathbf{r}_j \xi_j^\top \\ \mathbf{0}_{n \times \mathfrak{Q}_j} \end{pmatrix}, \quad \mathbf{v}_{\text{in},j} = \begin{pmatrix} \mathbf{R}_j \mathbf{\Gamma}_{j,0} \\ \mathbf{R}_j \mathbf{\Gamma}_{j,1} \\ -\mathbf{R}_j \end{pmatrix},$$

and sets

$$\begin{aligned} K = 0, \quad Z = 3n, \quad \text{ID} = \{0, 1, 2\}, \quad J_0 = J, \quad J_1 = J_2 = \sum_{j \in [J]} \mathfrak{Q}_j, \\ I_{0,0} = I_{1,0} = 1, \quad I_{2,0} = L, \quad I_{0,g} = I_{1,g} = I_{2,g} = 0, \quad \sigma_{\text{pre}} = \sigma_{\text{IPFE}}, \\ \mathbf{V}_{0,0} = (\mathbf{v}_{\text{msg},1}, \dots, \mathbf{v}_{\text{msg},J}), \quad \mathbf{V}_{1,0} = (\mathbf{v}_{t,1}, \dots, \mathbf{v}_{t,J}), \\ \mathbf{V}_{2,0} = (\mathbf{v}_{\text{in},1}, \dots, \mathbf{v}_{\text{in},J}), \quad \mathbf{V}_{0,g} = \mathbf{V}_{1,g} = \mathbf{V}_{2,g} = \perp. \end{aligned}$$

Lastly, the algorithm outputs the required components.

- $\mathcal{S}_U^\beta(r_{\text{pub}}; r_{\text{priv}})$ first runs the deterministic subprocedure $\mathcal{S}_{A_0}^\beta(r_{\text{pub}})$, which does the following. It first reruns $\mathcal{S}_V(r_{\text{pub}})$ to obtain q, n, \mathbf{x}, L . The algorithm finds matrices $\mathbf{A}_{\text{msg}}, \mathbf{A}_t, \{\mathbf{A}_\ell\}_{\ell \in [L]}$ such that for all $\tilde{\mathbf{s}}_0, \dots, \tilde{\mathbf{s}}_{L+1} \in \mathbb{Z}_q^n$ and $s_{\text{msg}} \in \mathbb{Z}_q$,

$$\begin{aligned} \mathbf{d}_0^\top \mathbf{A}_{\text{msg}} = \mathbf{u}_{\text{msg}}^\top = (\tilde{\mathbf{s}}_0^\top, \mathbf{0}_{1 \times n}, \beta s_{\text{msg}}, \mathbf{0}_{1 \times (n-1)}), \quad \mathbf{d}_0^\top \mathbf{A}_t = \mathbf{u}_t^\top = (\tilde{\mathbf{s}}_L^\top, \tilde{\mathbf{s}}_{L+1}^\top, \mathbf{0}_{1 \times n}), \\ \mathbf{d}_0^\top \mathbf{A}_\ell = \mathbf{u}_\ell^\top = ((1 - \mathbf{x}[\ell]) \tilde{\mathbf{s}}_\ell^\top, \mathbf{x}[\ell] \tilde{\mathbf{s}}_\ell^\top, \tilde{\mathbf{s}}_{\ell-1}^\top) \quad (\text{for all } \ell \in [L]), \end{aligned}$$

where $\mathbf{d}_0^\top = (s_{\text{msg}}, \tilde{\mathbf{s}}_0^\top, \dots, \tilde{\mathbf{s}}_{L+1}^\top)$. The algorithm \mathcal{S}_{A_0} outputs

$$1^{n'} = 1^{n(L+2)+1}, \quad \mathbf{A}_{0,0,0,1} = \mathbf{A}_{\text{msg}}, \quad \mathbf{A}_{1,0,0,1} = \mathbf{A}_t, \quad \mathbf{A}_{2,0,0,\ell} = \mathbf{A}_\ell \quad (\text{for all } \ell \in [L]).$$

Completing \mathcal{S}_{A_0} , the algorithm \mathcal{S}_U samples $\mathbf{d}_0 \xleftarrow{\$} \mathbb{Z}_q^{n'}$ using r_{priv} and outputs

$$\mathbf{u}_{\text{id},0}^\top = \begin{pmatrix} \mathbf{d}_0^\top \mathbf{A}_{\text{id},0,0,1} \\ \vdots \\ \mathbf{d}_0^\top \mathbf{A}_{\text{id},0,0,I_{\text{id},0}} \end{pmatrix} \quad \text{for all id} \in \text{ID}.$$

²²This condition is the same as in Theorem 19, instantiated with $B_{\text{short}} = 2$ and $B_{\text{NLG}} = 0$.

²³Precisely speaking, r_r must be efficiently sampleable conditioned on and given $\mathbf{R}_1, \mathbf{r}_1, \dots, \mathbf{R}_J, \mathbf{r}_J$.

By definition, S^β is a restricted sampler (Definition 6). We have the following:

Claim 26 (¶). S^β has pseudorandom noisy inner products, i.e., $\text{IPFEsec}_{\text{pre}}^{S^\beta}$ (Definition 5) holds for both $\beta \in \{0, 1\}$.

The very selective security of Construction 7 follows from Claim 26. Consider the following hybrids.

- H_0^β is $\text{Exp}_{\text{ABE}}^\beta$ for Construction 7. Note that s_{msg} in the challenge ciphertext is uniformly random over $\mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa}q, 2 \cdot 2^{-\kappa}q]$.
- In H_1^β , the scalar s_{msg} is changed to be uniformly random over \mathbb{Z}_q . Clearly, $H_0^\beta \approx_s H_1^\beta$. Moreover, H_1^β corresponds to the left distribution of $\text{IPFEsec}_{\text{post}}^{S^\beta}$.
- In H_2^β , the vectors inside IPFE ciphertexts (components of the ABE challenge ciphertext) are replaced by random, which corresponds to the right distribution of $\text{IPFEsec}_{\text{post}}^{S^\beta}$. By Claim 26 and the restricted- σ_{IPFE} -security of IPFE, we have $H_1^\beta \approx H_2^\beta$.

Lastly, $H_2^0 \equiv H_2^1$. By hybrid argument, we conclude $\text{Exp}_{\text{ABE}}^0 \approx \text{Exp}_{\text{ABE}}^1$, completing the proof. \square

Proof (Claim 26). Fix $\beta \in \{0, 1\}$ and consider the following hybrids.

- H_0 is the left distribution of $\text{IPFEsec}_{\text{pre}}^{S^\beta}$. Recall that the public randomness is $r_{\text{pub}} = (r_{\mathcal{A}}, r_{\text{qn}}, r_{\mathcal{R}})$. Adopting the notations of the ABE decryption algorithm, the noisy inner products are

$$\begin{aligned} w_{\text{msg},j} &= \mathbf{u}_{\text{msg}}^\top \mathbf{v}_{\text{msg},j} + e_{\text{out},j} = \widetilde{\mathbf{s}}_0^\top \mathbf{R}_j \cdot (-\mathbf{t}_1) + \beta s_{\text{msg}} + e_{\text{out},j}, \\ \mathbf{t}_j^\top &= \mathbf{u}_t^\top \mathbf{v}_{t,j} + \mathbf{e}_{t,j}^\top = -\widetilde{\mathbf{s}}_L^\top \mathbf{R}_j + \widetilde{\mathbf{s}}_{L+1}^\top \mathbf{r}_j \cdot \boldsymbol{\xi}_j^\top + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= \mathbf{u}_\ell^\top \mathbf{v}_{\ell} + \mathbf{e}_{j,\ell}^\top = \widetilde{\mathbf{s}}_\ell^\top \mathbf{R}_j \cdot \boldsymbol{\Gamma}_{\mathbf{x}[\ell]} - \widetilde{\mathbf{s}}_{\ell-1}^\top \mathbf{R}_j + \mathbf{e}_{j,\ell}^\top \quad (\text{for all } \ell \in [L]), \end{aligned}$$

where $j \in [J]$ and the entries of e , \mathbf{e} 's are independent $\mathcal{D}_{\mathbb{Z}, \sigma_{\text{IPFE}}}$.

- In H_1 , additional small noises are attached to $\widetilde{\mathbf{s}}^\top \mathbf{R}$, $\widetilde{\mathbf{s}}^\top \mathbf{r}$'s. For all $j \in [J]$, set the inner products to

$$\begin{aligned} w_{\text{msg},j} &= (\widetilde{\mathbf{s}}_0^\top \mathbf{R}_j + \mathbf{e}_{r,j,0}^\top)(-\mathbf{t}_1) + \beta s_{\text{msg}} + e_{\text{out},j}, \\ \mathbf{t}_j^\top &= -(\widetilde{\mathbf{s}}_L^\top \mathbf{R}_j + \mathbf{e}_{r,j,L}^\top) + (\widetilde{\mathbf{s}}_{L+1}^\top \mathbf{r}_j + \mathbf{e}_{r,j,L+1}^\top) \boldsymbol{\xi}_j^\top + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= (\widetilde{\mathbf{s}}_\ell^\top \mathbf{R}_j + \mathbf{e}_{r,j,\ell}^\top) \boldsymbol{\Gamma}_{\mathbf{x}[\ell]} - (\widetilde{\mathbf{s}}_{\ell-1}^\top \mathbf{R}_j + \mathbf{e}_{r,j,\ell-1}^\top) + \mathbf{e}_{j,\ell}^\top \quad (\text{for all } \ell \in [L]), \end{aligned}$$

where $\mathbf{e}_{r,j,\ell} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_r, \leq \sigma_r \sqrt{\kappa}}^{\boxplus j}$ for all $\ell \in \{0, \dots, L\}$ and $e_{r,j,L+1} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_r, \leq \sigma_r \sqrt{\kappa}}$. The errors introduced into the inner products are bounded by

$$2\sigma_r \sqrt{\kappa} \leq 2^{-\kappa-6} \sigma_{\text{IPFE}},$$

so they are flooded (Lemma 2) by $\{e_{\text{out},j}, \mathbf{e}_{t,j}, \{\mathbf{e}_{j,\ell}\}_{\ell \in [L]}\}_{j \in [J]}$, and $H_0 \approx_s H_1$.

- In H_2 , the \mathbf{e}_r 's are no longer truncated, i.e., for all $j \in [J]$, sample $\mathbf{e}_{r,j,\ell} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_r}^{\boxplus j}$ for all $\ell \in \{0, \dots, L\}$ and $e_{r,j,L+1} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}, \sigma_r}$. We have $H_1 \approx_s H_2$ by Lemma 1.

- In H_3 , the garblings use uniform randomness. For all $j \in [J]$, sample $\mathbf{s}_{j,\ell} \xleftarrow{\$} \mathbb{Z}_q^n$ for all $\ell \in \{0, \dots, L\}$ and $s_{j,L+1} \xleftarrow{\$} \mathbb{Z}_q$, and set the inner products to

$$\begin{aligned} w_{\text{msg},j} &= \mathbf{s}_{j,0}^\top \cdot (-\mathbf{t}_1) + \beta s_{\text{msg}} + e_{\text{out},j}, \\ \mathbf{t}_j^\top &= -\mathbf{s}_{j,L}^\top + s_{j,L+1} \boldsymbol{\xi}_j^\top + \mathbf{e}_{t,j}^\top, \\ \mathbf{w}_{j,\ell}^\top &= \mathbf{s}_{j,\ell}^\top \mathbf{\Gamma}_{\mathbf{x}[\ell]} - \mathbf{s}_{j,\ell-1}^\top + \mathbf{e}_{j,\ell}^\top \quad (\text{for all } \ell \in [L]), \end{aligned}$$

By $(L+1)$ hybrids of applying $\text{LWE}_{n,J_2,q,\sigma_r}$ with public matrix $(\mathbf{R}_1, \dots, \mathbf{R}_J)$ and secrets $(\tilde{\mathbf{s}}_0, \dots, \tilde{\mathbf{s}}_L)$, then $\text{LWE}_{n,J,q,\sigma_r}$ with public matrix $(\mathbf{r}_1, \dots, \mathbf{r}_J)$ and secret $\tilde{\mathbf{s}}_{L+1}$, we have

$$\left\{ \left\{ \tilde{\mathbf{s}}_\ell^\top \mathbf{R}_j + \mathbf{e}_{r,j,\ell}^\top \right\}_{\ell \in \{0, \dots, L\}}, \left\{ \tilde{\mathbf{s}}_{L+1}^\top \mathbf{r}_j + e_{r,j,L+1} \right\}_{j \in [J]} \right\} \approx \left\{ \left\{ \mathbf{s}_{j,\ell}^\top \right\}_{\ell \in \{0, \dots, L\}}, s_{j,L+1} \right\}_{j \in [J]}$$

which implies $H_2 \approx H_3$. In this step, we rely on the “straight-forwardness” of sampling \mathbf{R}, \mathbf{r} 's from r_r in \mathcal{S}^β .

- In H_4 , all the inner products are replaced by random, i.e., the right distribution of $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}^\beta}$. Note that the inner products in H_3 are well-randomized garblings (by \mathbf{s} 's) plus independent noises (e, \mathbf{e} 's). Therefore, by J hybrids of applying the security of NLG, we conclude $H_3 \equiv H_4$.

By hybrid argument, $H_0 \approx H_4$, i.e., $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}^\beta}$ holds. \square

7.4 CP-ABE for DFA and Summary

Our CP-ABE for DFA is similar to KP-ABE for DFA and general CP-ABE and features a mixture of how components are set in both schemes.

Ingredients of Construction 8. We rely on

- Construction 6, a specific noisy linear garbling NLG for DFA, and
- an identity-based evasive IPFE scheme IPFE for $\mathcal{I} \supseteq \{0, 1, 2\}$ that is B_{IPFE} -correct and restricted- σ_{IPFE} -secure (Definitions 4 and 6).

Construction 8 (CP-ABE for DFA). Let

$$\begin{aligned} \text{Params} &= \{ \bar{L} \mid \bar{L} \in \mathbb{N} \}, & X_{\bar{L}} &= \{ \mathbf{x} \in \{0, 1\}^L \mid L \leq \bar{L} \}, & Y_{\bar{L}} &= \{ \text{DFA } \Gamma \}, \\ P_{\bar{L}}(\mathbf{x}, \Gamma) &= \begin{cases} 1, & \text{if } \Gamma \text{ accepts } \mathbf{x}; \\ 0, & \text{if } \Gamma \text{ rejects } \mathbf{x}. \end{cases} \end{aligned}$$

Our CP-ABE for DFA works as follows.

- $\text{Setup}(\bar{L})$ picks suitable q, n in a straight-forward way. It sets $\text{pp} = (q, \bar{L}), K = 0, Z = 3n$. The algorithm runs $(\text{impk}, \text{msk}) \xleftarrow{\$} \text{IPFE.Setup}(q, 1^K, 1^Z)$ and outputs

$$\text{mpk} = (\text{pp}, 1^n, 1^K, 1^Z, \text{impk}), \quad \text{msk} = (\text{mpk}, \text{msk}).$$

- $\text{KeyGen}(\text{msk}, \mathbf{x})$, given $\mathbf{x} \in \{0, 1\}^L$, samples and sets

$$\mathbf{r}_0, \dots, \mathbf{r}_L \xleftarrow{\$} \mathbb{Z}_q^n, \quad r_{L+1} \xleftarrow{\$} \mathbb{Z}_q,$$

$$\mathbf{v}_{\text{msg}} = \begin{pmatrix} \mathbf{r}_0 \\ 1 \\ \mathbf{0}_{2n-1} \end{pmatrix}, \quad \mathbf{v}_t = \begin{pmatrix} \mathbf{r}_L \\ r_{L+1} \\ \mathbf{0}_{2n-1} \end{pmatrix}, \quad \mathbf{v}_\ell = \begin{pmatrix} (1 - \mathbf{x}[\ell])\mathbf{r}_\ell \\ \mathbf{x}[\ell]\mathbf{r}_\ell \\ \mathbf{r}_{\ell-1} \end{pmatrix} \quad \text{for all } \ell \in [L].$$

It generates IPFE secret keys

$$\begin{aligned} \text{isk}_{\text{msg}} &\stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, 0, \mathbf{v}_{\text{msg}}, \perp), \\ \text{isk}_t &\stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, 1, \mathbf{v}_t, \perp), \\ \text{isk}_\ell &\stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, 2, \mathbf{v}_\ell, \perp) \quad \text{for all } \ell \in [L]. \end{aligned}$$

The algorithm outputs $\text{sk} = (\mathbf{r}_0, \dots, \mathbf{r}_L, r_{L+1}, \text{isk}_{\text{msg}}, \text{isk}_t, \text{isk}_1, \dots, \text{isk}_L)$.

- $\text{Enc}(\text{mpk}, \Gamma, \mu)$ parses $\Gamma = (1^\Omega, \mathbf{\Gamma}_0, \mathbf{\Gamma}_1, \boldsymbol{\xi})$. It samples

$$\mathbf{S} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{\Omega \times n}, \quad s_{-1} \stackrel{\$}{\leftarrow} \mathbb{Z}_q, \quad s_{\text{msg}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa}q, 2 \cdot 2^{-\kappa}q],$$

and sets

$$\mathbf{u}_{\text{msg}}^\top = (-\mathbf{t}_1^\top \mathbf{S}, \mu s_{\text{msg}}, \mathbf{0}_{1 \times (2n-1)}), \quad \mathbf{U}_t^\top = (-\mathbf{S}, s_{-1} \boldsymbol{\xi}, \mathbf{0}_{\Omega \times (2n-1)}), \quad \mathbf{U}_{\text{in}}^\top = (\mathbf{\Gamma}_0^\top \mathbf{S}, \mathbf{\Gamma}_1^\top \mathbf{S}, -\mathbf{S}).$$

The algorithm generates IPFE ciphertexts

$$\begin{aligned} \text{ict}_{\text{msg}} &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 0, \mathbf{v}, \mathbf{u}_{\text{msg}}^\top), \\ \text{ict}_t &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 1, \mathbf{v}, \mathbf{U}_t^\top), \\ \text{ict}_{\text{in}} &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 2, \mathbf{v}, \mathbf{U}_{\text{in}}^\top), \end{aligned}$$

and outputs $\text{ct} = (\text{ict}_{\text{msg}}, \text{ict}_t, \text{ict}_{\text{in}})$.

- $\text{Dec}(\text{mpk}, \mathbf{x}, \text{sk}, \Gamma, \text{ct})$ outputs \perp and terminates if Γ rejects \mathbf{x} . Otherwise, it parses sk, ct as in $\text{KeyGen}, \text{Enc}$. Suppose $\mathbf{x} \in \{0, 1\}^L$, the algorithm recomputes $\mathbf{v}_{\text{msg}}, \mathbf{v}_t, \mathbf{v}_1, \dots, \mathbf{v}_L$ as in KeyGen and performs IPFE decryptions

$$\begin{aligned} w_{\text{msg}} &\leftarrow \text{IPFE.Dec}(\text{impk}, 0, \mathbf{v}_{\text{msg}}, \perp, \text{isk}_{\text{msg}}, \mathbf{v}, \text{ict}_{\text{msg}}), \\ \mathbf{t} &\leftarrow \text{IPFE.Dec}(\text{impk}, 1, \mathbf{v}_t, \perp, \text{isk}_t, \mathbf{v}, \text{ict}_t), \\ \mathbf{w}_\ell &\leftarrow \text{IPFE.Dec}(\text{impk}, 2, \mathbf{v}_\ell, \perp, \text{isk}_\ell, \mathbf{v}, \text{ict}_{\text{in}}) \quad \text{for all } \ell \in [L]. \end{aligned}$$

It evaluates the DFA garbling

$$w_{\text{out}} \leftarrow \text{NLG.Eval}(\text{pp}, (\Gamma, 1^L), \perp, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)$$

and outputs

$$\mu' \leftarrow \begin{cases} 0, & \text{if } w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa}q, 2^{-\kappa}q]; \\ 1, & \text{otherwise.} \end{cases}$$

Theorem 27. *Construction 8 is correct (Definition 1) if $(\bar{L} + 2)B_{\text{IPFE}}(q, 0, 3n) \leq 2^{-\kappa}q$.*

Theorem 28. *Construction 8 is very selectively secure (Definition 2) if $\text{LWE}_{n, \text{poly}(\lambda), q, \sigma_r}$ (Assumption 1) holds for some $\sigma_r \leq \frac{2^{-\kappa-6} \sigma_{\text{IPFE}}}{2\sqrt{\kappa}}$.*

Summary. By instantiating our ABE schemes with the IPFE scheme of Construction 2, we obtain KP- and CP-ABE for DFA.

Corollary 29. Under LWE (Assumption 1) and evasive LWE (Assumption 4), there exists KP-ABE for DFA with

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}_\Gamma| = (|\Gamma| + 1) \text{poly}(\lambda), \quad |\text{ct}_\mathbf{x}| = (|\mathbf{x}| + 1) \text{poly}(\lambda).$$

and CP-ABE for DFA with

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}_\mathbf{x}| = (|\mathbf{x}| + 1) \text{poly}(\lambda), \quad |\text{ct}_\Gamma| = (|\Gamma| + 1) \text{poly}(\lambda).$$

We remark that, starting from the linear secret sharing scheme [LL20a] for L, we also obtain ABE for L via the same construction. However, ABE for NFA or NL cannot be obtained so. The schemes for DFA and L crucially rely on the fact that, when noises are added to the LSSS shares, the output noise grows *linearly* with the computation size. On the other hand, the secret sharing schemes for NFA and NL, once cast as noisy linear garbling, incur exponential noise growth during reconstruction. The latter fact makes the resultant ABE bounded.

8 CP-ABE with Succinct Ciphertexts

In this section, we show that with the help of evasive IPFE with structured noises, we can construct CP-ABE from general noisy linear garbling scheme (not necessarily short). When instantiated with the noisy linear garbling implicit in [BGG⁺14], this yields a CP-ABE scheme with ciphertext size dependent only on circuit depth instead of circuit size.

8.1 Succinct Noisy Linear Garbling for Circuits

The work of [BGG⁺14] presented a KP-ABE scheme with succinct keys for circuits based on lattices. At the core of their construction are a pair of deterministic homomorphic evaluation algorithms $\text{EvalC}, \text{EvalCX}$, which, for circuit $C : \{0, 1\}^L \rightarrow \{0, 1\}$, input $x \in \{0, 1\}^L$, and public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m(L+1)}$, satisfies the relation

$$\begin{aligned} \text{EvalC}(\mathbf{A}, C) &= \mathbf{A}_C \in \mathbb{Z}_q^{n \times m}, \\ \text{EvalCX}(\mathbf{s}^\top(\mathbf{A} - (\mathbf{1}, \mathbf{x}^\top) \otimes \mathbf{G}) + \mathbf{e}^\top, f, \mathbf{x}) &= \mathbf{s}^\top(\mathbf{A}_C - C(\mathbf{x}) \cdot \mathbf{G}) + \mathbf{e}_C^\top, \end{aligned}$$

where $\|\mathbf{e}_C\|$, relative to $\|\mathbf{e}\|$, grows exponentially with the depth of C . (In [BGG⁺14], it also works for bounded-arithmetic circuits, but we only present the Boolean version.) It naturally yields a noisy linear garbling scheme.

Construction 9 (garbling adapted from [BGG⁺14]). Let

$$\begin{aligned} \text{Params} &= \{1^d \mid d \in \mathbb{N}\}, \\ F_{1^d} &= \{f_C \mid C \text{ is a Boolean circuit of depth no more than } d\}, \\ f_C(\mathbf{x}) &= \begin{cases} \perp, & \text{if } \mathbf{x} \notin \{0, 1\}^L; \\ -C(\mathbf{x}), & \text{if } \mathbf{x} \in \{0, 1\}^L; \end{cases} \quad \text{for } C : \{0, 1\}^L \rightarrow \{0, 1\} \text{ and } \mathbf{x} \in \mathbb{Z}^L. \end{aligned}$$

The noisy linear garbling scheme works as follows.

- $\text{Setup}(1^d)$ picks suitable q, n, m (with $m \geq n \lceil \log_2 q \rceil$) and outputs $\text{pp} = (q, 1^n, 1^m)$.
- $\text{GenF}(\text{pp}, C)$ samples $\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_L \xleftarrow{\$} \mathbb{Z}_q^{n \times m(L+1)}$ and $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n$. It computes $\mathbf{A}_C \leftarrow \text{EvalC}((\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_L), C)$. The algorithm outputs

$$1^{n_C} = 1^n, \quad \mathbf{w}_{\text{out}} = \mathbf{A}_C \mathbf{G}^{-1}(\mathbf{a}), \quad \mathbf{W}_{\ell,0} = \mathbf{A}_\ell, \quad \mathbf{W}_{\ell,\times} = -\mathbf{G}, \quad (\text{for all } \ell \in [L])$$

$$\mathbf{T} = \mathbf{A}_0 - \mathbf{G}, \quad \mathbf{R} = \mathbf{a}.$$

All the samplings are done in a straight-forward way.

- $\text{Eval}(\text{pp}, C, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)$ computes and outputs

$$\text{EvalCX}((\mathbf{t}^\top, \mathbf{w}_1^\top, \dots, \mathbf{w}_L^\top), C, \mathbf{x}) \cdot \mathbf{G}^{-1}(\mathbf{a}).$$

Lemma 30 ([BGG⁺14]). *Construction 9 is correct (Definition 7) if*

$$(m+2)^{d+1} B_{\text{in}}(d) \leq B_{\text{out}}(d).$$

It has fixed randomness dimension (Definition 9). Suppose

$$\text{GenGaussian}'(1^d, (q, n, m), C, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}, \mathbf{x})$$

samples (truncated) Gaussian noises with appropriate width σ of suitable shape

$$(e_{\text{out}}, \mathbf{e}_1, \dots, \mathbf{e}_L, \mathbf{e}_t) \xleftarrow{\$} \mathcal{D}_{\mathbb{Z}, \sigma, \leq \sigma \sqrt{\kappa}}^{1+(L+1)m}$$

then Construction 9 is GenGaussian'-secure (Definition 10) if $\text{LWE}_{n, \text{poly}(\lambda), q, \sigma'}$ (Assumption 2) holds for some $\sigma' \leq \frac{\sigma}{2^{\kappa+6} (m+2)^{d+1} \sqrt{\kappa}}$.

8.2 CP-ABE from General Noisy Linear Garbling

We present a construction of CP-ABE from noisy linear garbling with fixed randomness dimension that is not necessarily short. It has a lot in common with Construction 5. Both compute rerandomized garbling using IPFE. In Construction 5, security relies on flooding to introduce noises to argue that the garblings use good pseudorandomness. Here, we use structured noises, together with some algebraic tweaks (“noisy secret sharing”, see Section 1.2), to attach the same noise to the garbling pseudorandomness.

Ingredients of Construction 10. We rely on

- a noisy linear garbling scheme NLG supporting $F = \{F_{\text{param}'}\}_{\text{param}' \in \text{Params}'}$ that is $(B_{\text{in}}, B_{\text{out}})$ -correct, has fixed randomness dimension, and is GenNoise-secure such that the output of GenNoise is in $[-B_{\text{NLG}}, B_{\text{NLG}}]^*$ (Definitions 7, 9, and 10), and
- an identity-based evasive IPFE with structured noises scheme IPFE for identity space $\mathcal{I} \supseteq \{-1, 0, 1, \dots, L\}$ that is B_{IPFE} -correct and restricted- σ_{IPFE} -secure (Definitions 4 and 6).

Construction 10 (CP-ABE from general garbling). Define

$$\begin{aligned} \text{Params} &= \{ \text{param} = (\text{param}', 1^L) \mid \text{param}' \in \text{Params}', L \in \mathbb{N} \}, & X_{\text{param}', 1^L} &= \{0, 1\}^L, \\ Y_{\text{param}', 1^L} &= \{ f : \mathbb{Z}^L \rightarrow \{0, 1, \perp\} \mid f \in \mathcal{F}_{\text{param}'} \}, & P_{\text{param}', 1^L}(\mathbf{x}, f) &= f(\mathbf{x}). \end{aligned}$$

(Note that \mathbf{x} is confined to Boolean.) Our CP-ABE scheme works as follows.

- Setup(param', 1^L) runs

$$\text{pp} = (q, 1^n, \dots) \stackrel{\$}{\leftarrow} \text{NLG.Setup}(\text{param}').$$

The algorithm sets $K = \lceil \log_2 q \rceil$ to be the dimension of $\tilde{\mathbf{g}}$ (for structured noises), and $Z = 2(n + n^2K)$. It runs (impk, imsk) $\stackrel{\$}{\leftarrow}$ IPFE.Setup($q, 1^K, 1^Z$) and outputs

$$\text{mpk} = (\text{pp}, 1^K, 1^Z, \text{impk}), \quad \text{msk} = (\text{mpk}, \text{imsk}).$$

- KeyGen(msk, \mathbf{x}) samples $\mathbf{r} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^n$ and $\mathbf{Q} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times nK}$. Recall that $\text{col}(\mathbf{Q})$ is a column of dimension n^2K concatenating the columns of \mathbf{Q} . The algorithm sets

$$\mathbf{V}_Q = \begin{pmatrix} \mathbf{Q} \\ \mathbf{0}_{(Z-n) \times nK} \end{pmatrix}, \quad \mathbf{v}_0 = \begin{pmatrix} \mathbf{r} \\ \text{col}(\mathbf{Q}) \\ 1 \\ \mathbf{0}_{n-1+n^2K} \end{pmatrix}, \quad \mathbf{v}_\ell = \begin{pmatrix} (1 - \mathbf{x}[\ell])\mathbf{r} \\ (1 - \mathbf{x}[\ell])\text{col}(\mathbf{Q}) \\ \mathbf{x}[\ell]\mathbf{r} \\ \mathbf{x}[\ell]\text{col}(\mathbf{Q}) \end{pmatrix} \text{ for all } \ell \in [L].$$

It generates IPFE secret keys

$$\begin{aligned} \text{isk}_{-1} &\stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, -1, \mathbf{0}, \mathbf{V}_Q), \\ \text{isk}_0 &\stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, 0, \mathbf{v}_0, \mathbf{0}), \\ \text{isk}_\ell &\stackrel{\$}{\leftarrow} \text{IPFE.KeyGen}(\text{imsk}, \ell, \mathbf{v}_\ell, \mathbf{0}) \quad \text{for all } \ell \in [L], \end{aligned}$$

and outputs $\text{sk} = (\mathbf{r}, \mathbf{Q}, \text{isk}_{-1}, \text{isk}_0, \text{isk}_1, \dots, \text{isk}_L)$.

- Enc(mpk, f, μ) runs

$$(1^n, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}) \stackrel{\$}{\leftarrow} \text{GenF}(\text{pp}, f).$$

It samples $\mathbf{S} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{n \times n}$, $\mathbf{s}_Q \stackrel{\$}{\leftarrow} \mathbb{Z}_q^n$, $s_{\text{msg}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa}q, 2 \cdot 2^{-\kappa}q]$, and sets

$$\begin{aligned} \mathbf{u}_Q^\top &= (\mathbf{s}_Q^\top, \mathbf{0}_{1 \times (Z-n)}), \\ \mathbf{u}_{\text{msg}}^\top &= (\mathbf{w}_{\text{out}}^\top \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}) \otimes \mathbf{s}_Q)^\top, \mu s_{\text{msg}}, \mathbf{0}_{1 \times (n-1+n^2K)}), \\ \mathbf{U}_t^\top &= (\mathbf{T}^\top \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{T}) \otimes \mathbf{s}_Q)^\top, \mathbf{0}_{\star \times (n+n^2K)}), \\ \mathbf{U}_\ell^\top &= (\mathbf{W}_{\ell,0}^\top \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0}) \otimes \mathbf{s}_Q)^\top, (\mathbf{W}_{\ell,0}^\top + \mathbf{W}_{\ell,\times}^\top) \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{W}_{\ell,\times}) \otimes \mathbf{s}_Q)^\top), \end{aligned}$$

where $\ell \in [L]$. The algorithm generates IPFE ciphertexts

$$\begin{aligned} \text{ict}_Q &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, -1, \mathbf{g}, \mathbf{u}_Q^\top), \\ \text{ict}_{\text{msg}} &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 0, \mathbf{v}, \mathbf{u}_{\text{msg}}^\top), & \text{ict}_t &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, 0, \mathbf{v}, \mathbf{U}_t^\top), \\ \text{ict}_\ell &\stackrel{\$}{\leftarrow} \text{IPFE.Enc}(\text{impk}, \ell, \mathbf{v}, \mathbf{U}_\ell^\top) \quad \text{for all } \ell \in [L], \end{aligned}$$

and outputs $\text{ct} = (\mathbf{R}, \text{ict}_Q, \text{ict}_{\text{msg}}, \text{ict}_t, \text{ict}_1, \dots, \text{ict}_L)$.

- $\text{Dec}(\text{mpk}, \mathbf{x}, \text{sk}, f, \text{ct})$ outputs \perp and terminates if $f(\mathbf{x}) \neq 1$. Otherwise, it parses sk, ct as in $\text{KeyGen}, \text{Enc}$, recomputes $\mathbf{V}_Q, \mathbf{v}_0, \{\mathbf{v}_\ell\}_{\ell \in [L]}$, and performs IPFE decryptions for the shares

$$\begin{aligned} \text{sh}_Q^\top &\leftarrow \text{IPFE.Dec}(\text{impk}, -1, \mathbf{0}, \mathbf{V}_Q, \text{isk}_{-1}, \mathbf{g}, \text{ict}_Q), \\ \text{sh}_{\text{msg}} &\leftarrow \text{IPFE.Dec}(\text{impk}, 0, \mathbf{v}_0, \mathbf{0}, \text{isk}_0, \mathbf{0}, \text{ict}_{\text{msg}}), \\ \text{sh}_t &\leftarrow \text{IPFE.Dec}(\text{impk}, 0, \mathbf{v}_0, \mathbf{0}, \text{isk}_0, \mathbf{0}, \text{ict}_t), \\ \text{sh}_\ell &\leftarrow \text{IPFE.Dec}(\text{impk}, \ell, \mathbf{v}_\ell, \mathbf{0}, \text{isk}_\ell, \mathbf{0}, \text{ict}_\ell) \quad \text{for all } \ell \in [L]. \end{aligned}$$

The algorithm reconstructs the garbling from the shares as

$$\begin{aligned} w_{\text{msg}} &\leftarrow \text{sh}_{\text{msg}} + \text{sh}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), & \mathbf{t}^\top &\leftarrow \text{sh}_t^\top + \text{sh}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\ \mathbf{w}_\ell^\top &\leftarrow \text{sh}_\ell^\top + \text{sh}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,x}) & & \text{for all } \ell \in [L]. \end{aligned}$$

It evaluates the garbling by

$$w_{\text{out}} \leftarrow \text{NLG.Eval}(\text{pp}, f, \mathbf{R}, \mathbf{x}, \{\mathbf{w}_\ell^\top\}_{\ell \in [L]}, \mathbf{t}^\top)$$

and outputs

$$\mu' \leftarrow \begin{cases} 0, & \text{if } w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa}q, 2^\kappa q]; \\ 1, & \text{otherwise.} \end{cases}$$

Theorem 31 (¶). *Construction 5 is correct (Definition 1) if*

- the modulus q output by $\text{NLG.Setup}(\text{param}')$ is always a prime and always satisfies $B_{\text{out}}(\text{param}') \leq 2^{-\kappa}q$, and
- it holds that $(nK + 1)B_{\text{IPFE}}(q, K, 2(n + n^2K)) \leq B_{\text{in}}(\text{param}')$.

Proof (Theorem 31). By the correctness of IPFE and how we set the vectors, we have

$$\text{sh}_Q^\top = \mathbf{s}_Q^\top \mathbf{Q} + \tilde{\mathbf{g}}^\top \otimes \mathbf{e}_{\text{sr}}^\top + \mathbf{e}_Q^\top = \mathbf{s}_Q^\top \mathbf{Q} + \mathbf{e}_{\text{sr}}^\top \tilde{\mathbf{G}} + \mathbf{e}_Q^\top$$

for some B_{IPFE} -bounded errors $\mathbf{e}_{\text{sr}}, \mathbf{e}_Q$. For any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \star}$,

$$\text{sh}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{A}) = (\mathbf{s}_Q^\top \mathbf{Q} + \mathbf{e}_{\text{sr}}^\top \tilde{\mathbf{G}} + \mathbf{e}_Q^\top) \tilde{\mathbf{G}}^{-1}(\mathbf{A}) = \underbrace{\mathbf{s}_Q^\top \mathbf{Q} \tilde{\mathbf{G}}^{-1}(\mathbf{A})}_{\text{one-time pad}} + \underbrace{\mathbf{e}_{\text{sr}}^\top \mathbf{A}}_{\text{"attached" error}} + \underbrace{\mathbf{e}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{A})}_{\text{small error}}.$$

For the table shares, we have

$$\begin{aligned} \text{sh}_t &= \mathbf{U}_t^\top \mathbf{v}_0 + \mathbf{e}_t = \mathbf{T}^\top \mathbf{S}_r - ((\tilde{\mathbf{G}}^{-1}(\mathbf{T}))^\top \otimes \mathbf{s}_Q^\top) \text{col}(\mathbf{Q}) + \mathbf{e}_t \\ &= \mathbf{T}^\top \mathbf{S}_r - \text{col}(\mathbf{s}_Q^\top \mathbf{Q} \tilde{\mathbf{G}}^{-1}(\mathbf{T})) + \mathbf{e}_t = ((\mathbf{S}_r)^\top \mathbf{T} + \mathbf{e}_t^\top - \mathbf{s}_Q^\top \mathbf{Q} \tilde{\mathbf{G}}^{-1}(\mathbf{T}))^\top, \end{aligned}$$

where the third equality follows from $\text{col}(\mathbf{ABC}) = (\mathbf{C}^\top \otimes \mathbf{A}) \text{col}(\mathbf{B})$ and the last equality holds because $\mathbf{s}_Q^\top \mathbf{Q} \tilde{\mathbf{G}}^{-1}(\mathbf{T})$ is a row vector. The reconstructed garbled table is hence

$$\mathbf{t}^\top = \underbrace{(\mathbf{S}_r)^\top \mathbf{T} + \mathbf{e}_t^\top - \mathbf{s}_Q^\top \mathbf{Q} \tilde{\mathbf{G}}^{-1}(\mathbf{T})}_{\text{sh}_t^\top} + \underbrace{\mathbf{s}_Q^\top \mathbf{Q} \tilde{\mathbf{G}}^{-1}(\mathbf{T}) + \mathbf{e}_{\text{sr}}^\top \mathbf{T} + \mathbf{e}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T})}_{\text{sh}_Q^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T})} = (\mathbf{S}_r + \mathbf{e}_{\text{sr}})^\top \mathbf{T} + \tilde{\mathbf{e}}_t^\top,$$

where $\tilde{\mathbf{e}}_t = \mathbf{e}_t + (\tilde{\mathbf{G}}^{-1}(\mathbf{T}))^\top \mathbf{e}_Q$ is $(nK + 1)B_{\text{IPFE}}$ -bounded. Similarly,

$$\begin{aligned} w_{\text{msg}} &= (\mathbf{S}\mathbf{r} + \mathbf{e}_{\text{sr}})^\top \mathbf{w}_{\text{out}} + \mu s_{\text{msg}} + \tilde{e}_{\text{out}}, \\ \mathbf{w}_\ell^\top &= (\mathbf{S}\mathbf{r} + \mathbf{e}_{\text{sr}})^\top (\mathbf{W}_{\ell,0} + \mathbf{x}[\ell] \mathbf{W}_{\ell,\times}) + \tilde{\mathbf{e}}_\ell^\top \quad \text{for all } \ell \in [L], \end{aligned}$$

with $\tilde{e}_{\text{out}}, \tilde{\mathbf{e}}_\ell$'s bounded by $(nK + 1)B_{\text{IPFE}}$. The garbling uses secret $\mathbf{S}\mathbf{r} + \mathbf{e}_{\text{sr}} \in \mathbb{Z}_q^n$ and its noises ($\tilde{\mathbf{e}}$'s) are bounded by $(nK + 1)B_{\text{IPFE}}(q, K, Z) \leq B_{\text{in}}(\text{param}')$. Therefore, by the correctness of NLG,

$$w_{\text{out}} - ((\mathbf{S}\mathbf{r} + \mathbf{e}_{\text{sr}})^\top \mathbf{w}_{\text{out}} + e'_{\text{out}}) \in [-B_{\text{out}}(\text{param}'), B_{\text{out}}(\text{param}')] \subseteq [-2^{-\kappa}q, 2^{-\kappa}q].$$

If $\mu = 0$, it always holds that $w_{\text{msg}} - w_{\text{out}} \in [-2^{-\kappa}q, 2^{-\kappa}q]$, which never holds when $\mu = 1$ by our choice of s_{msg} . We conclude that Construction 10 is correct. \square

8.3 Security and Summary

Theorem 32 (\spadesuit). *Construction 10 is very selectively secure (Definition 2) if*

- $\text{LWE}_{n, \text{poly}(\lambda), q, \sigma}$ (Assumption 1) holds for some $\sigma \leq \frac{2^{-\kappa-6}\sigma_{\text{IPFE}}}{nK\sqrt{\kappa}}$, and
- it holds that $2^{\kappa+6}B_{\text{NLG}}(\text{param}') \leq \sigma_{\text{IPFE}}$.

Proof (Theorem 32). Let \mathcal{A} be an efficient adversary and assume that it always chooses challenges such that $f(\mathbf{x}_j) = 0$ for all $j \in [J]$. Consider the following evasive IPFE sampler $\mathcal{S}^\beta = (\mathcal{S}_V, \mathcal{S}_U^\beta)$ per Definition 5.

- $\mathcal{S}_V(r_{\text{pub}})$ parses $r_{\text{pub}} = (r_{\mathcal{A}}, r_{\text{NLG.Setup}}, r_{\text{rq}}, r_{\text{NLG.GenF}})$. It runs $\mathcal{A}(r_{\mathcal{A}})$ to obtain

$$\text{param} = (\text{param}', 1^L), \quad \{\mathbf{x}_j\}_{j \in [J]} \quad (\mathbf{x}_j \in \{0, 1\}^L), \quad f : \mathbb{Z}^L \rightarrow \{0, 1, \perp\}.$$

The algorithm next sets up the noisy linear garbling by

$$\text{pp} = (q, 1^n, \dots) \leftarrow \text{NLG.Setup}(\text{param}'; r_{\text{NLG.Setup}}),$$

and computes K, Z as in Setup of Construction 10. It then uses r_{rq} to sample $\mathbf{r}_1, \dots, \mathbf{r}_J \in \mathbb{Z}_q^n$, $\mathbf{Q}_1, \dots, \mathbf{Q}_J \in \mathbb{Z}_q^{n \times nK}$ uniformly at random in a straight-forward way.²⁴ The algorithm sets for all $j \in [J]$,

$$\mathbf{V}_{Q,j} = \begin{pmatrix} \mathbf{Q}_j \\ \mathbf{0}_{(Z-n) \times nK} \end{pmatrix}, \quad \mathbf{v}_{j,0} = \begin{pmatrix} \mathbf{r}_j \\ \text{col}(\mathbf{Q}_j) \\ 1 \\ \mathbf{0}_{n-1+n^2K} \end{pmatrix}, \quad \mathbf{v}_{j,\ell} = \begin{pmatrix} (1 - \mathbf{x}_j[\ell])\mathbf{r}_j \\ (1 - \mathbf{x}_j[\ell])\text{col}(\mathbf{Q}_j) \\ \mathbf{x}_j[\ell]\mathbf{r}_j \\ \mathbf{x}_j[\ell]\text{col}(\mathbf{Q}_j) \end{pmatrix} \quad \text{for all } \ell \in [L].$$

It runs

$$(1^n, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}) \leftarrow \text{NLG.GenF}(\text{pp}, f; r_{\text{NLG.GenF}}).$$

Suppose $\mathbf{W}_{\ell,0} \in \mathbb{Z}_q^{n \times m_\ell}$ and $\mathbf{T} \in \mathbb{Z}_q^{n \times m_t}$, the algorithm outputs

$$\begin{aligned} q, 1^K, 1^Z, & & \text{ID} = \{-1, 0, 1, \dots, L\}, & & 1^{J_{\text{id}}} = 1^J, \\ 1^{I_{-1,0}} = 1^0, & & 1^{I_{0,0}} = 1^{m_t+1}, & & 1^{I_{\ell,0}} = 1^{m_\ell}, \\ 1^{I_{-1,9}} = 1^1, & & 1^{I_{0,9}} = 1^{I_{\ell,9}} = 1^0, & & \sigma_{\text{pre}} = \sigma_{\text{IPFE}}, \end{aligned}$$

²⁴Precisely speaking, r_{rq} must be efficiently sampleable conditioned on and given $\mathbf{r}_1, \dots, \mathbf{r}_J$ and $\mathbf{Q}_1, \dots, \mathbf{Q}_J$.

$$\begin{aligned}
\mathbf{v}_{-1,\mathbf{g}} &= (\mathbf{v}_{Q,1}, \dots, \mathbf{v}_{Q,J}), & \mathbf{v}_{-1,\mathbf{o}} &= \mathbf{0}, \\
\mathbf{V}_{0,\mathbf{o}} &= (\mathbf{v}_{0,0}, \dots, \mathbf{v}_{J,0}), & \mathbf{V}_{0,\mathbf{g}} &= \mathbf{0}, \\
\mathbf{V}_{\ell,\mathbf{o}} &= (\mathbf{v}_{1,\ell}, \dots, \mathbf{v}_{J,\ell}), & \mathbf{V}_{\ell,\mathbf{g}} &= \mathbf{0},
\end{aligned}$$

where the indices are $\text{id} \in \text{ID}$ and $\ell \in [L]$.

- $\mathcal{S}_U^\beta(r_{\text{pub}}; r_{\text{priv}})$ first runs the deterministic subprocedure $\mathcal{S}_{A_0}^\beta(r_{\text{pub}})$, which does the following. It first reruns $\mathcal{S}_V(r_{\text{pub}})$ to obtain $\mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}$. The algorithm then rearranges them into $\mathbf{A}_Q, \mathbf{A}_{\text{msg}}, \{\mathbf{A}_{t,i}\}_{i \in [m_t]}, \{\mathbf{A}_{\ell,i}\}_{\ell \in [L], i \in [m_\ell]}$ such that for all $\mathbf{s}_Q, \tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_n \in \mathbb{Z}_q^n$ and $s_{\text{msg}} \in \mathbb{Z}_q$,

$$\begin{aligned}
\mathbf{d}_0^\top \mathbf{A}_Q &= \mathbf{u}_Q^\top = (\mathbf{s}_Q^\top, \mathbf{0}_{1 \times (Z-n)}), \\
\mathbf{d}_0^\top \mathbf{A}_{\text{msg}} &= \mathbf{u}_{\text{msg}}^\top = (\mathbf{w}_{\text{out}}^\top \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}) \otimes \mathbf{s}_Q)^\top, \beta s_{\text{msg}}, \mathbf{0}_{1 \times (n-1+n^2K)}), \\
(\mathbf{A}_{t,1}^\top \mathbf{d}_0, \dots, \mathbf{A}_{t,m_t}^\top \mathbf{d}_0)^\top &= \mathbf{U}_t^\top = (\mathbf{T}^\top \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{T}) \otimes \mathbf{s}_Q)^\top, \mathbf{0}_{m_t \times (n+n^2K)}), \\
(\mathbf{A}_{\ell,1}^\top \mathbf{d}_0, \dots, \mathbf{A}_{\ell,m_\ell}^\top \mathbf{d}_0)^\top &= \mathbf{U}_\ell^\top \\
&= (\mathbf{W}_{\ell,0}^\top \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0}) \otimes \mathbf{s}_Q)^\top, (\mathbf{W}_{\ell,0}^\top + \mathbf{W}_{\ell,\times}^\top) \mathbf{S}, -(\tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{W}_{\ell,\times}) \otimes \mathbf{s}_Q)^\top),
\end{aligned}$$

where $\ell \in [L]$ and $\mathbf{d}_0^\top = (s_{\text{msg}}, \mathbf{s}_Q^\top, \tilde{\mathbf{s}}_1^\top, \dots, \tilde{\mathbf{s}}_n^\top)$ and $\mathbf{S} = (\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_n)$. This is possible since every entry on the right-hand side is linear in \mathbf{d}_0 . The algorithm \mathcal{S}_{A_0} outputs

$$\begin{aligned}
1^{n'} &= 1^{n(n+1)+1}, & \mathbf{A}_{-1,\mathbf{g},0,1} &= \mathbf{A}_Q, \\
\mathbf{A}_{0,\mathbf{o},0,m_t+1} &= \mathbf{A}_{\text{msg}}, & \mathbf{A}_{0,\mathbf{o},0,i} &= \mathbf{A}_{t,i} \quad (\text{for all } i \in [m_t]), \\
& & \mathbf{A}_{\ell,\mathbf{o},0,i} &= \mathbf{A}_{\ell,i} \quad (\text{for all } \ell \in [L], i \in [m_\ell]).
\end{aligned}$$

Completing \mathcal{S}_{A_0} , the algorithm \mathcal{S}_U samples $\mathbf{d}_0 \xleftarrow{\$} \mathbb{Z}_q^{n'}$ using r_{priv} and outputs

$$\mathbf{U}_{\text{id},\mathbf{s}}^\top = \begin{pmatrix} \mathbf{d}_0^\top \mathbf{A}_{\text{id},\mathbf{s},0,1} \\ \vdots \\ \mathbf{d}_0^\top \mathbf{A}_{\text{id},\mathbf{s},0,I_{\text{id},\mathbf{s}}} \end{pmatrix} \quad \text{for all } \text{id} \in \text{ID}, \mathbf{s} \in \{\mathbf{o}, \mathbf{g}\}.$$

By definition, \mathcal{S}^β is a restricted sampler (Definition 6). We have the following:

Claim 33 (¶). \mathcal{S}^β has pseudorandom structurally noisy inner products, i.e., $\text{IPFEsec}_{\text{pre}}^{\mathcal{S}^\beta}$ (Definition 5) holds for both $\beta \in \{0, 1\}$.

The very selective security of Construction 10 follows from Claim 33. Consider the following hybrids.

- H_0^β is $\text{Exp}_{\text{ABE}}^\beta$ for Construction 10. Note that s_{msg} in the challenge ciphertext is uniformly random over $\mathbb{Z}_q \setminus [-2 \cdot 2^{-\kappa} q, 2 \cdot 2^{-\kappa} q]$.
- In H_1^β , the scalar s_{msg} is changed to be uniformly random over \mathbb{Z}_q . Clearly, $H_0^\beta \approx_s H_1^\beta$. Moreover, H_1^β corresponds to the left distribution of $\text{IPFEsec}_{\text{post}}^{\mathcal{S}^\beta}$.
- In H_2^β , the vectors inside IPFE ciphertexts (components of the ABE challenge ciphertext) are replaced by random, which corresponds to the right distribution of $\text{IPFEsec}_{\text{post}}^{\mathcal{S}^\beta}$. By Claim 33 and the restricted- σ_{IPFE} -security of IPFE, we have $H_1^\beta \approx H_2^\beta$.

Lastly, $H_2^0 \equiv H_2^1$. By hybrid argument, we conclude $\text{Exp}_{\text{ABE}}^0 \approx \text{Exp}_{\text{ABE}}^1$, completing the proof. \square

Proof (Claim 33). Fix $\beta \in \{0, 1\}$ and consider the following hybrids.

- H_0 is the left distribution of $\text{IPFESec}_{\text{pre}}^{S^\beta}$. Recall that the public randomness is $r_{\text{pub}} = (r_{\mathcal{A}}, r_{\text{NLG.Setup}}, r_{\text{rq}}, r_{\text{NLG.GenF}})$. Adopting the notations of the ABE decryption algorithm, the noisy inner products are

$$\begin{aligned} \text{sh}_{Q,j}^\top &= \mathbf{s}_Q^\top \mathbf{Q}_j + \mathbf{e}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{Q,j}^\top, \\ \text{sh}_{\text{msg},j} &= (\mathbf{S}r_j)^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \beta s_{\text{msg}} - \mathbf{s}_Q^\top \mathbf{Q}_j \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), \\ \text{sh}_{t,j}^\top &= (\mathbf{S}r_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top - \mathbf{s}_Q^\top \mathbf{Q}_j \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\ \text{sh}_{j,\ell}^\top &= (\mathbf{S}r_j)^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top - \mathbf{s}_Q^\top \mathbf{Q}_j \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) \quad (\ell \in [L]), \end{aligned}$$

where $j \in [J]$ and the entries of \mathbf{e} 's are independent $\mathcal{D}_{Z, \sigma_{\text{IPFE}}}$.

- In H_1 , additional small noises are attached. For all $j \in [J]$, the inner products become

$$\begin{aligned} \text{sh}_{Q,j}^\top &= \mathbf{s}_Q^\top \mathbf{Q}_j + \tilde{\mathbf{e}}_{Q,j}^\top + \tilde{\mathbf{e}}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{Q,j}^\top, \\ \text{sh}_{\text{msg},j} &= (\mathbf{S}r_j)^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \beta s_{\text{msg}} - (\mathbf{s}_Q^\top \mathbf{Q}_j + \tilde{\mathbf{e}}_{Q,j}^\top) \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), \\ \text{sh}_{t,j}^\top &= (\mathbf{S}r_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top - (\mathbf{s}_Q^\top \mathbf{Q}_j + \tilde{\mathbf{e}}_{Q,j}^\top) \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\ \text{sh}_{j,\ell}^\top &= (\mathbf{S}r_j)^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top - (\mathbf{s}_Q^\top \mathbf{Q}_j + \tilde{\mathbf{e}}_{Q,j}^\top) \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}), \end{aligned}$$

where $\ell \in [L]$ and the entries of $\tilde{\mathbf{e}}$ are independent $\mathcal{D}_{Z, \sigma, \leq \sigma \sqrt{k}}$. In $\text{sh}_{Q,j}$'s, the $\tilde{\mathbf{e}}_{Q,j}$'s and $\tilde{\mathbf{e}}_{\text{sr},j}$'s are flooded by $\mathbf{e}_{Q,j}$'s and $\mathbf{e}_{\text{sr},j}$'s, respectively. In the other inner products, each $\tilde{\mathbf{e}}_{Q,j}^\top \tilde{\mathbf{G}}^{-1}(\dots)$ is flooded by the e or \mathbf{e} in that term. Since

$$\begin{aligned} (\text{for } \text{sh}_{Q,j}\text{'s}) \quad \sigma \sqrt{k} &\leq \frac{2^{-\kappa-6} \sigma_{\text{IPFE}}}{nK} \leq 2^{-\kappa-6} \sigma_{\text{IPFE}}, \\ (\text{for the rest}) \quad nK \cdot \sigma \sqrt{k} &\leq nK \cdot \frac{2^{-\kappa-6} \sigma_{\text{IPFE}}}{nK} = 2^{-\kappa-6} \sigma_{\text{IPFE}}, \end{aligned}$$

Lemma 2 applies and $H_0 \approx_s H_1$.

- In H_2 , the $\tilde{\mathbf{e}}$'s are no longer truncated, i.e., their entries follow $\mathcal{D}_{Z, \sigma}$. We have $H_1 \approx_s H_2$ by Lemma 1.
- In H_3 , the LWE samples $(\mathbf{s}_Q^\top \mathbf{Q}_j + \tilde{\mathbf{e}}_{Q,j}^\top)$ are replaced by random. For all $j \in [J]$, sample $\mathbf{q}_j \xleftarrow{\$} \mathbb{Z}_q^n$ and the inner products are

$$\begin{aligned} \text{sh}_{Q,j}^\top &= \mathbf{q}_j^\top + \tilde{\mathbf{e}}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{Q,j}^\top, \\ \text{sh}_{\text{msg},j} &= (\mathbf{S}r_j)^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \beta s_{\text{msg}} - \mathbf{q}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), \\ \text{sh}_{t,j}^\top &= (\mathbf{S}r_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top - \mathbf{q}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\ \text{sh}_{j,\ell}^\top &= (\mathbf{S}r_j)^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top - \mathbf{q}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) \quad (\ell \in [L]). \end{aligned}$$

By $\text{LWE}_{n,nJ,q,\sigma}$, we have $H_2 \approx H_3$. In this step, we rely on \mathbf{Q} 's sampling being "straight-forward" (for reduction to sample r_{rq} conditioned on and given \mathbf{Q} 's).

- In H_4 , we perform a change of variable. For all $j \in [J]$, sample $\tilde{\mathbf{q}}_j \xleftarrow{\$} \mathbb{Z}_q^n$ and set $\mathbf{q}_j^\top = \tilde{\mathbf{q}}_j^\top - \tilde{\mathbf{e}}_{Q,j}^\top \tilde{\mathbf{G}}$, then the inner products are

$$\begin{aligned}
\text{sh}_{Q,j}^\top &= \tilde{\mathbf{q}}_j^\top + \mathbf{e}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{Q,j}^\top, \\
\text{sh}_{\text{msg},j} &= (\mathbf{S}\mathbf{r}_j)^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \beta s_{\text{msg}} - (\tilde{\mathbf{q}}_j^\top - \tilde{\mathbf{e}}_{Q,j}^\top \tilde{\mathbf{G}}) \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}) \\
&= (\mathbf{S}\mathbf{r}_j + \tilde{\mathbf{e}}_{Q,j})^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \beta s_{\text{msg}} - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), \\
\text{sh}_{t,j}^\top &= (\mathbf{S}\mathbf{r}_j)^\top \mathbf{T} + \mathbf{e}_{t,j}^\top - (\tilde{\mathbf{q}}_j^\top - \tilde{\mathbf{e}}_{Q,j}^\top \tilde{\mathbf{G}}) \tilde{\mathbf{G}}^{-1}(\mathbf{T}) \\
&= (\mathbf{S}\mathbf{r}_j + \tilde{\mathbf{e}}_{Q,j})^\top \mathbf{T} + \mathbf{e}_{t,j}^\top - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\
\text{sh}_{j,\ell}^\top &= (\mathbf{S}\mathbf{r}_j)^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top - (\tilde{\mathbf{q}}_j^\top - \tilde{\mathbf{e}}_{Q,j}^\top \tilde{\mathbf{G}}) \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) \\
&= (\mathbf{S}\mathbf{r}_j + \tilde{\mathbf{e}}_{Q,j})^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}),
\end{aligned}$$

where $\ell \in [L]$. Clearly, $H_3 \equiv H_4$.

- In H_5 , the LWE samples $(\mathbf{S}\mathbf{r}_j + \tilde{\mathbf{e}}_{Q,j})$ are replaced by random. For all $j \in [J]$, sample $\mathbf{s}_j \xleftarrow{\$} \mathbb{Z}_q^n$ and the inner products are

$$\begin{aligned}
\text{sh}_{Q,j}^\top &= \tilde{\mathbf{q}}_j^\top + \mathbf{e}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{Q,j}^\top, \\
\text{sh}_{\text{msg},j} &= \mathbf{s}_j^\top \mathbf{w}_{\text{out}} + e_{\text{out},j} + \beta s_{\text{msg}} - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), \\
\text{sh}_{t,j}^\top &= \mathbf{s}_j^\top \mathbf{T} + \mathbf{e}_{t,j}^\top - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\
\text{sh}_{j,\ell}^\top &= \mathbf{s}_j^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \mathbf{e}_{j,\ell}^\top - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) \quad (\ell \in [L]).
\end{aligned}$$

By n hybrids of $\text{LWE}_{n,J,q,\sigma}$ over the rows of \mathbf{S} with public matrix $(\mathbf{r}_1, \dots, \mathbf{r}_J)$, we have $H_4 \approx H_5$.

- In H_6 , additional noises generated by GenNoise are attached to the garblings. For all $j \in [J]$, run

$$(\bar{e}_{\text{out},j}, \{\bar{\mathbf{e}}_{j,\ell}\}_{\ell \in [L]}, \bar{\mathbf{e}}_{t,j}) \xleftarrow{\$} \text{GenNoise}(\text{param}', \text{pp}, f, \mathbf{w}_{\text{out}}, \{\mathbf{W}_{\ell,0}, \mathbf{W}_{\ell,\times}\}_{\ell \in [L]}, \mathbf{T}, \mathbf{R}, \mathbf{x}_j),$$

and set the inner products to

$$\begin{aligned}
\text{sh}_{Q,j}^\top &= \tilde{\mathbf{q}}_j^\top + \mathbf{e}_{\text{sr},j}^\top \tilde{\mathbf{G}} + \mathbf{e}_{Q,j}^\top, \\
\text{sh}_{\text{msg},j} &= \mathbf{s}_j^\top \mathbf{w}_{\text{out}} + \bar{e}_{\text{out},j} + e_{\text{out},j} + \beta s_{\text{msg}} - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{w}_{\text{out}}), \\
\text{sh}_{t,j}^\top &= \mathbf{s}_j^\top \mathbf{T} + \bar{\mathbf{e}}_{t,j}^\top + \mathbf{e}_{t,j}^\top - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{T}), \\
\text{sh}_{j,\ell}^\top &= \mathbf{s}_j^\top (\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) + \bar{\mathbf{e}}_{j,\ell}^\top + \mathbf{e}_{j,\ell}^\top - \tilde{\mathbf{q}}_j^\top \tilde{\mathbf{G}}^{-1}(\mathbf{W}_{\ell,0} + \mathbf{x}_j[\ell] \mathbf{W}_{\ell,\times}) \quad (\ell \in [L]).
\end{aligned}$$

The \bar{e} , $\bar{\mathbf{e}}$'s introduced into the inner products are bounded by $B_{\text{NLG}} \leq 2^{-\kappa-6} \sigma_{\text{IPFE}}$, so they are flooded (Lemma 2) by e , \mathbf{e} 's. We have $H_5 \approx_s H_6$.

- In H_7 , all the inner products are replaced by random, i.e., the right distribution of $\text{IPFEsec}_{\text{pre}}^{S^\beta}$. Note that in H_6 , non- $\text{sh}_{Q,j}$'s are well-randomized garblings (by $\mathbf{s}, \bar{e}, \bar{\mathbf{e}}$'s) plus independent noises and values derived from $\text{sh}_{Q,j}$. Therefore, by J hybrids of applying GenNoise-security of NLG, the non- $\text{sh}_{Q,j}$'s can be replaced by random, after which $\tilde{\mathbf{q}}_j$'s ensure $\text{sh}_{Q,j}$'s are independent random. Therefore, $H_6 \approx H_7$.

By hybrid argument, $H_0 \approx H_7$, i.e., $\text{IPFEsec}_{\text{pre}}^{S^\beta}$ holds. \square

Summary. Combining Corollary 14, Construction 9, and Construction 10, we obtain a CP-ABE scheme for Boolean circuits with succinct ciphertexts.

Corollary 34. *Under LWE (Assumption 1) and evasive learning with structured errors (Assumption 3), there exists a CP-ABE scheme for Boolean circuits with*

$$|\text{mpk}| = \text{poly}(\lambda, d), \quad |\text{sk}_{\mathbf{x}}| = (L + 1) \text{poly}(\lambda, d), \quad |\text{ct}_C| = (L + 1) \text{poly}(\lambda, d),$$

where $L = |\mathbf{x}|$ is the attribute length and d is the maximum depth.

Acknowledgments. The authors were supported by NSF grants CNS-1936825 (CAREER), CNS-2026774, a JP Morgan AI Research Award, a Cisco Research Award, and a Simons Collaboration on the Theory of Algorithmic Fairness. The views expressed are those of the authors and do not reflect the official policy or position of the funding agencies. The authors thank the anonymous reviewers of Eurocrypt 2024 for their valuable comments.

References

- [ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Heidelberg, May / June 2010.
- [ABDP15] Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 733–751. Springer, Heidelberg, March / April 2015.
- [ACF⁺18] Michel Abdalla, Dario Catalano, Dario Fiore, Romain Gay, and Bogdan Ursu. Multi-input functional encryption for inner products: Function-hiding realizations and constructions without pairings. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 597–627. Springer, Heidelberg, August 2018.
- [Agr19] Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 191–225. Springer, Heidelberg, May 2019.
- [AIK11] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. How to garble arithmetic circuits. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 120–129. IEEE Computer Society Press, October 2011.
- [ALMT20] Shweta Agrawal, Benoît Libert, Monosij Maitra, and Radu Titiiu. Adaptive simulation security for inner product functional encryption. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 34–64. Springer, Heidelberg, May 2020.

- [ALS16] Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 333–362. Springer, Heidelberg, August 2016.
- [AMY19] Shweta Agrawal, Monosij Maitra, and Shota Yamada. Attribute based encryption (and more) for nondeterministic finite automata from LWE. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 765–797. Springer, Heidelberg, August 2019.
- [AP20] Shweta Agrawal and Alice Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 110–140. Springer, Heidelberg, May 2020.
- [AS17] Shweta Agrawal and Ishaan Preet Singh. Reusable garbled deterministic finite automata from learning with errors. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *ICALP 2017*, volume 80 of *LIPICs*, pages 36:1–36:13. Schloss Dagstuhl, July 2017.
- [AWY20] Shweta Agrawal, Daniel Wichs, and Shota Yamada. Optimal broadcast encryption from LWE and pairings in the standard model. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 149–178. Springer, Heidelberg, November 2020.
- [AY20] Shweta Agrawal and Shota Yamada. Optimal broadcast encryption from pairings and LWE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 13–43. Springer, Heidelberg, May 2020.
- [BDHM92] Gerhard Buntrock, Carsten Damm, Ulrich Hertrampf, and Christoph Meinel. Structure and importance of logspace-MOD class. *Mathematical Systems Theory*, 25(3):223–237, 1992.
- [Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Heidelberg, May 2014.
- [CHL⁺15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12. Springer, Heidelberg, April 2015.

- [CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 577–607. Springer, Heidelberg, August 2018.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- [GKP⁺13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2013.
- [GKW16] Rishab Goyal, Venkata Koppula, and Brent Waters. Semi-adaptive security and bundling functionalities made generic and easy. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 361–388. Springer, Heidelberg, October / November 2016.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013.
- [GW20] Junqing Gong and Hoeteck Wee. Adaptively secure ABE for DFA from k -Lin and more. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 278–308. Springer, Heidelberg, May 2020.
- [GWW19] Junqing Gong, Brent Waters, and Hoeteck Wee. ABE for DFA from k -Lin. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 732–764. Springer, Heidelberg, August 2019.
- [HJL21] Samuel B. Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying iO. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 673–700, Virtual Event, August 2021. Springer, Heidelberg.
- [HLL23a] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In *64th FOCS*, pages 415–434. IEEE Computer Society Press, November 2023.

- [HLL23b] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices: Garbled circuits of optimal size, laconic functional evaluation, and more. Cryptology ePrint Archive, Report 2023/1716, 2023. <https://eprint.iacr.org/2023/1716>.
- [HLL24] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. A general framework for lattice-based ABE using evasive inner-product functional encryption. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 433–464. Springer, Cham, May 2024.
- [IW14] Yuval Ishai and Hoeteck Wee. Partial garbling schemes and their applications. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP 2014, Part I*, volume 8572 of *LNCS*, pages 650–662. Springer, Heidelberg, July 2014.
- [JLL23] Aayush Jain, Huijia Lin, and Ji Luo. On the optimal succinctness and efficiency of functional encryption and attribute-based encryption. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 479–510. Springer, Heidelberg, April 2023.
- [JLLS23] Aayush Jain, Huijia Lin, Paul Lou, and Amit Sahai. Polynomial-time cryptanalysis of the subspace flooding assumption for post-quantum $i\mathcal{O}$. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 205–235. Springer, Heidelberg, April 2023.
- [KW93] Mauricio Karchmer and Avi Wigderson. On span programs. In *SCT 1993*, pages 102–111. IEEE, May 1993.
- [LL20a] Huijia Lin and Ji Luo. Compact adaptively secure ABE from k -Lin: Beyond NC^1 and towards NL. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 247–277. Springer, Heidelberg, May 2020.
- [LL20b] Huijia Lin and Ji Luo. Succinct and adaptively secure ABE for ABP from k -lin. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 437–466. Springer, Heidelberg, December 2020.
- [LLL22] Hanjun Li, Huijia Lin, and Ji Luo. ABE for circuits with constant-size secret keys and adaptive security. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 680–710. Springer, Heidelberg, November 2022.
- [MP11] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. Cryptology ePrint Archive, Report 2011/501, 2011. <https://eprint.iacr.org/2011/501>.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.

- [Mul87] Ketan Mulmuley. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. *Combinatorica*, 7(1):101–104, 1987.
- [QWW18] Willy Quach, Hoeteck Wee, and Daniel Wichs. Laconic function evaluation and applications. In Mikkel Thorup, editor, *59th FOCS*, pages 859–870. IEEE Computer Society Press, October 2018.
- [QWW21] Willy Quach, Brent Waters, and Daniel Wichs. Targeted lossy functions and applications. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 424–453, Virtual Event, August 2021. Springer, Heidelberg.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- [RW15] Yannis Rouselakis and Brent Waters. Efficient statically-secure large-universe multi-authority attribute-based encryption. In Rainer Böhme and Tatsuaki Okamoto, editors, *FC 2015*, volume 8975 of *LNCS*, pages 315–332. Springer, Heidelberg, January 2015.
- [SW05] Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
- [TCH12] Terence Tao, Ernest Croot, and Harald Helfgott. Deterministic methods to find primes. *Mathematics of Computation*, 81(278):1233–1246, 2012.
- [Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 535–559. Springer, Heidelberg, August 2022.
- [VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 195–221. Springer, Heidelberg, December 2022.
- [Wat12] Brent Waters. Functional encryption for regular languages. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 218–235. Springer, Heidelberg, August 2012.
- [Wee17] Hoeteck Wee. Attribute-hiding predicate encryption in bilinear groups, revisited. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 206–233. Springer, Heidelberg, November 2017.
- [Wee21] Hoeteck Wee. ABE for DFA from LWE against bounded collusions, revisited. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 288–309. Springer, Heidelberg, November 2021.

- [Wee22] Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Heidelberg, May / June 2022.
- [WWW22] Brent Waters, Hoeteck Wee, and David J. Wu. Multi-authority ABE from lattices without random oracles. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 651–679. Springer, Heidelberg, November 2022.

A Noisy Linear Garbling for Boolean Circuits

In this section, we present an alternative construction of noisy linear garbling. It is closer to the Yao's garbling scheme and supports bounded depth boolean circuits.

Construction 11 (noisy linear garbling for Boolean circuits). Let

$$\begin{aligned} \text{Params} &= \{1^d \mid d \in \mathbb{N}\}, \\ F_{1^d} &= \{f_C \mid C \text{ is a circuit (using only NAND gates) of depth no more than } d\}, \\ f_C(\mathbf{x}) &= \begin{cases} \perp, & \text{if } \mathbf{x} \notin \{0, 1\}^L; \\ -C(\mathbf{x}), & \text{if } \mathbf{x} \in \{0, 1\}^L; \end{cases} \quad \text{for } C : \{0, 1\}^L \rightarrow \{0, 1\} \text{ and } \mathbf{x} \in \mathbb{Z}^L. \end{aligned}$$

The scheme works as follows.

- Setup(1^d) picks and outputs suitable pp = $(q, 1^N)$.
- GenF(pp, C) first parses C into L input gates $1, \dots, L$ and $(|C| - L)$ NAND gates $L + 1, \dots, |C|$, sorted topologically. Let $\text{gin}[i, \gamma] < i$ be the index of the γ^{th} input to gate i , where $\gamma \in \{1, 2\}$ and $L < i \leq |C|$. The algorithm samples the shrunken labels and sets the label functions

$$\mathbf{W}_{i,0}, \mathbf{W}_{i,1} \stackrel{\$}{\leftarrow} \{0, 1\}^{N \times N} \quad \text{for all } i \in [|C|], \quad \mathbf{W}_{\ell, \times} = \mathbf{W}_{\ell,1} - \mathbf{W}_{\ell,0} \quad \text{for all } \ell \in [L].$$

For each NAND gate $i > L$, it samples the expanded labels

$$\widetilde{\mathbf{W}}_{\text{gin}[i,\gamma], x'}^{(i,\gamma,x'')} \stackrel{\$}{\leftarrow} \{0, 1\}^{N \times N} \quad \text{for all } \gamma \in \{1, 2\}, x', x'' \in \{0, 1\},$$

and prepares four garbled table entries

$$\mathbf{T}_{i,x_1,x_2} = \mathbf{W}_{i,\neg(x_1 \wedge x_2)} + \widetilde{\mathbf{W}}_{\text{gin}[i,1], x_1}^{(i,1,x_2)} + \widetilde{\mathbf{W}}_{\text{gin}[i,2], x_2}^{(i,2,x_1)} \quad \text{for all } x_1, x_2 \in \{0, 1\}.$$

The subscript of $\widetilde{\mathbf{W}}$ indicates the gate and its purported value. Its superscript indicates why the block is added — because the gate is the γ^{th} input to gate i and the other input to gate i evaluates to a particular value. The table entry \mathbf{T}_{i,x_1,x_2} enables conversion from the expanded labels of gates $\text{gin}[i, 1], \text{gin}[i, 2]$ to the shrunken label of gate i if the two inputs to gate i are x_1, x_2 , respectively. For the output gate, the algorithm samples the expanded labels and sets

$$\widetilde{\mathbf{W}}_{|C|,0}, \widetilde{\mathbf{W}}_{|C|,1} \stackrel{\$}{\leftarrow} \{0, 1\}^N, \quad \mathbf{w}_{\text{out}} = \widetilde{\mathbf{W}}_{|C|,0}.$$

It prepares garbled table entries for converting shrunken labels to expanded labels. For each tuple (i', i, γ) with $\text{gin}[i, \gamma] = i'$, the algorithm samples and sets

$$\mathbf{R}_{i',x'}^{(i,\gamma,x'')} \stackrel{\$}{\leftarrow} \{0, 1\}^{N \times N}, \quad \mathbf{T}_{i',x'}^{(i,\gamma,x'')} = \mathbf{W}_{i',x'} \mathbf{R}_{i',x'}^{(i,\gamma,x'')} + \widetilde{\mathbf{W}}_{i',x'}^{(i,\gamma,x'')} \quad \text{for all } x', x'' \in \{0, 1\}.$$

It samples and sets

$$\mathbf{R}_{|C|,x} \stackrel{\$}{\leftarrow} \{0, 1\}^N, \quad \mathbf{T}_{|C|,x} = \mathbf{W}_{|C|,x} \mathbf{R}_{|C|,x} + \widetilde{\mathbf{W}}_{|C|,x} \quad \text{for all } x \in \{0, 1\}.$$

The algorithm collects \mathbf{T}, \mathbf{R} and outputs the components in the required format, with $n_C = n$.

- Eval(pp, C , \mathbf{R} , \mathbf{x} , $\{\mathbf{w}_\ell^\top\}_{\ell \in [L]}$, \mathbf{t}^\top) parses C as in GenF and evaluates $C(\mathbf{x})$ for the wire values $\{x_i\}_{i \in [|C|]}$ by

$$x_i \leftarrow \begin{cases} \mathbf{x}[\ell], & \text{if gate } i \text{ is an input gate } (i = \ell \in [L]); \\ \neg(x_{\text{gin}[i,1]} \wedge x_{\text{gin}[i,2]}), & \text{if gate } i \text{ is NAND.} \end{cases}$$

It computes the label of each non-input gate in increasing order by its index i .

- The algorithm decrypts for the expanded labels using

$$(\widetilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1,x_{\text{gin}[i,2]})})^\top \leftarrow (\mathbf{t}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1,x_{\text{gin}[i,2]})})^\top - \mathbf{w}_{\text{gin}[i,1]}^\top \mathbf{R}_{\text{gin}[i,1],x_{\text{gin}[i,1]}}^{(i,1,x_{\text{gin}[i,2]})}$$

and similarly for $\widetilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2,x_{\text{gin}[i,1]})}$.

- The algorithm recovers the shrunken label using

$$\mathbf{w}_i \leftarrow \mathbf{t}_{i,x_{\text{gin}[i,1]},x_{\text{gin}[i,2]}} - \widetilde{\mathbf{w}}_{\text{gin}[i,1]}^{(i,1,x_{\text{gin}[i,2]})} - \widetilde{\mathbf{w}}_{\text{gin}[i,2]}^{(i,2,x_{\text{gin}[i,1]})}.$$

Lastly, the algorithm computes and outputs the secret

$$w_{\text{out}} \leftarrow \mathbf{t}_{|C|,x_{|C|}}^\top - \mathbf{w}_{|C|}^\top \mathbf{R}_{|C|,x_{|C|}}.$$