Circuit ABE with poly(depth, λ)-sized Ciphertexts and Keys from Lattices

Hoeteck Wee

NTT Research

Abstract. We present new lattice-based attribute-based encryption (ABE) and laconic function evaluation (LFE) schemes for circuits with *sublinear* ciphertext overhead. For depth d circuits over ℓ-bit inputs, we obtain

- an ABE with ciphertext and secret key size $O(1)$;
- a LFE with ciphertext size $\ell + O(1)$ and digest size $O(1)$;
- an ABE with public key and ciphertext size $O(\ell^{2/3})$ and secret key size $O(1)$,

where $O(\cdot)$ hides $\mathsf{poly}(d, \lambda)$ factors. The first two results achieve almost optimal ciphertext and secret key / digest sizes, up to the $\mathsf{poly}(d)$ dependencies. The security of our schemes relies on ℓ-succinct LWE, a falsifiable assumption which is implied by evasive LWE. At the core of our results is a new technique for compressing LWE samples $\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$ as well as the matrix $\mathbf{A}$.

1 Introduction

Let $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$ be a matrix where $m = O(n \log q)$. Given $\mathbf{A}$ and a circuit $f : \{0,1\}^\ell \to \{0,1\}$, we can derive a matrix $\mathbf{A}_f \in \mathbb{Z}_q^{n \times m}$ such that for any $\mathbf{x} \in \{0,1\}^\ell$, we can compute a low-norm matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ satisfying

$$(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(\mathbf{x})\mathbf{G} \tag{1}$$

This remarkable relation, first discovered in the context of attribute-based and fully homomorphic encryption [13,31], enabled a spectacular array of cryptographic advances in the past decade that fall under the broad theme of "encrypted computation": fully homomorphic signatures [35], constrained pseudorandom functions [17], predicate encryption [34], laconic function evaluation [43,24], correlation-intractable hashing and NIZK [42,20], and many more. Moreover, these schemes support expressive computation on circuits and achieve security under the standard LWE assumption.

1.1 Our Results

We present a new technique for compressing LWE samples for the matrix $\mathbf{A} - \mathbf{x} \otimes \mathbf{G}$ in (1) as well as the matrix $\mathbf{A}$ itself. This yields new constructions of attribute-based encryption (ABE) and laconic function evaluation (LFE) for circuits with sublinear ciphertext overhead; security relies on ℓ-succinct LWE, a new falsifiable variant of LWE put forth in this work, which in turn follows from evasive LWE [49,45].

Attribute-based encryption. In attribute-based encryption (ABE), ciphertexts ct are associated with an attribute $x \in \{0,1\}^\ell$ and a message μ and keys sk with a predicate f, and decryption returns μ when x satisfies f (i.e, $f(x) = 0$). We require security against unbounded collusions, so that an adversary that sees a ciphertext along with secret keys for an arbitrary number of predicates learns nothing about μ as long as x satisfies none of these predicates.

Prior work. In 2014, Boneh, Gentry, Gorbunov, Halevi, Nikolaenko, Segev, Vaikuntanathan, and Vinayagamurthy [13], henceforth BGGHNSVV, constructed an ABE scheme for circuits with small keys from LWE, improving on [33]. For depth d, size s circuits over ℓ-bit inputs where ℓ and d are fixed at set-up, the scheme achieves

$$|\mathsf{mpk}| = O(\ell), \quad |\mathsf{ct}| = O(\ell), \quad |\mathsf{sk}| = O(1)$$

where $O(\cdot)$ hides $\mathsf{poly}(d, \lambda)$ factors. Roughly speaking, $\mathsf{mpk}, \mathsf{ct}, \mathsf{sk}$ correspond to $\mathbf{A}, \mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}), \mathbf{A}_f$ respectively, where $\mathbf{x} \in \{0,1\}^\ell$ is the attribute. In spite of the substantial progress and improvements in lattice-based ABE since BGGHNSVV [18,6,4,5,19,49,40,38,26], all known schemes inherit the limitation that $|\mathsf{ct}| + |\mathsf{sk}| = \Omega(\ell)$ as well as $|\mathsf{mpk}| + |\mathsf{sk}| = \Omega(\ell)$, cf. Fig 1. In particular, in the setting where $|\mathsf{sk}| = O(1)$, the state of the art requires both $|\mathsf{ct}|, |\mathsf{mpk}| = \Omega(\ell)$.

This work. For any $1/3 \le \alpha \le 1$, we construct an ABE scheme with parameters

$$|\mathsf{mpk}| = O(\ell^{2\alpha}), \quad |\mathsf{ct}| = O(\ell^{1-\alpha}), \quad |\mathsf{sk}| = O(1);$$

where $O(\cdot)$ hides $\mathsf{poly}(d, \lambda)$ factors. We obtain as special cases corresponding to $\alpha = 1$ and $\alpha = 1/3$:

- the first lattice-based ABE to simultaneously achieve $O(1)$-sized ciphertexts and secret keys —almost optimal, up to $\mathsf{poly}(d)$ factors in $O(\cdot)$— answering a natural question left open in BGGHNSVV[1];
- an ABE with $|\mathsf{mpk}| = |\mathsf{ct}| = O(\ell^{2/3}), |\mathsf{sk}| = O(1)$, simultaneously breaking the $\Omega(\ell)$ barrier for both $|\mathsf{ct}| + |\mathsf{sk}|$ and $|\mathsf{mpk}| + |\mathsf{sk}|$.

Laconic function evaluation. In laconic function evaluation (LFE), a server publishes a short digest dig to a function f. Anyone can use dig to efficiently encrypt an input $x \in \{0,1\}^\ell$. Given f, the ciphertext ct can then be decrypted to recover $f(x)$, but hides everything else about x. Building on the BGGHNSVV ABE, Quach, Wee and Wichs [43], QWW for short, constructed an LFE scheme for circuits with parameters

$$|\mathsf{crs}| = O(\ell), \quad |\mathsf{ct}| = O(\ell), \quad |\mathsf{dig}| = O(1),$$
$$\text{encryption time } O(\ell)$$

As before, ℓ and d are fixed at set-up, and $O(\cdot)$ hides $\mathsf{poly}(d, \lambda)$ factors. Roughly speaking, $\mathsf{mpk}, \mathsf{ct}, \mathsf{dig}$ correspond to $\mathbf{A}, (\hat{\mathbf{x}}, \mathbf{s}(\mathbf{A} - \hat{\mathbf{x}} \otimes \mathbf{G})), \mathbf{A}_f$ respectively, where $\hat{\mathbf{x}}$ is a FHE encryption of the input $x \in \{0,1\}^\ell$.

This work. For any $1/3 \le \alpha \le 1$, we construct a LFE scheme with parameters

$$|\mathsf{crs}| = O(\ell^{2\alpha}), \quad |\mathsf{ct}| = \ell + O(\ell^{1-\alpha}), \quad |\mathsf{dig}| = O(1),$$
$$\text{encryption time } O(\ell)$$

That is, we reduce $|\mathsf{ct}|$ from $\ell \cdot \mathsf{poly}(d, \lambda)$ in QWW to $\ell + \ell^{1-\alpha} \cdot \mathsf{poly}(d, \lambda)$. We obtain as special cases corresponding to $\alpha = 1$ and $\alpha = 1/3$:

- the first lattice-based LFE to simultaneously achieve $\ell + O(1)$-sized ciphertext and $O(1)$-sized digest —almost optimal, up to $\mathsf{poly}(d)$ factors in $O(\cdot)$;
- a LFE with $|\mathsf{crs}| = O(\ell^{2/3}), |\mathsf{ct}| = \ell + O(\ell^{2/3})$, simultaneously achieving sublinear crs and ciphertext overhead; prior to this work, even achieving $|\mathsf{crs}| + |\mathsf{dig}| = o(\ell)$ was open.

We refer to Fig 2 for additional comparison with prior works for LFE.

The ℓ-succinct LWE assumption. The ℓ-succinct LWE assumption is a strengthening of the standard LWE assumption where the distinguisher additionally receives short Gaussian pre-images of size $O(\ell^2)$ from a fixed distribution. As with prior works on encrypted computation from LWE for circuits, we require hardness of ℓ-succinct LWE with a sub-exponential modulus-to-noise ratio. In our schemes, the LWE parameters depend on d, λ, and the parameter ℓ corresponds to the input length for the circuit. The gap between ℓ-succinct LWE and evasive LWE is analogous to that of q-type assumptions and the generic group model (GGM) in pairing-based cryptography; in both cases, the former is falsifiable and thus more desirable from both a theoretical and cryptanalytic stand-point. We defer an informal statement of ℓ-succinct LWE and additional discussion and justification to Sections 1.3 and 1.4.

We regard the introduction and use of falsifiable "q-type" LWE assumptions for advanced encryption primitives, where the assumption can in turn be justified using evasive LWE, as an additional conceptual contribution of this work. We reiterate that none of our results was known even from the stronger and non-falsifiable evasive LWE assumption.

[1] BGGHNSVV constructed a *second* ABE for circuits with $|\mathsf{ct}| = O(1)$ and $|\mathsf{sk}| = O(s)$, assuming multi-linear maps.

Reference	$\|\mathsf{mpk}\|$	$\|\mathsf{ct}\|$	$\|\mathsf{sk}\|$	Assumption
GVW13 [33]	$\ell \cdot \mathsf{poly}(d)$	$\ell \cdot \mathsf{poly}(d)$	$s \cdot \mathsf{poly}(d)$	LWE ✓
BGGHNSVV [13]	$\ell \cdot \mathsf{poly}(d)$	$\ell \cdot \mathsf{poly}(d)$	$\mathsf{poly}(d)$	LWE ✓
BV16 [18]	$\mathsf{poly}(d)$	$\ell \cdot \mathsf{poly}(d)$	$\ell + \mathsf{poly}(d)$	LWE ✓
CW23 [26,40]	$\ell \cdot \mathsf{poly}(d)$	$\ell \cdot \mathsf{poly}(d)$	1	LWE ✓
W22 [49,19]	$\ell \cdot \mathsf{poly}(d)$	$\mathsf{poly}(d)$	$\ell \cdot \mathsf{poly}(d)$	evasive LWE + tensor LWE
HLL23 [38]	ℓ	ℓ	1	evasive + circular LWE
this work	$\ell^2 \cdot \mathsf{poly}(d)$	$\mathsf{poly}(d)$	$\mathsf{poly}(d)$	ℓ-succinct LWE ✓
	$\ell^{2/3} \cdot \mathsf{poly}(d)$	$\ell^{2/3} \cdot \mathsf{poly}(d)$	$\mathsf{poly}(d)$	ℓ-succinct LWE ✓

Fig. 1. Comparison with prior lattice-based ABE for circuits of size s and depth d. The quantities $|\mathsf{ct}|, |\mathsf{sk}|$ refer to the cryptographic overhead beyond transmitting x and f in the clear, ignoring $\mathsf{poly}(\lambda)$ factors. When restricted to NC^1, the $\mathsf{poly}(d)$ factors can be omitted. A ✓ indicates a falsifiable assumption. Note that W22 is a ciphertext-policy scheme where the predicate f is associated with ct.

Reference	$\|\mathsf{crs}\|$	$\|\mathsf{ct}\|$	$\|\mathsf{dig}\|$	Assumption
QWW18 [43]	$\ell \cdot \mathsf{poly}(d, \lambda)$	$\ell \cdot \mathsf{poly}(d, \lambda)$	$\mathsf{poly}(d, \lambda)$	LWE
HLL23 [38]	$\ell \cdot \mathsf{poly}(\lambda)$	$\ell \cdot \mathsf{poly}(\lambda)$	$\mathsf{poly}(\lambda)$	circular LWE
this work	$\ell^2 \cdot \mathsf{poly}(d, \lambda)$	$\ell + \mathsf{poly}(d, \lambda)$	$\mathsf{poly}(d, \lambda)$	ℓ-succinct LWE
	$\ell^{2/3} \cdot \mathsf{poly}(d, \lambda)$	$\ell + \ell^{2/3} \cdot \mathsf{poly}(d, \lambda)$	$\mathsf{poly}(d, \lambda)$	ℓ-succinct LWE

Fig. 2. Comparison with prior LFE for circuits of size s and depth d.

1.2 High-level Overview

In the overview, we focus on our new ABE scheme with $O(1)$-sized ciphertexts and keys. Our approach is inspired by recent advances in lattice-based succinct arguments and functional commitments in [8,51]. We begin our overview with the results in the latter, using the notion of homomorphic instead of functional commitments. We continue to use $O(\cdot)$ to hide $\mathsf{poly}(d, \lambda)$ factors.

Homomomorphic commitments. Homomorphic commitments (HC) enable computing on a commitment com to $x \in \{0,1\}^\ell$ to derive a commitment com_f to $f(x)$; moreover, given the opening to x, we can also derive an opening to $f(x)$. In 2015, Gorbunov, Vaikuntanathan and Wichs [35] (GVW) constructed homomorphic commitments for circuits with $|\mathsf{com}| = O(\ell), |\mathsf{com}_f| = O(1)$, whose security relies on SIS. The GVW construction builds on the BGGHNSVV ABE, where $\mathsf{com}, \mathsf{com}_f$ correspond roughly to the ABE ciphertexts and keys respectively.[2]

A recent work of Wee and Wu (WW) [51], building on [8], improves on the GVW construction to achieve $|\mathsf{com}| = |\mathsf{com}_f| = O(1)$. The key innovation in WW is to compress the GVW commitment down to $O(1)$ bits using a trapdoor basis; homomorphic computation first decompresses –or, expands– the compressed commitment to recover a GVW commitment, and then proceeds as before in GVW. Security relies on $\mathsf{BASIS}_{\mathsf{struct}}$, a non-standard and falsifiable variant of SIS introduced in WW, which asserts that SIS is hard even given the trapdoor basis.

Our approach. Our high-level approach is to "lift" the WW homomorphic commitment into an ABE with $O(1)$-sized ciphertexts and keys *à la* GVW; in particular, we show how to compress BGGHNSVV ABE ciphertexts —i.e., LWE samples $\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$— *à la* WW. Our key technical contribution is an error-friendly variant of WW compression dualWW where decompression entails multiplication by *low-norm* matrices. In a nutshell,

[2] More precisely, $\mathsf{com}, \mathsf{com}_f$ correspond to $\mathbf{A}, \mathbf{A}_f$ in (1).

– The WW compressed commitment is derived from a linear combination of sub-matrices of the trapdoor basis, and decompression entails *left*-multiplication by *random* matrices $\mathbf{V}_1,\ldots,\mathbf{V}_\ell \leftarrow \mathbb{Z}_q^{n\times n}$ in the public parameters[3]; the latter is incompatible with ABE ciphertexts due to the blow-up in the error term.
– Our compressed ciphertext in dualWW is derived from a linear combination of $\mathbf{V}_1,\ldots,\mathbf{V}_\ell$, multiplied by an LWE secret on the *left*. Decompression entails *right*-multiplication by *low-norm* sub-matrices of the trapdoor basis. In fact, our compressed ciphertext for attribute $(x_1,\ldots,x_\ell)\in\{0,1\}^\ell$ is quite simply:

$$\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + \sum x_i \mathbf{V}_i \mathbf{G}] + \mathbf{e} \in \mathbb{Z}_q^{2m}$$

where $\mathbf{B},\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{V}_i \leftarrow \mathbb{Z}_q^{n\times n}$ are specified in the public key and $\mathbf{G}\in\mathbb{Z}_q^{n\times m}$ is the gadget matrix.

Applying dualWW compression to the BGGHNSVV ABE, we obtain an ABE for circuits with $\mathsf{poly}(\lambda,d)$-sized ciphertexts and keys. Security relies on the LWE analogue of $\mathsf{BASIS}_{\mathsf{struct}}$, which we refer to as $\mathsf{BALWE}_{\mathsf{struct}}$. We can also apply dualWW compression to the QWW LFE for circuits to reduce the ciphertext size from $O(\ell)$ to $\ell + O(1)$, while preserving digest size $O(1)$ and encryption time $O(\ell)$. Security of the ensuing LFE also relies on $\mathsf{BALWE}_{\mathsf{struct}}$.

Additional improvements. At this point, we inherit two limitations of the WW scheme. The first is a large mpk of size $O(\ell^2)$. To mitigate this issue, we show how to reduce the mpk size to $O(\ell^{2\alpha}+\ell^{1-\alpha})$, at the cost of increasing the ciphertext size to $O(\ell^{1-\alpha})$, for any $0\le\alpha\le 1$. The basic idea is to break up $x\in\{0,1\}^\ell$ into $\ell^{1-\alpha}$ blocks of size ℓ^α; we additionally show how to compress the matrix $\mathbf{A}$ in order to achieve sublinear mpk. The second limitation is that $\mathsf{BALWE}_{\mathsf{struct}}$ assumption does not follow from evasive LWE. To this end, we replace $\mathbf{V}_i \leftarrow \mathbb{Z}_q^{n\times n}$ (more precisely, $\mathbf{V}_i\mathbf{G}$) with $\mathbf{W}_i \leftarrow \mathbb{Z}_q^{n\times m}$ in both $\mathsf{BALWE}_{\mathsf{struct}}$ and our scheme. We refer to the ensuing assumption as ℓ-succinct LWE assumption to emphasize that the assumption is parameterized by ℓ, and we show that ℓ-succinct LWE is implied by evasive LWE (up to a small polynomial loss in parameters).

1.3 Technical Overview

Fix LWE parameters $n,q,m=O(n\log q)$. For notational simplicity, we often omit LWE error terms, or replace them with curly underlines. We proceed to present a self-contained description of our schemes and defer a detailed comparison with WW to Section A.

Trapdoor basis and ℓ-succinct LWE. We start by specifying the trapdoor basis we use in this work. Given a "compression" parameter ℓ, we sample

$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{W} \in \mathbb{Z}_q^{\ell n\times m}$$

along with a random Gaussian $\mathbf{T} = \begin{pmatrix}\overline{\mathbf{T}}\\ \underline{\mathbf{T}}\end{pmatrix} \in \mathbb{Z}^{(\ell+1)m\times \ell m}$ where $\overline{\mathbf{T}}\in\mathbb{Z}^{\ell m\times \ell m}, \underline{\mathbf{T}}\in\mathbb{Z}^{m\times \ell m}$ such that

$$\overbrace{[\mathbf{I}_\ell\otimes\mathbf{B}\mid\mathbf{W}]\cdot\mathbf{T}}^{=(\mathbf{I}_\ell\otimes\mathbf{B})\cdot\overline{\mathbf{T}}+\mathbf{W}\cdot\underline{\mathbf{T}}} = \mathbf{I}_\ell\otimes\mathbf{G} \tag{2}$$

That is, $\mathbf{T}$ is a random gadget trapdoor [41] for $[\mathbf{I}_\ell\otimes\mathbf{B}\mid\mathbf{W}]$. Pictorially, we have

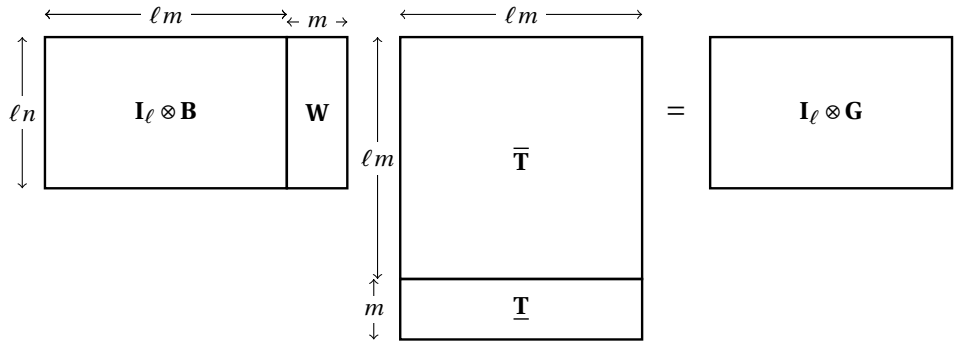

[3] The square matrix $\mathbf{V}_i$ corresponds to $\mathbf{W}_i^{-1}$ in [51] and are used to partially "randomize" $\mathbf{B}\in\mathbb{Z}_q^{n\times m}$ to produce $\mathbf{V}_1^{-1}\mathbf{B},\ldots,\mathbf{V}_\ell^{-1}\mathbf{B}$.

The ℓ-succinct LWE assumption stipulates that

$$(\mathbf{B}, \mathbf{sB}+\mathbf{e}, \mathbf{W}, \mathbf{T}) \approx_c (\mathbf{B}, \mathbf{c}, \mathbf{W}, \mathbf{T})$$

where $\mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^m, \mathbf{c} \leftarrow \mathbb{Z}_q^m$. As a quick sanity check, observe that 1-succinct LWE follows readily from LWE [21]: the reduction samples $\mathbf{W} \leftarrow \mathbb{Z}_q^{n\times m}$ along with a trapdoor, which is used to derive a trapdoor for $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$. Moreover, the assumption becomes stronger as ℓ increases, since we can derive a trapdoor basis for a smaller ℓ from a larger one.

Compressing $\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})$. In the BGGHNSVV ABE, the public key specifies a uniformly random $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times \ell m}$ and the ciphertext for an attribute $\mathbf{x} \in \{0,1\}^\ell$ contains

$$\underbrace{\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})}_{} \in \mathbb{Z}_q^{\ell m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n$$

We show how to recover the above quantity using the above trapdoor basis starting from a compressed ciphertext in $\mathbb{Z}_q^{2m}$. First, we sample an additional matrix $\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n\times m}$, and our compressed ciphertext is given by:

$$\underbrace{\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}]}_{} \in \mathbb{Z}_q^{2m} \tag{3}$$

To recover $\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})$ from (3), we start by multiplying both sides of (2) on the *left* by $\mathbf{x}\otimes\mathbf{I}_n$ and use the fact that $\mathbf{x}\otimes\mathbf{I}_n$ "commutes" with $\mathbf{I}_\ell\otimes\mathbf{B}$ —i.e., $(\mathbf{x}\otimes\mathbf{I}_n)(\mathbf{I}_\ell\otimes\mathbf{B}) = \mathbf{B}(\mathbf{x}\otimes\mathbf{I}_m)$— to obtain:

$$\mathbf{B}\cdot(\mathbf{x}\otimes\mathbf{I}_m)\overline{\mathbf{T}} + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}\cdot\underline{\mathbf{T}} = \mathbf{x}\otimes\mathbf{G} \tag{4}$$

Next, we add $\mathbf{B}_1\underline{\mathbf{T}}$ to both sides of (4) and flip the signs to obtain:[4]

$$[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}] \cdot \overbrace{\begin{pmatrix} -(\mathbf{x}\otimes\mathbf{I}_m)\overline{\mathbf{T}} \\ -\underline{\mathbf{T}} \end{pmatrix}}^{\mathbf{T}_\mathbf{x}\text{ small}} = \overbrace{-\mathbf{B}_1\underline{\mathbf{T}}}^{\mathbf{A}} - \mathbf{x}\otimes\mathbf{G} \tag{5}$$

We can now define $\mathbf{A} := -\mathbf{B}_1\underline{\mathbf{T}}$ and $\mathbf{T}_\mathbf{x} := \left(\begin{smallmatrix} -(\mathbf{x}\otimes\mathbf{I}_m)\overline{\mathbf{T}} \\ -\underline{\mathbf{T}} \end{smallmatrix}\right)$. Multiplying both sides of (5) by $\mathbf{s}$ on the left yields the desired decompression:

$$\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G}) \approx \underbrace{\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}]}_{} \cdot \mathbf{T}_\mathbf{x} \tag{6}$$

Next, we show that replacing a uniformly random $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times \ell m}$ in the BGGHNSVV ABE with $\mathbf{A} := -\mathbf{B}_1\underline{\mathbf{T}}$ does not affect security. Looking ahead to the setting with general α, the fact that $\mathbf{A}$ is deterministically derived from $\mathbf{B}_1, \mathbf{T}$ is crucial for obtaining $o(\ell)$ total parameter size.

Security analysis. Recall that in the proof of selective security for the BGGHNSVV ABE, the reduction receives an LWE challenge $(\mathbf{B}, \underbrace{\mathbf{sB}}_{})$ and an attribute $\mathbf{x}$, samples a low-norm $\mathbf{R} \leftarrow \{0,1\}^{m\times \ell m}$, and programs $\mathbf{A} := \mathbf{BR} + \mathbf{x}\otimes\mathbf{G}$. This allows the reduction to simulate $\underbrace{\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})}_{}$ in the ciphertext given $\underbrace{\mathbf{sB}}_{}$, and answer key queries using a trapdoor derived from $\mathbf{R}$.

In our setting, security will instead rely on ℓ-succinct LWE. The reduction receives a challenge $(\mathbf{B}, \underbrace{\mathbf{sB}}_{}, \mathbf{W}, \mathbf{T})$ and an attribute $\mathbf{x}$, samples a low-norm $\mathbf{U} \leftarrow \{0,1\}^{m\times m}$ and programs $\mathbf{B}_1 := \mathbf{BU} - (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}$. This allows the reduction to simulate $\underbrace{\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}]}_{}$ in the ciphertext given $\underbrace{\mathbf{sB}}_{}$. Next, observe that the matrix $\mathbf{A}$ in our scheme satisfies:

$$\mathbf{A} = \mathbf{B}\cdot\overbrace{[\mathbf{I} \mid \mathbf{U}]\cdot\mathbf{T}_\mathbf{x}}^{\text{small}} + \mathbf{x}\otimes\mathbf{G}$$

This follows from replacing $\mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}$ in (5) with $\mathbf{BU}$. We can then answer key queries as in the BGGHNSVV ABE security reduction with $[\mathbf{I} \mid \mathbf{U}]\cdot\mathbf{T}_\mathbf{x}$ in place of $\mathbf{R}$.

[4] If we parse $\mathbf{W}, \overline{\mathbf{T}}$ respectively as $\mathbf{W}_1, \ldots, \mathbf{W}_\ell \in \mathbb{Z}_q^{n\times m}, \overline{\mathbf{T}}_1, \ldots, \overline{\mathbf{T}}_\ell \in \mathbb{Z}^{m\times m}$ (stacked vertically) as well as $\mathbf{x} = (x_1, \ldots, x_\ell)$, then we can also write (5) as

$$[\mathbf{B} \mid \mathbf{B}_1 + \textstyle\sum x_i\mathbf{W}_i] \cdot \begin{pmatrix} -\sum x_i\overline{\mathbf{T}}_i \\ -\underline{\mathbf{T}} \end{pmatrix} = -\mathbf{B}_1\underline{\mathbf{T}} - \mathbf{x}\otimes\mathbf{G}$$

ABE and LFE with short ciphertexts. Our ABE with $\mathsf{poly}(\lambda, d)$-sized ciphertext follow from applying our compression mechanism to the BGGHNSVV ABE:

- We append $(\mathbf{B}, \mathbf{W}, \mathbf{T}, \mathbf{B}_1)$ to mpk, and replace $\mathbf{A} \leftarrow \mathbb{Z}_q^{\ell n \times m}$ with $\mathbf{A} := -\mathbf{B}_1 \underline{\mathbf{T}}$;
- Key generation is exactly as before, except with the new $\mathbf{A}$;
- We replace $\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$ in the ciphertext with $\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)\mathbf{W}]$;
- Decryption runs decompression as in (6) and proceeds as before;
- In the security proofs, we replace programming $\mathbf{A} = \mathbf{B} \cdot \mathbf{R} + \mathbf{x} \otimes \mathbf{G}$ with programming $\mathbf{B}_1$ as decsribed above, and proceed as before.

Similarly, we obtain an LFE with $\ell + \mathsf{poly}(\lambda, d)$-sized ciphertexts by applying our compression mechanism to the QWW LFE.

Parameter trade-offs. In the rest of this section, we use $O(\cdot)$ to supress $\mathsf{poly}(n, \log q)$ factors. So far, we have $|\mathsf{mpk}| = O(\ell^2)$, dominated by the matrix $\mathbf{T}$. Fix ℓ_0, ℓ_1 such that $\ell_0 \cdot \ell_1 = \ell$. We show how to reduce the size of the public parameter mpk from $O(\ell^2)$ to $O(\ell_0^2 + \ell_1)$, at the cost of increasing the size of the compressed LWE sample ct from $O(1)$ to $O(\ell_1)$; this also gives a way to compress the matrix $\mathbf{A}$ in addition to compressing $\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$. The results in Section 1.1 correspond to setting $\ell_0 = \ell^\alpha, \ell_1 = \ell^{1-\alpha}$.

The basic idea is to divide $\mathbf{x} \in \{0,1\}^\ell$ as well as $\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \in \mathbb{Z}_q^{\ell m}$ into ℓ_1 blocks of size ℓ_0, and run ℓ_1 copies of our base scheme with input length ℓ_0. Naively implementing this idea yields

$$|\mathsf{mpk}| = O(\ell_1 \cdot {\ell_0}^2), \quad |\mathsf{ct}| = O(\ell_1)$$

To get to $|\mathsf{mpk}| = O(\ell_0^2 + \ell_1)$, we reuse $(\mathbf{W}, \mathbf{T})$ for all ℓ_1 blocks (contributing $O(\ell_0^2)$), while sampling a fresh $\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n \times m}$ for each block (contributing $O(\ell_1)$). It is straight-forward to verify that this does not affect functionality. To see why reusing $\mathbf{W}, \mathbf{T}$ is fine for security, observe that the reduction from ℓ-succinct LWE programs $\mathbf{B}_1$ but not $\mathbf{W}$; the latter also means that we need a fresh $\mathbf{B}_1$ for each block for security. As mentioned earlier, we exploit the fact that $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$ is derived from $\mathbf{T}$ and the $\mathbf{B}_1$ matrices to avoid an additive $O(\ell)$ blow-up.

In a bit more detail, we sample

$$\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n \times \ell_1 m}, \mathbf{W} \leftarrow \mathbb{Z}_q^{\ell_0 n \times m}, \mathbf{T} \leftarrow [\mathbf{I}_{\ell_0} \otimes \mathbf{B} \mid \mathbf{W}]^{-1}(\mathbf{I}_{\ell_0} \otimes \mathbf{G})$$

and output as the compressed LWE sample

$$\underbrace{\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})]}_{} \in \mathbb{Z}_q^{(\ell_1+1)m}$$

We can then adapt (5) to obtain

$$[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})] \cdot \overbrace{\begin{pmatrix} -(\mathbf{x} \otimes \mathbf{I}_m)(\mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}}) \\ -\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix}}^{\mathbf{T}_\mathbf{x}\,\text{small}} = \overbrace{-\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}})}^{\mathbf{A}} - \mathbf{x} \otimes \mathbf{G}$$

1.4 Discussion and perspectives

On the use of non-standard lattice assumptions. As mentioned earlier in the introduction, our ABE with short ciphertexts are a substantial improvement over the state of the art of ABE from standard LWE. In fact, starting from standard LWE, we do not even know how to build ABE with $o(\ell)$-sized ciphertexts for very simple circuits, such as linear or NC^0 functionalities (e.g., the index function, which corresponds to broadcast encryption); this is the case even if we allow large secret keys. One way to understand this phenomenon is to look at pairing-based ABE, where all known approaches for $o(\ell)$-sized ciphertexts require q-type assumptions, short of relying on dual system encryption [47].

If we move beyond LWE, then one natural question is, what is the simplest useful non-standard lattice assumption? A natural starting point would be to rely on evasive LWE for a specific distribution of "hints" $\mathbf{B}^{-1}(\mathbf{P})$. Following k-R-SIS

[8], we can consider $\mathbf{P}$ corresponding to polynomials in degree D over ℓ random square matrices $\mathbf{V}_1, \ldots, \mathbf{V}_\ell \leftarrow \mathbb{Z}_q^{n \times n}$; the assumption becomes stronger as D increases. If $D = 1$, then the assumption is equivalent to LWE. $D = 2$ is essentially $\mathsf{BALWE}_{\mathsf{struct}}$ (following [51, § 6.1]), and our results indicate that even $D = 2$ already enables broad feasibility results far beyond what we know from LWE. The ℓ-succinct LWE assumption further relaxes $D = 2$ to allow for products of random wide matrices $\mathbf{W}_i \leftarrow \mathbb{Z}_q^{n \times m}$ with random Gaussian $\mathbf{R}_j \leftarrow \mathbb{Z}^{m \times m}$ (instead of random square matrices).

On ℓ-succinct LWE vs evasive LWE. As mentioned earlier in the introduction, ℓ-succinct LWE being a falsifiable and instance-independent assumption, is better than relying on evasive LWE. Here, we point out two additional technical and conceptual benefits in the context of ABE. First, using ℓ-succinct LWE, we are able to achieve standard selective security for ABE, where only the challenge attribute is fixed in advance. On the other hand, known ABE from evasive LWE in [49,38] only achieve weakly selective security, where the key queries must additionally be fixed in advanced. This is because the key queries determine the distribution $\mathbf{P}$, which must be fixed in advance in evasive LWE (short of defining an interactive variant of evasive LWE). Second, ABE security proofs from evasive LWE side-steps the issue of designing a "crippled" trapdoor for simulating key queries, something we do need to address when basing security from LWE or ℓ-succinct LWE. We hope that all of these considerations, together with the results and techniques in this work, would prompt the research community to move towards the use of falsifiable lattice assumptions like ℓ-succinct LWE, instead of evasive LWE.

Gaining confidence in ℓ-succinct LWE. We begin by noting that all known attacks on LWE have a SIS analogue, so we will treat ℓ-succinct LWE and ℓ-succinct SIS somewhat interchangeably. Given that ℓ-succinct LWE and ℓ-succinct SIS are respectively weaker than evasive LWE and $\mathsf{BASIS}_{\mathsf{struct}}$ (and the closely related k-R-SIS) used in prior works ([49,45,46,48,2,3] for the former, [51,8,10,25,29,50] for the latter), up to polynomial losses in parameters, this gives us significant confidence in ℓ-succinct LWE.

Also, crypt-analysts have started looking at $\mathsf{BASIS}_{\mathsf{struct}}$ and other similar assumptions that fall under the broad umbrella of "SIS with hints" [7] (also, evasive SIS [49]), with the only attack so far being for the knowledge-variant of these assumptions (which are non-falsifiable and much stronger). No non-trivial attacks —beyond ignoring the hints/trapdoor and attacking SIS/LWE directly— have been otherwise discovered so far. We refer to Section 6.3 for concrete intermediate targets for cryptanalysis.

1.5 Additional related works

Related works based on obfuscation. [39] showed that assuming iO, we can get "optimal" ABE (and even FE) for circuits with $|\mathsf{mpk}|, |\mathsf{ct}|, |\mathsf{sk}| = \mathsf{poly}(\lambda)$. Two recent works on LFE from obfuscation (for Turing machines and RAM program respectively) [28,27] achieve ciphertext size $\ell \cdot \mathsf{poly}(\lambda)$. Interestingly, the former refers to their scheme as "asymptotically optimal". To the best of our knowledge, our work is the first to explore LFE with rate one ciphertexts of size $(1 + o(1)) \cdot \ell$.

Pairing-based schemes. We note that our compressed LWE sample shares a similar algebraic structure to the constant-size ciphertext in the pairing-based ABE schemes in [12,14,9] as well as the constant-size ciphertext in the second BGGHNSVV ABE based on multi-linear maps. However, the way we exploit this structure is very different and has no analogue in the pairings setting. In fact, naively translating our technique to the group-based setting would require at least trilinear maps, since the ciphertext, trapdoor basis, and secret keys would need to be encoded in three separate groups.

Improving on the depth dependency. Two recent works [26,38] improved on the dependency on d in existing ABE and LFE schemes, replacing several $\mathsf{poly}(d, \lambda)$ factors with $\mathsf{poly}(\lambda)$ factors; these improvements are orthogonal to the ones in this work, which focuses on the dependency on ℓ, and rely on completely different techniques. In particular, for NC^1 circuits, these works do not improve on the state-of-the-art from LWE, whereas we do. It is easy to see that we can combine our ciphertext compression technique with the LFE and ABE schemes for unbounded-depth circuits in [38]

to remove the $\mathsf{poly}(d)$ factors in our schemes; in particular, this yields an ABE scheme for unbounded-depth circuits with $\mathsf{poly}(\lambda)$-sized ciphertexts and keys. We expect the security proofs to also go through: for LFE, this would require a circular small-secret variant of ℓ-succinct LWE, whereas for ABE, the evasive circular small-secret LWE assumption suffices.

Follow-up works. A follow-up work[5] of Wee and Wu [50] gave new constructions of functional commitments for circuits with fast verification based on ℓ-succinct LWE, starting from the compression mechanism introduced in this work. A more recent work of Champion and Wu [22] built upon our broadcast encryption scheme in Section 4.3 to obtain distributed broadcast encryption schemes from ℓ-succinct LWE.

2 Preliminaries

Notations. We use boldface lower case for row vectors (e.g. $\mathbf{v}$) and boldface upper case for matrices (e.g. $\mathbf{V}$). For integral vectors and matrices (i.e., those over $\mathbb{Z}$), we use the notation $|\mathbf{v}|, |\mathbf{V}|$ to denote the maximum absolute value over all the entries. We use $v \leftarrow \mathscr{D}$ to denote a random sample from a distribution $\mathscr{D}$, as well as $v \leftarrow S$ to denote a uniformly random sample from a set S. We use $\approx_s$ and $\approx_c$ as the abbreviation for statistically close and computationally indistinguishable.

Tensor product. The tensor product (Kronecker product) for matrices $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}^{\ell \times m}$, $\mathbf{B} \in \mathbb{Z}^{n \times p}$ is defined as

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{1,1}\mathbf{B}, & \ldots, & a_{1,m}\mathbf{B} \\ \ldots, & \ldots, & \ldots \\ a_{\ell,1}\mathbf{B}, & \ldots, & a_{\ell,m}\mathbf{B} \end{bmatrix} \in \mathbb{Z}^{\ell n \times mp}.$$

The mixed-product property for tensor product says that

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD})$$

2.1 Lattices background

We use $\mathscr{D}_{\mathbb{Z},\chi}$ to denote the discrete Gaussian distribution over $\mathbb{Z}$ with standard deviation χ.

Learning with errors (LWE). Given $n, m, q, \chi \in \mathbb{N}$, the $\mathsf{LWE}_{n,m,q,\chi}$ assumption states that

$$(\mathbf{B}, \mathbf{sB} + \mathbf{e}) \approx_c (\mathbf{B}, \mathbf{c})$$

where

$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \mathscr{D}_{\mathbb{Z}^m,\chi}, \mathbf{c} \leftarrow \mathbb{Z}_q^m$$

Trapdoor and preimage sampling [41,30]. Given any $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$, $\sigma > 0$, we use $\mathbf{B}^{-1}(\mathbf{Z}, \sigma)$ to denote the distribution of a matrix $\mathbf{Y}$ sampled from $\mathscr{D}_{\mathbb{Z}^{m \times n'},\sigma}$ conditioned on $\mathbf{BY} = \mathbf{Z} \pmod q$. We sometimes suppress σ when the context is clear.

There is an efficient algorithm $\mathsf{TrapGen}(1^n, 1^m, q)$ that, given the modulus $q \geq 2$ and dimension n and $m \geq 2n \log q$, outputs $\mathbf{B} \approx_s U(\mathbb{Z}_q^{n \times 2n \log q})$ with a trapdoor $\mathbf{T}$ such that $\mathbf{BT} = \mathbf{G}$. Moreover, there is an efficient algorithm $\mathsf{SamplePre}(\mathbf{B}, \mathbf{T}, \mathbf{Z}, \sigma)$ that given $\mathbf{B}$ and any $\mathbf{T}$ such that $\mathbf{BT} = \mathbf{G}$, $\sigma \geq 2\sqrt{n \log q} \cdot |\mathbf{T}|$ and $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$, outputs a sample from $\mathbf{B}^{-1}(\mathbf{Z}, \sigma)$. Note that given $\mathbf{B}, \mathbf{T}$ such that $\mathbf{BT} = \mathbf{G}$, we have $[\mathbf{B} \mid \mathbf{B}'] \binom{\mathbf{T}}{\mathbf{0}} = \mathbf{G}$; we will sometimes abuse notation and write $\mathbf{T}$ as a trapdoor for $[\mathbf{B} \mid \mathbf{B}']$.

[5] An earlier version of this work of this work containing only the results for $\alpha = 1$ was submitted to EUROCRYPT 2023.

2.2 Homomorphic Computation on Matrices

Lemma 1 (EvalF, EvalFX [13,31]). *Fix lattice parameters n, q and $m \geq 2n\log q$. Let $\mathscr{F}_{\ell,d,s}$ denote the family of functions $f : \{0,1\}^\ell \to \{0,1\}$ computable by circuits of depth d and size s. There exist a pair of efficient algorithms* (EvalF, EvalFX) *where*

- $\mathsf{EvalF}(\mathbf{A}, f) \to \mathbf{A}_f$*: On input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$ and a function $f \in \mathscr{F}_{\ell,d,s}$, outputs a matrix $\mathbf{A}_f \in \mathbb{Z}_q^{n\times m}$;*
- $\mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x}) \to \mathbf{H}_{\mathbf{A},f,\mathbf{x}}$*: On input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times \ell m}$, a function $f \in \mathscr{F}_{\ell,d,s}$, and an input $\mathbf{x} \in \{0,1\}^\ell$, outputs a matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \in \mathbb{Z}^{\ell m \times m}$.*

For all $\mathbf{A} \in \mathbb{Z}_q^{n\times \ell m}, f \in \mathscr{F}_{\ell,d,s}, \mathbf{x} \in \{0,1\}^\ell$, the matrices $\mathbf{A}_f \leftarrow \mathsf{EvalF}(\mathbf{A}, f)$ and $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \leftarrow \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x})$ satisfy

$$\begin{aligned}(\mathbf{A} - \mathbf{x}\otimes\mathbf{G})\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}} &= \mathbf{A}_f - f(\mathbf{x})\mathbf{G} \\ |\mathbf{H}_{\mathbf{A},f,\mathbf{x}}| &= m^{O(d)}\cdot s\end{aligned} \tag{7}$$

3 ℓ-Succinct Lattice Assumptions

In this section, we introduce the ℓ-succinct LWE assumption as well as its (weaker) SIS analogue ℓ-succinct SIS. The results in this work rely on the former; we state the latter in part to highlight the connection to the $\mathsf{BASIS}_{\mathsf{struct}}$ assumption in [51], which heavily inspired ℓ-succinct LWE. We defer reductions and evidence for hardness of ℓ-succinct LWE and ℓ-succinct SIS to Section 6.

Assumption 1 (ℓ-succinct LWE) *Fix security parameter λ and LWE parameters n, m, q, χ where $m \geq 2n\log q$. The $(\ell, \hat{m}, \sigma)$-succinct LWE assumption where $m \leq \hat{m} \leq \ell m$ stipulates that*

$$(\mathbf{B}, \mathbf{sB} + \mathbf{e}, \mathbf{W}, \mathbf{T}) \approx_c (\mathbf{B}, \mathbf{c}, \mathbf{W}, \mathbf{T})$$

where

$$\begin{aligned}&\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^m, \mathbf{c} \leftarrow \mathbb{Z}_q^m \\ &\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n\times \hat{m}}, \mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)\end{aligned}$$

That is, $\mathbf{T}$ is a random gadget trapdoor [41] with quality σ for the matrix $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$.

We abbreviate the assumption to ℓ-succinct LWE when $\hat{m} = m$ and $\sigma = \mathsf{poly}(\lambda, \ell, m)$. The results in this work primarily rely on polynomial-time hardness of ℓ-succinct LWE for modulus-to-noise ratio $q/\chi \approx 2^{n^\epsilon}$, for some $0 < \epsilon < 1$.

Remark 1. It is easy to see that LWE implies $(\ell, \ell m, \mathsf{poly}(\lambda, \ell, m))$-succinct LWE and in particular 1-succinct LWE: the reduction (following [21]) samples $\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n\times \ell m}$ along with a trapdoor, which is used to derive a trapdoor for $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$ with norm $\mathsf{poly}(\lambda, \ell, m)$.

SIS variant. The SIS assumption for parameters n, q, m, β says that given $\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}$, it is hard to find a non-zero $\mathbf{v} \in \mathbb{Z}_q^m$ such that $\mathbf{Bv} = \mathbf{0} \bmod q$ and $|\mathbf{v}| \leq \beta$. We also introduce the SIS analogue of ℓ-succinct LWE:

Assumption 2 (ℓ-succinct SIS) *Fix SIS parameters n, m, q, β. The succcinct SIS assumption with parameters (ℓ, σ) asserts that SIS is hard w.r.t. $\mathbf{B}$ (i.e., it is hard to find a non-zero $\mathbf{v} \in \mathbb{Z}_q^m$ such that $\mathbf{Bv} = \mathbf{0} \bmod q$ and $|\mathbf{v}| \leq \beta$) given $\mathbf{W}, \mathbf{T}$ where*

$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n\times m}, \mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$$

For the same reason LWE implies SIS, we also have ℓ-succinct LWE implies ℓ-succinct SIS.

4 Attribute-Based Encryption

4.1 Attribute-based encryption

Definition 1 (ABE [44,37]). *A (key-policy) attribute-based encryption (ABE) scheme for some class $\mathscr{F}$ consists of four algorithms:*

- $\mathsf{Setup}(1^\lambda, \mathscr{F}) \to (\mathsf{mpk}, \mathsf{msk})$. *The setup algorithm gets as input the security parameter 1^λ and class description $\mathscr{F}$. It outputs the master public key* mpk *and the master secret key* msk.
- $\mathsf{Enc}(\mathsf{mpk}, x, \mu) \to \mathsf{ct}$. *The encryption algorithm gets as input* mpk*, an input x and a message $\mu \in \{0,1\}^\lambda$. It outputs a ciphertext* ct.
- $\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, f) \to \mathsf{sk}$. *The key generation algorithm gets as input* mpk, msk *and $f \in \mathscr{F}$. It outputs a secret key* sk.
- $\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}, f, \mathsf{ct}, x) \to \mu$. *The decryption algorithm gets as input* sk, f, ct, x *for which $f(x) = 0$ along with* mpk.[6] *It outputs a message μ.*

Correctness. *For all inputs x and f with $f(x) = 0$ and all $\mu \in \{0,1\}^\lambda$, we require*

$$\Pr\left[\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}, f, \mathsf{ct}, x) = \mu : \begin{array}{c}(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathscr{F}) \\ \mathsf{sk} \leftarrow \mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, f) \\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{mpk}, x, \mu)\end{array}\right] = 1 - \mathit{negl}(\lambda).$$

Security. *For a stateful adversary $\mathscr{A}$, we define the advantage function*

$$\mathsf{Adv}^{\mathrm{ABE}}_{\mathscr{A}}(\lambda) := \Pr\left[b = b' : \begin{array}{c} x \leftarrow \mathscr{A}(1^\lambda) \\ (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathscr{F}) \\ (\mu_0, \mu_1) \leftarrow \mathscr{A}^{\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, \cdot)}(\mathsf{mpk}) \\ b \leftarrow \{0,1\};\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{mpk}, x, \mu_b) \\ b' \leftarrow \mathscr{A}^{\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, \cdot)}(\mathsf{ct})\end{array}\right] - \frac{1}{2}$$

with the restriction that all queries f that $\mathscr{A}$ sent to $\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, \cdot)$ *satisfy $f(x) \neq 0$. An ABE scheme is* selectively secure *if for all PPT adversaries $\mathscr{A}$, the advantage $\mathsf{Adv}^{\mathrm{ABE}}_{\mathscr{A}}(\lambda)$ is a negligible function in λ.*

4.2 ABE for Circuits

Construction 1 (ABE for circuits) *We construct an ABE scheme for the family $\mathscr{F}_{\ell,d,s}$ of circuits of depth d and size s, with parameters ℓ_0, ℓ_1 such that $\ell_0 \cdot \ell_1 = \ell$, as follows:*

– $\mathsf{Setup}(1^n, \mathscr{F}_{\ell,d,s})$*: Sample*

$$(\mathbf{B}, \mathbf{T_B}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q),\ \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n \times \ell_1 m},\ \mathbf{W} \leftarrow \mathbb{Z}_q^{\ell_0 n \times m},\ \mathbf{P} \leftarrow \mathbb{Z}_q^{n \times \lambda}$$

$$\mathbf{T} = \begin{pmatrix}\overline{\mathbf{T}} \\ \underline{\mathbf{T}}\end{pmatrix} \leftarrow \mathsf{SamplePre}([\mathbf{I}_{\ell_0} \otimes \mathbf{B} \mid \mathbf{W}], \mathbf{I}_{\ell_0} \otimes \mathbf{T_B}, \mathbf{I}_{\ell_0} \otimes \mathbf{G}, \sigma_0)$$

where $\overline{\mathbf{T}} \in \mathbb{Z}^{\ell_0 m \times \ell_0 m}, \underline{\mathbf{T}} \in \mathbb{Z}^{m \times \ell_0 m}$*. Output*

$$\mathsf{mpk} := \left(\mathbf{B}, \mathbf{B}_1, \mathbf{W}, \mathbf{T}, \mathbf{P}\right) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{n \times \ell_1 m} \times \mathbb{Z}_q^{\ell_0 n \times m} \times \mathbb{Z}_q^{(\ell_0+1)m \times \ell_0 m} \times \mathbb{Z}_q^{n \times \lambda}$$

$$\mathsf{msk} := (\mathbf{T_B})$$

– $\mathsf{Enc}(\mathsf{mpk}, \mathbf{x}, \mathbf{m})$*. Sample*

$$\mathbf{s} \leftarrow \mathbb{Z}_q^n,\ \mathbf{e}_0 \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^m, \mathbf{e}_1 \leftarrow \mathscr{D}_{\mathbb{Z},\chi'}^{\ell_1 m}, \mathbf{e}_2 \leftarrow \mathscr{D}_{\mathbb{Z},\chi'}^{\lambda},$$

Output

$$\mathsf{ct} := \left(\overbrace{\mathbf{sB} + \mathbf{e}_0}^{\mathbf{c}_0}, \overbrace{\mathbf{s}(\mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})) + \mathbf{e}_1}^{\mathbf{c}_1}, \overbrace{\mathbf{sP} + \mathbf{e}_2 + \mathbf{m} \cdot \lfloor \frac{q}{2} \rfloor}^{\mathbf{c}_2}\right) \in \mathbb{Z}_q^m \times \mathbb{Z}_q^{\ell_1 m} \times \mathbb{Z}_q^{\lambda}$$

[6] We follow the convention in [13] where $f(x) = 0$ corresponds to "authorized".

– $\mathsf{KeyGen}(\mathsf{msk}, f)$: *Compute* $\mathbf{A} := -\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}})$ *and* $\mathbf{A}_f := \mathsf{EvalF}(\mathbf{A}, f)$. *Sample*

$$\mathbf{D} \leftarrow \mathsf{SamplePre}([\mathbf{B} \mid \mathbf{A}_f], \mathbf{T}_{\mathbf{B}}, \mathbf{P}, \sigma_1)$$

Output

$$\mathsf{sk} := \mathbf{D} \in \mathbb{Z}^{2m \times \lambda}$$

– $\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk} = \mathbf{D}, f, \mathsf{ct} = (\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2), \mathbf{x})$: *Compute*

$$\begin{aligned}
\mathbf{A} &:= -\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}}), \\
\mathbf{H}_{\mathbf{A},f,\mathbf{x}} &:= \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x}) \\
\mathbf{T}_{\mathbf{x}} &:= \begin{pmatrix} -(\mathbf{x} \otimes \mathbf{I}_m)(\mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}}) \\ -\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix} \\
\mathbf{c}_3 &:= [\mathbf{c}_0 \mid \mathbf{c}_1] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}}.
\end{aligned}$$

Output

$$\left\lfloor \tfrac{2}{q} \cdot \left(\mathbf{c}_2 - [\mathbf{c}_0 \mid \mathbf{c}_3] \cdot \mathbf{D} \bmod q \right) \right\rceil \in \{0,1\}^{\lambda}$$

Parameters. Fix $0 < \epsilon < 1$, where (ℓ_0, m, σ_0)-succinct LWE is hard for a 2^{n^ϵ} modulus-to-noise ratio. We set LWE parameters

$$\begin{aligned}
n &= d^{1/\epsilon} \cdot \mathsf{poly}(\lambda, \log \ell, \log s) \\
m &= O(n^{1+\epsilon}) \\
q &= m^{O(d)} s \cdot \mathsf{poly}(\ell) \cdot \lambda^{\omega(1)} \\
\chi &= \mathsf{poly}(n, \lambda)
\end{aligned}$$

to satisfy

$$\begin{aligned}
q/4 &\geq (\chi + \chi') \cdot \sigma_0 \cdot \sigma_1 \cdot m^{O(d)} s \cdot \mathsf{poly}(m, \lambda) && \text{(correctness)} \\
2^{n^\epsilon} &\geq q/\chi && \text{(modulus-to-noise ratio)} \\
m &\geq 2n \log q \\
\sigma_0 &= \mathsf{poly}(\ell, m, \lambda) && (\ell\text{-succinct LWE}) \\
\sigma_1 &\geq \sigma_0 \cdot m^{O(d)} s \cdot \mathsf{poly}(m, \lambda) && (\mathsf{H}_2 \approx_s \mathsf{H}_3) \\
\chi' &\geq \chi \cdot \sigma_0 \cdot \lambda^{\omega(1)} && (\mathsf{H}_1 \approx_s \mathsf{H}_2)
\end{aligned}$$

where $\mathsf{H}_1, \mathsf{H}_2, \mathsf{H}_3, \mathsf{H}_4$ are defined in the proof below. This yields the following parameter sizes for our ABE scheme:

$$|\mathsf{mpk}| = O_{\lambda,d}({\ell_0}^2 + \ell_1), \quad |\mathsf{ct}| = O_{\lambda,d}(\ell_1), \quad |\mathsf{sk}| = O_{\lambda,d}(1)$$

where $O_{\lambda,d}(\cdot)$ hides factors polynomial in $\lambda, d^{1/\epsilon}$. In particular, setting $\ell_0 = \ell^\alpha, \ell_1 = \ell^{1-\alpha}$ yields

$$|\mathsf{mpk}| = O_{\lambda,d}(\ell^{2\alpha} + \ell^{1-\alpha}), \quad |\mathsf{ct}| = O_{\lambda,d}(\ell^{1-\alpha}), \quad |\mathsf{sk}| = O_{\lambda,d}(1)$$

Remark 2 (Running times). The running times for encryption and decryption are essentially the same as that of the BGGHNSVV ABE. Encryption takes time $\widetilde{O}(\ell_0)$. Decryption takes time $\widetilde{O}(s + {\ell_0}^2)$ in our scheme and $\widetilde{O}(s)$ in BGGHNSVV: they are both dominated by the time taken to compute $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$. Here, $\widetilde{O}(\cdot)$ hides factors polynomial in the lattice parameters and the circuit depth, but it is the same polynomial in both schemes, and basically the same lattice parameters (q could be a polynomial factor larger in our scheme, but the running times only depend on $\log q$).

Remark 3 (Polynomial hardness for NC^1). For NC^1 circuits, we can hope to improve the result to only rely ℓ_0-succinct LWE with a polynomial instead of a sub-exponential modulus-to-noise ratio. To achieve this, we rely on the variant of Lemma 1 in [36,16] for NC^1 circuits achieving $|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}| = O(2^d \cdot s)$. In addition, we can avoid noise flooding in the ciphertexts by taking $\mathbf{e}_1 = \mathbf{e}_0\mathbf{U}, \mathbf{e}_2 = \mathbf{e}_0\mathbf{U}$, as in [1]. We omit this optimization from the current work.

Remark 4 (Compatibility with BGGHNSVV keys). Suppose we have a deployment of the BGGHNSVV scheme with $\mathsf{mpk} = (\mathbf{B}, \mathbf{A}, \mathbf{P})$, $\mathsf{msk} = \mathbf{T_B}$ and secret keys for f satisfying $\mathbf{D} \leftarrow [\mathbf{B} \mid \mathbf{A}_f]^{-1}(\mathbf{P}, \sigma_1)$. We can then use the same $\mathbf{D}$ as a secret key for f in our ABE scheme with the following modifications:

- Append $\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{A} + \mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}}), \sigma_0)$ along with $(\mathbf{B}_1, \mathbf{W}, \mathbf{T})$ to mpk.
- Decryption proceeds as above except with

$$\mathbf{T}_{\mathbf{x}} = \begin{pmatrix} \boxed{\mathbf{K}} - (\mathbf{x} \otimes \mathbf{I}_m)(\mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}}) \\ -\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix}$$

- In the security proof, we sample $\mathbf{K} \leftarrow \mathscr{D}_{\mathbb{Z}, \sigma_0}^{m \times \ell m}$ and program $\mathbf{A} := \mathbf{B}\mathbf{K} + \mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}})$

A technical claim. We begin by proving the equation we use for decompression:

Claim. Suppose $\ell_0 \cdot \ell_1 = \ell$ and $[\mathbf{I}_{\ell_0} \otimes \mathbf{B} \mid \mathbf{W}] \cdot \binom{\overline{\mathbf{T}}}{\underline{\mathbf{T}}} = \mathbf{I}_{\ell_0} \otimes \mathbf{G}$. Then, for all $\mathbf{x} \in \{0,1\}^{\ell}$, we have:

$$[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})] \cdot \overbrace{\begin{pmatrix} -(\mathbf{x} \otimes \mathbf{I}_m)(\mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}}) \\ -\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix}}^{\mathbf{T}_{\mathbf{x}}} = \overbrace{-\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}})}^{\mathbf{A}} - \mathbf{x} \otimes \mathbf{G} \tag{8}$$

Proof. Observe that

$$[\mathbf{I}_{\ell} \otimes \mathbf{B} \mid \mathbf{I}_{\ell_1} \otimes \mathbf{W}] \cdot \begin{pmatrix} \mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}} \\ \mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix} = \mathbf{I}_{\ell} \otimes \mathbf{G} \tag{9}$$

Multiplying both sides of (9) on the left by $-\mathbf{x} \otimes \mathbf{I}_n$, and observing $(\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell} \otimes \mathbf{B}) = \mathbf{B}(\mathbf{x} \otimes \mathbf{I}_m)$, we obtain

$$[\mathbf{B} \mid (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})] \begin{pmatrix} -(\mathbf{x} \otimes \mathbf{I}_m)(\mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}}) \\ -\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix} = -\mathbf{x} \otimes \mathbf{G}$$

Adding $-\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}})$ to both sides yields the claim above. □

Correctness. Combining (8) with (7), we have

$$[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A}, f, \mathbf{x}} = \mathbf{A}_f - f(\mathbf{x})\mathbf{G} \tag{10}$$

This means that whenever $f(\mathbf{x}) = 0$,

$$\begin{aligned} \mathbf{c}_3 &\approx \mathbf{s}(\mathbf{A}_f - f(x)\mathbf{G}) = \mathbf{s}\mathbf{A}_f \\ [\mathbf{c}_0 \mid \mathbf{c}_3] \cdot \mathbf{D} &\approx \mathbf{s}[\mathbf{B} \mid \mathbf{A}_f] \cdot \mathbf{D} = \mathbf{s}\mathbf{P} \\ \mathbf{c}_2 - [\mathbf{c}_0 \mid \mathbf{c}_3] \cdot \mathbf{D} &\approx \mathbf{m} \cdot \lfloor \tfrac{q}{2} \rfloor \end{aligned} \tag{11}$$

The error term in the final $\approx$ is given by

$$\mathbf{e}_2 - [\mathbf{e}_0 \mid ([\mathbf{e}_0 \mid \mathbf{e}_1] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A}, f, \mathbf{x}})] \cdot \mathbf{D}$$

whose norm is bounded by

$$\overbrace{(\chi + \chi')}^{\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2} \cdot \overbrace{\sigma_0}^{\overline{\mathbf{T}}, \underline{\mathbf{T}}} \cdot \overbrace{\sigma_1}^{\mathbf{D}} \cdot \overbrace{m^{O(d)}}^{\mathbf{H}_{\mathbf{A}, f, \mathbf{x}}} s \cdot \mathsf{poly}(m, \lambda)$$

Correctness follows as long as the preceding quantity is bounded by $q/4$.

Theorem 2. *Under the (ℓ_0, m, σ_0)-succinct LWE assumption, Construction 1 is a selectively secure ABE scheme.*

Proof. We define a series of games:

– H_0: This is the real ABE security game.

– H_1: Same as H_1, except the challenger samples $\mathbf{B}_1, \mathbf{P}$ as follows:
  1. samples $\mathbf{U} \leftarrow \{0,1\}^{m\times m}$, and programs $\mathbf{B}_1 := \mathbf{BU} - (\mathbf{x}\otimes\mathbf{I}_n)(\mathbf{I}_{\ell_1}\otimes\mathbf{W})$
  2. samples $\mathbf{U}_0 \leftarrow \{0,1\}^{m\times\lambda}$, and programs $\mathbf{P} := \mathbf{BU}_0$.

  $\mathsf{H}_0 \approx_s \mathsf{H}_1$ follows readily from left-over hash lemma.

– H_2: Same as H_1, except the challenger in Enc samples $\mathbf{c}_1 := \mathbf{c}_0\mathbf{U} + \mathbf{e}_1, \mathbf{c}_2 := \mathbf{c}_0\mathbf{U}_0 + \mathbf{e}_2$.

  $\mathsf{H}_1 \approx_s \mathsf{H}_2$ follows readily from noise-flooding, along with $\mathbf{c}_0\mathbf{U} \approx \mathbf{sBU} = \mathbf{s}(\mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)(\mathbf{I}_{\ell_1}\otimes\mathbf{W}))$ and $\mathbf{c}_0\mathbf{U}_0 \approx \mathbf{sBU}_0 = \mathbf{sP}$.

– H_3: Same as H_2, except the challenger in KeyGen samples $\mathbf{D}$ using $\mathsf{SamplePre}([\mathbf{B} \mid \mathbf{A}_f], \left(\begin{smallmatrix}-[\mathbf{I}|\mathbf{U}]\cdot\mathbf{T}_\mathbf{x}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\\ \mathbf{I}_m\end{smallmatrix}\right), \mathbf{P}, \sigma_1)$ instead of $\mathsf{SamplePre}([\mathbf{B} \mid \mathbf{A}_f], \mathbf{T}_\mathbf{B}, \mathbf{P}, \sigma_1)$.

  $\mathsf{H}_2 \approx_s \mathsf{H}_3$ follows from trapdoor sampling together with the following:
  - substituting $\mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)(\mathbf{I}_{\ell_1}\otimes\mathbf{W}) = \mathbf{BU}$ into (10) yields

$$[\mathbf{B} \mid \mathbf{BU}]\cdot\mathbf{T}_\mathbf{x}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{B}\cdot[\mathbf{I} \mid \mathbf{U}]\cdot\mathbf{T}_\mathbf{x}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(x)\mathbf{G} \tag{12}$$

    and thus $[\mathbf{B} \mid \mathbf{A}_f]\cdot\left(\begin{smallmatrix}-[\mathbf{I}|\mathbf{U}]\cdot\mathbf{T}_\mathbf{x}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\\ \mathbf{I}_m\end{smallmatrix}\right) = f(x)\mathbf{G}, f(x)\neq 0$.
  - $\left|[\mathbf{I} \mid \mathbf{U}]\cdot\mathbf{T}_\mathbf{x}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\right| = \sigma_0\cdot m^{O(d)}s\cdot\mathsf{poly}(m,\lambda)$.

– H_4: Same as H_3, except the challenger samples $\mathbf{c}_0 \leftarrow \mathbb{Z}_q^m$.

  $\mathsf{H}_3 \approx_c \mathsf{H}_4$ follows from (ℓ_0, m, σ_0)-succinct LWE.

– H_5: Same as H_4, except the challenger samples $\mathbf{c}_2 \leftarrow \mathbb{Z}_q^\lambda$.

  $\mathsf{H}_4 \approx_s \mathsf{H}_5$ follows from left-over hash lemma, which tells us $(\mathbf{B}, \mathbf{c}_0, \mathbf{BU}_0, \mathbf{c}_0\mathbf{U}_0)$ is statistically close to uniform.

In H_5, the challenge bit b is perfectly hidden, so the advantage is 0. □

4.3 Broadcast Encryption

ABE for circuits captures broadcast encryption for $N = \ell$ users as a special case: the input $\mathbf{x} \in \{0,1\}^\ell$ corresponds to characteristic vector for the broadcast set, and we can check membership with circuits in $\mathscr{F}_{N,1,N}$ via an inner product.[7] This way, we can rely on ℓ-succinct LWE with a $n^{\omega(1)}$ modulus-to-noise ratio.

Corollary 1 (Broadcast encryption). *Assuming (ℓ, m)-succinct LWE with $n^{\omega(1)}$ modulus-to-noise ratio, we have a broadcast encryption scheme for N users with parameters*

$$|\mathsf{mpk}| = N^2\cdot\mathsf{poly}(\lambda, \log N), \quad |\mathsf{ct}| = \mathsf{poly}(\lambda, \log N), \quad |\mathsf{sk}| = \mathsf{poly}(\lambda, \log N)$$

This is the first post-quantum broadcast encryption scheme with sub-linear size ciphertext based on a simple, falsifiable assumption. We refer to Appendix B for an additional ABE for inner product and for broadcast encryption with a smaller mpk but a larger sk.

4.4 Reusable Garbled Circuits

Goldwasser et al. [32], with improvements from Boneh et al. [13], showed that starting from (i) an ABE scheme for $\mathscr{F}_{\ell,d,s}$ with mpk, ciphertext and key sizes $P(\ell,d,s), C(\ell,d,s), K(\ell,d,s)$, and (ii) the LWE assumption (used for FHE with rate one ciphertexts), we can construct a reusable garbling scheme for $\mathscr{F}_{\ell,d,s}$ in the CRS model where

– the CRS has size $P(\ell', d', s')$;
– the garbled input has size $\ell' + \mathsf{poly}(\lambda)\cdot C(\ell', d', s')$;
– the garbled circuit has size $s + \mathsf{poly}(\lambda)\cdot K(\ell', d', s')$;

[7] For inner product with vectors $\mathbf{y} \in \{-1,0,1\}^\ell$ (which suffices for broadcast encryption), Lemma 1 simply asserts that $[\mathbf{A} - \mathbf{x}\otimes\mathbf{G}]\cdot(\mathbf{y}^\top\otimes\mathbf{I}_m) = \mathbf{A}(\mathbf{y}^\top\otimes\mathbf{I}_m) - \mathbf{x}\mathbf{y}^\top\otimes\mathbf{G}$. For arbitrary $\mathbf{y} \in \mathbb{Z}_q^\ell$, we can simply replace $\mathbf{y}^\top\otimes\mathbf{I}_m$ with $(\mathbf{I}_\ell\otimes\mathbf{G})^{-1}(\mathbf{y}^\top\otimes\mathbf{G})$.

where $\ell' = \ell + \mathsf{poly}(\lambda, d), d' = d \cdot \mathsf{poly}(\lambda), s' = s \cdot \mathsf{poly}(\lambda, d)$. Here, ℓ' is the size of a FHE encryption of $x \in \{0,1\}^\ell$ and d', s' correspond to the depth and the size of the circuit performing FHE homomorphic evaluation of f plus symmetric-key decryption. Combined with our ABE scheme in Construction 1, we have the following corollary:

Corollary 2 (Reusable garbling scheme). *Assuming ℓ_0-succinct LWE with 2^{n^ϵ} modulus-to-noise ratio, we have a reusable garbling scheme for $\mathscr{F}_{\ell,d,s}$ in the CRS model where*

- *the CRS has size $O_{\lambda,d}(\ell^{2\alpha} + \ell^{1-\alpha})$*
- *the garbled input has size $\ell + O_{\lambda,d}(\ell^{1-\alpha})$, and*
- *the garbled circuit has size $s + O_{\lambda,d}(1)$.*

Here, $O_{\lambda,d}(\cdot)$ hides factors polynomial in $\lambda, d^{1/\epsilon}$.

5 Laconic Function Evaluation

5.1 Definition of LFE

Definition 2 (LFE [43,24]). *A laconic function evaluation (LFE) scheme for some class $\mathscr{F}$ consists of four algorithms* Setup, Compress, Enc, Dec.

$\mathsf{Setup}(1^\lambda, \mathscr{F})$ *takes as input the security parameter 1^λ and circuit parameters $\mathscr{F}$ and outputs a common reference string* crs.
$\mathsf{Compress}(\mathsf{crs}, f)$ *is a deterministic algorithm that takes as input* crs *and $f \in \mathscr{F}$ and outputs a digest* dig.
$\mathsf{Enc}(\mathsf{crs}, \mathsf{dig}, x)$ *takes as input* crs*, a digest* dig *and a message x and outputs a ciphertext* ct.
$\mathsf{Dec}(\mathsf{crs}, f, \mathsf{ct})$ *takes as input* crs*, $f \in \mathscr{F}$, and a ciphertext* ct *and outputs a message y.*

Correctness. *We require that for all $\lambda, \mathscr{F}$ and $f \in \mathscr{F}$:*

$$\Pr\left[y = f(x) \;\middle|\; \begin{array}{lll} \mathsf{crs} & \leftarrow \mathsf{Setup}(1^\lambda, \mathscr{F}) \\ \mathsf{dig} & = \mathsf{Compress}(\mathsf{crs}, f) \\ \mathsf{ct} & \leftarrow \mathsf{Enc}(\mathsf{crs}, \mathsf{dig}, x) \\ y & \leftarrow \mathsf{Dec}(\mathsf{crs}, f, \mathsf{ct}) \end{array} \right] = 1.$$

Selective security. *We require that there exists a PPT simulator* Sim *such that for all stateful PPT adversary $\mathscr{A}$, we have:*

$$\left| \Pr\left[\mathrm{EXP}_{LFE}^{Real}(1^\lambda) = 1 \right] - \Pr\left[\mathrm{EXP}_{LFE}^{Ideal}(1^\lambda) \right] \right| \leq \mathsf{negl}(\lambda)$$

for the experiments $\mathrm{EXP}_{LFE}^{Real}(1^\lambda)$ and $\mathrm{EXP}_{LFE}^{Ideal}(1^\lambda)$ defined below:

$\mathrm{EXP}_{LFE}^{Real}(1^\lambda)$:

0. $(\mathscr{F}, x) \leftarrow \mathscr{A}(1^\lambda)$
1. $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, \mathscr{F})$
2. $f \leftarrow \mathscr{A}(\mathsf{crs})$:
3. $\mathsf{dig} = \mathsf{Compress}(\mathsf{crs}, f)$
4. $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{crs}, \mathsf{dig}, x)$
5. *Output* $\mathscr{A}(\mathsf{ct})$

$\mathrm{EXP}_{LFE}^{Ideal}(1^\lambda)$:

0. $(\mathscr{F}, x) \leftarrow \mathscr{A}(1^\lambda)$
1. $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, \mathscr{F})$
2. $f \leftarrow \mathscr{A}(\mathsf{crs})$:
3. $\mathsf{dig} = \mathsf{Compress}(\mathsf{crs}, f)$
4. $\mathsf{ct} \leftarrow \mathsf{Sim}(\mathsf{crs}, \mathsf{dig}, f, f(x))$
5. *Output* $\mathscr{A}(\mathsf{ct})$

5.2 LFE for Circuits

Following QWW [43], we start by constructing AB-LFE for circuits, which corresponds to LFE for the following functionality:

$$(\mathbf{x}, \mathbf{m}_0, \mathbf{m}_1) \in \{0,1\}^{\ell} \times \{0,1\}^{\lambda} \times \{0,1\}^{\lambda} \overset{f \in \mathscr{F}_{\ell,d,s}}{\longmapsto} (\mathbf{x}, \mathbf{m}_{f(\mathbf{x})})$$

Our formalization of AB-LFE corresponds to the two-outcome variant in [43, Section 4.4]. As in QWW:

- the digest is simply $\mathbf{A}_f$;
- the ciphertext contains $\mathbf{x}$ along with a compression of $\underline{\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$, and we additionally use $\underline{\mathbf{s}\mathbf{A}_f}$ to mask $\mathbf{m}_0$, and $\underline{\mathbf{s}(\mathbf{A}_f - \mathbf{G})}$ to mask $\mathbf{m}_1$.

Construction 3 (AB-LFE for circuits) *We construct an AB-LFE scheme for the family $\mathscr{F}_{\ell_0,d,s}$ of circuits of depth d and size s, with parameters ℓ_0, ℓ_1 such that $\ell_0 \cdot \ell_1 = \ell$, as follows:*

- $\mathsf{Setup}(1^n, \mathscr{F}_{\ell_0,d,s})$*: Sample*

$$(\mathbf{B}, \mathbf{T}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q),\ \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n \times \ell_1 m},\ \mathbf{W} \leftarrow \mathbb{Z}_q^{\ell_0 n \times m},$$
$$\mathbf{T} = \begin{pmatrix} \overline{\mathbf{T}} \\ \underline{\mathbf{T}} \end{pmatrix} \leftarrow \mathsf{SamplePre}([\mathbf{I}_{\ell_0} \otimes \mathbf{B} \mid \mathbf{W}], \mathbf{I}_{\ell_0} \otimes \mathbf{T}_{\mathbf{B}}, \mathbf{I}_{\ell_0} \otimes \mathbf{G}, \sigma_0)$$

  where $\overline{\mathbf{T}} \in \mathbb{Z}^{\ell_0 m \times \ell_0 m}, \underline{\mathbf{T}} \in \mathbb{Z}^{m \times \ell_0 m}$*. Output*

$$\mathsf{crs} := \left(\mathbf{B}, \mathbf{B}_1, \mathbf{W}, \mathbf{T}\right) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{n \times \ell_1 m} \times \mathbb{Z}_q^{\ell_0 n \times m} \times \mathbb{Z}_q^{(\ell_0+1) m \times \ell_0 m}$$

- $\mathsf{Compress}(\mathsf{crs}, f)$*: Compute* $\mathbf{A} := -\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}})$ *and* $\mathbf{A}_f := \mathsf{EvalF}(\mathbf{A}, f)$*. Output*

$$\mathsf{dig} := \mathbf{A}_f \in \mathbb{Z}^{n \times m}$$

- $\mathsf{Enc}(\mathsf{crs}, \mathbf{A}_f, (\mathbf{x}, \mathbf{m}_0, \mathbf{m}_1))$*. Sample*

$$\mathbf{s} \leftarrow \mathbb{Z}_q^n,\ \mathbf{e}_0 \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^m, \mathbf{e}_1 \leftarrow \mathscr{D}_{\mathbb{Z},\chi'}^{\ell_1 m}, \mathbf{e}_{2,0}, \mathbf{e}_{2,1} \leftarrow \mathscr{D}_{\mathbb{Z},\chi''}^{\lambda}, \mathbf{P}_0, \mathbf{P}_1 \leftarrow \mathbb{Z}_q^{n \times \lambda}$$

  Compute[8]

$$\begin{aligned}
\mathbf{c}_0 &:= \mathbf{s}\mathbf{B} + \mathbf{e}_0 \\
\mathbf{c}_1 &:= \mathbf{s}(\mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})) + \mathbf{e}_1 \\
\mathbf{c}_{2,0} &:= \mathbf{s}\mathbf{A}_f \cdot \mathbf{G}^{-1}(\mathbf{P}_0) + \mathbf{m}_0 \cdot \lfloor \tfrac{q}{2} \rfloor + \mathbf{e}_{2,0} \\
\mathbf{c}_{2,1} &:= \mathbf{s}(\mathbf{A}_f - \mathbf{G}) \cdot \mathbf{G}^{-1}(\mathbf{P}_1) + \mathbf{m}_1 \cdot \lfloor \tfrac{q}{2} \rfloor + \mathbf{e}_{2,1}
\end{aligned}$$

  Output

$$\mathsf{ct} := \left(\mathbf{x}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_{2,0}, \mathbf{c}_{2,1}, \mathbf{P}_0, \mathbf{P}_1\right) \in \{0,1\}^{\ell} \times \mathbb{Z}_q^m \times \mathbb{Z}_q^{\ell_1 m} \times (\mathbb{Z}_q^{\lambda})^2 \times (\mathbb{Z}_q^{n \times \lambda})^2$$

- $\mathsf{Dec}(\mathsf{crs} = (\mathbf{A}_f), f, \mathsf{ct} = (\mathbf{x}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_{2,0}, \mathbf{c}_{2,1}, \mathbf{P}_0, \mathbf{P}_1))$*: Compute*

$$\begin{aligned}
\mathbf{A} &:= -\mathbf{B}_1(\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}}), \\
\mathbf{H}_{\mathbf{A},f,\mathbf{x}} &:= \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x}) \\
\mathbf{T}_{\mathbf{x}} &:= \begin{pmatrix} -(\mathbf{x} \otimes \mathbf{I}_m)(\mathbf{I}_{\ell_1} \otimes \overline{\mathbf{T}}) \\ -\mathbf{I}_{\ell_1} \otimes \underline{\mathbf{T}} \end{pmatrix} \\
\mathbf{c}_3 &:= [\mathbf{c}_0 \mid \mathbf{c}_1] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}}.
\end{aligned}$$

  Output

$$\left\lfloor \tfrac{2}{q} \cdot \left(\mathbf{c}_{2,f(\mathbf{x})} - \mathbf{c}_3 \cdot \mathbf{G}^{-1}(\mathbf{P}_{f(\mathbf{x})}) \bmod q\right) \right\rceil \in \{0,1\}^{\lambda}$$

[8] Here, $\mathbf{G}^{-1}(\cdot)$ denotes the standard deterministic entry-wise bit decomposition.

Parameters. Fix $0 < \epsilon < 1$, where (ℓ_0, m, σ_0)-succinct LWE is hard for a 2^{n^ϵ} modulus-to-noise ratio. We will set LWE parameters as in our ABE scheme in Section 4.2

$$\begin{aligned} n &= d^{1/\epsilon} \cdot \mathsf{poly}(\lambda, \log \ell, \log s) \\ m &= O(n^{1+\epsilon}) \\ q &= m^{O(d)} s \cdot \mathsf{poly}(\ell) \cdot \lambda^{\omega(1)} \\ \chi &= \mathsf{poly}(n, \lambda) \end{aligned}$$

which also satisfy the following minor modifications to the constraints pertaining to χ'' (in place of σ_1):

$$\begin{aligned} q/4 &\geq (\chi'' + (\chi + \chi') \cdot \sigma_0 \cdot m^{O(d)} s) \cdot \mathsf{poly}(m, \lambda) && \text{(correctness)} \\ \chi'' &\geq (\chi + \chi') \cdot \sigma_0 \cdot m^{O(d)} s) \cdot \mathsf{poly}(m, \lambda) \cdot \lambda^{\omega(1)} && (\mathsf{H}_0 \approx_s \mathsf{H}_1 \text{ in proof below}) \end{aligned}$$

This yields the following parameter sizes for our AB-LFE scheme:

$$|\mathsf{crs}| = O_{\lambda,d}({\ell_0}^2 + \ell_1), \quad |\mathsf{dig}| = O_{\lambda,d}(1), \quad |\mathsf{ct}| = \ell + O_{\lambda,d}(\ell_1)$$

and the encryption running time is $O_{\lambda,d}(\ell)$. Here, $O_{\lambda,d}(\cdot)$ hides factors polynomial in $\lambda, d^{1/\epsilon}$. In particular, setting $\ell_0 = \ell^\alpha, \ell_1 = \ell^{1-\alpha}$ yields

$$|\mathsf{crs}| = O_{\lambda,d}(\ell^{2\alpha} + \ell^{1-\alpha}), \quad |\mathsf{dig}| = O_{\lambda,d}(1), \quad |\mathsf{ct}| = \ell + O_{\lambda,d}(\ell^{1-\alpha})$$

Correctness. As in Section 4.2, we have from (11) that $\mathbf{c}_3 \approx \mathbf{s}(\mathbf{A}_f - f(x)\mathbf{G})$. Therefore,

$$\mathbf{c}_{2,f(\mathbf{x})} \approx \mathbf{c}_3 \cdot \mathbf{G}^{-1}(\mathbf{P}_{f(\mathbf{x})}) + \mathbf{m}_{f(\mathbf{x})} \cdot \lfloor \tfrac{q}{2} \rfloor \tag{13}$$

The error term in the above $\approx$ is given by

$$\mathbf{e}_{2,f(\mathbf{x})} - ([\mathbf{e}_0 \mid \mathbf{e}_1] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}}) \cdot \mathbf{G}^{-1}(\mathbf{P}_{f(\mathbf{x})})$$

whose norm is bounded by

$$(\overbrace{\chi''}^{\mathbf{e}_{2,f(\mathbf{x})}} + \overbrace{(\chi + \chi')}^{\mathbf{e}_0,\mathbf{e}_1} \cdot \overbrace{\sigma_0}^{\overline{\mathbf{T}},\underline{\mathbf{T}}} \cdot \overbrace{m^{O(d)} s}^{\mathbf{H}_{\mathbf{A},f,\mathbf{x}}}) \cdot \mathsf{poly}(m, \lambda)$$

Correctness follows as long as the preceding quantity is bounded by $q/4$.

Theorem 4. *Under the (ℓ_0, m, σ_0)-succinct LWE assumption, Construction 3 is selectively secure.*

Proof. We begin by specifying the simulator:

– $\mathsf{Sim}(\mathsf{crs}, \mathsf{dig}, f, (\mathbf{x}, \mathbf{z}))$: Compute $f(\mathbf{x}) \in \{0, 1\}$, and sample

$$\mathbf{c}_0 \leftarrow \mathbb{Z}_q^m, \mathbf{c}_1 \leftarrow \mathbb{Z}_q^m, \mathbf{e}_{2,f(\mathbf{x})} \leftarrow \mathscr{D}_{\mathbb{Z},\chi''}^\lambda, \mathbf{c}_{2,1-f(\mathbf{x})} \leftarrow \mathbb{Z}_q^\lambda, \mathbf{P}_0, \mathbf{P}_1 \leftarrow \mathbb{Z}_q^{n \times \lambda}$$

Compute

$$\begin{aligned} \mathbf{c}_3 &:= [\mathbf{c}_0 \mid \mathbf{c}_1] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} \quad \text{(same as in Dec)} \\ \mathbf{c}_{2,f(\mathbf{x})} &:= \mathbf{c}_3 \cdot \mathbf{G}^{-1}(\mathbf{P}_{f(\mathbf{x})}) + \mathbf{z} \cdot \lfloor \tfrac{q}{2} \rfloor + \mathbf{e}_{2,f(\mathbf{x})} \end{aligned}$$

Output

$$\mathsf{ct} := \left(\mathbf{x}, \mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_{2,0}, \mathbf{c}_{2,1}, \mathbf{P}_0, \mathbf{P}_1\right)$$

We define a series of games:

– H_0: This is the real AB-LFE security game.

– H_1: Same as H_0, except the challenger computes $\mathbf{c}_3$ as in Sim, and $\mathbf{c}_{2,0}, \mathbf{c}_{2,1}$ as follows:

$$\mathbf{c}_{2,b} := \mathbf{c}_3 \cdot \mathbf{G}^{-1}(\mathbf{P}_b) + \mathbf{m}_b \cdot \lfloor \tfrac{q}{2} \rfloor + (f(\mathbf{x}) - b) \cdot \mathbf{sP}_b + \mathbf{e}_{2,b}, \forall b \in \{0,1\}$$

$\mathsf{H}_0 \approx_s \mathsf{H}_1$ follows from

- a straight-forward adaptation of (13) which tells us $\mathbf{c}_{2,b}$ in H_0 satisfies:

$$\mathbf{c}_{2,b} \approx \mathbf{c}_3 \cdot \mathbf{G}^{-1}(\mathbf{P}_b) + \mathbf{m}_b \cdot \lfloor \tfrac{q}{2} \rfloor + (f(\mathbf{x}) - b) \cdot \mathbf{sP}_b, \forall b \in \{0,1\}$$

- noise-flooding using $\mathbf{e}_{2,b}$ to flood the error term

$$([\mathbf{e}_0 \mid \mathbf{e}_1] \cdot \mathbf{T}_{\mathbf{x}} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}}) \cdot \mathbf{G}^{-1}(\mathbf{P}_b)$$

– H_2: Same as H_1, except the challenger samples $\mathbf{B}_1, \mathbf{P}_{1-f(\mathbf{x})}$ as follows:
1. samples $\mathbf{U} \leftarrow \{0,1\}^{m \times m}$, and programs $\mathbf{B}_1 := \mathbf{BU} - (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W})$
2. samples $\mathbf{U}_{1-f(\mathbf{x})} \leftarrow \{0,1\}^{m \times \lambda}$, and programs $\mathbf{P}_{1-f(\mathbf{x})} = \mathbf{BU}_{1-f(\mathbf{x})}$.

$\mathsf{H}_1 \approx_s \mathsf{H}_2$ follows readily from left-over hash lemma.

– H_3: Same as H_2, except the challenger in Enc samples $\mathbf{c}_1, \mathbf{c}_{2,1-f(\mathbf{x})}$ as follows:

$$\mathbf{c}_1 := \mathbf{c}_0\mathbf{U} + \mathbf{e}_1$$
$$\mathbf{c}_{2,b} := \mathbf{c}_3 \cdot \mathbf{G}^{-1}(\mathbf{P}_b) + \mathbf{m}_b \cdot \lfloor \tfrac{q}{2} \rfloor + (f(\mathbf{x}) - b) \cdot \mathbf{c}_0\mathbf{U}_b + \mathbf{e}_{2,b}, \;\; b = 1 - f(\mathbf{x})$$

$\mathsf{H}_2 \approx_s \mathsf{H}_3$ follows readily from noise-flooding along with $\mathbf{c}_0\mathbf{U} \approx \mathbf{sBU} = \mathbf{s}(\mathbf{B}_1 + (\mathbf{x} \otimes \mathbf{I}_n)(\mathbf{I}_{\ell_1} \otimes \mathbf{W}))$ and $\mathbf{c}_0\mathbf{U}_{1-f(\mathbf{x})} \approx \mathbf{sBU}_{1-f(\mathbf{x})} = \mathbf{sP}_{1-f(\mathbf{x})}$.

– H_4: Same as H_3, except the challenger samples $\mathbf{c}_0 \leftarrow \mathbb{Z}_q^m$.

$\mathsf{H}_3 \approx_c \mathsf{H}_4$ follows from (ℓ_0, m, σ_0)-succinct LWE.

– H_5: Same as H_4, except the challenger samples $\mathbf{c}_1 \leftarrow \mathbb{Z}_q^m, \mathbf{c}_{2,1-f(\mathbf{x})} \leftarrow \mathbb{Z}_q^\lambda$.

$\mathsf{H}_4 \approx_s \mathsf{H}_5$ follows from left-over hash lemma, which tells us $(\mathbf{B}, \mathbf{c}_0, \mathbf{BU}, \mathbf{c}_0\mathbf{U}, \mathbf{BU}_{1-f(\mathbf{x})}, \mathbf{c}_0\mathbf{U}_{1-f(\mathbf{x})})$ is statistically close to uniform.

Observe that H_5 is exactly the output of Sim, since $\mathbf{z} = \mathbf{m}_{f(\mathbf{x})}$. □

From AB-LFE to LFE. Prior work [43] showed —via a construction similar to that in Section 4.4— that starting from (i) an AB-LFE scheme for $\mathscr{F}_{\ell,d,s}$ with CRS, ciphertext and digest sizes $P(\ell,d,s), \ell + C(\ell,d,s), K(\ell,d,s)$, and (ii) the LWE assumption (used for FHE with rate one ciphertexts), we can construct an LFE scheme for $\mathscr{F}_{\ell,d,s}$ where

$$|\mathsf{crs}| = P(\ell',d',s'), \quad |\mathsf{dig}| = \mathsf{poly}(\lambda) \cdot K(\ell',d',s'), \quad |\mathsf{ct}| = \ell' + \mathsf{poly}(\lambda) \cdot C(\ell',d',s')$$

where $\ell' = \ell + \mathsf{poly}(\lambda, d), d' = \mathsf{poly}(\lambda, d), s' = s \cdot \mathsf{poly}(\lambda, d)$. Combined with our AB-LFE scheme in Construction 3, we have the following corollary:

Corollary 3 (LFE for circuits). *Assuming ℓ^α-succinct LWE with 2^{n^ϵ} modulus-to-noise ratio, we have an LFE scheme for $\mathscr{F}_{\ell,d,s}$ where*

$$|\mathsf{crs}| = O_{\lambda,d}(\ell^{2\alpha} + \ell^{1-\alpha}), \quad |\mathsf{dig}| = O_{\lambda,d}(1), \quad |\mathsf{ct}| = \ell + O_{\lambda,d}(\ell^{1-\alpha})$$

and the encryption running time is $O_{\lambda,d}(\ell)$. Here, $O_{\lambda,d}(\cdot)$ hides factors polynomial in $\lambda, d^{1/\epsilon}$.

6 Reductions for ℓ-Succinct LWE

In this section, we relate ℓ-succinct LWE to other recently introduced lattice assumptions in [49,45,51], in order to gain confidence in our assumption. That is, our goal is to establish that our assumptions are *qualitatively* weaker. We clarify that we think we believe the best way to set parameters for ℓ-succinct LWE is based on direct crypt-analysis on the assumption, and we provide concrete targets in Section 6.3. As such, we do not try to optimize on the parameters in these reductions showing hardness for ℓ-succinct LWE and its SIS variant. We further clarify that the known devastating attacks on non-standard lattice assumptions used in multi-linear maps and obfuscation tend to be "complete breaks" that are not affected by small parameter losses in reductions.

6.1 Relation to $\mathsf{BASIS}_{\mathsf{struct}}$ in [51]

Next, we state the $\mathsf{BASIS}_{\mathsf{struct}}$ assumption in [51] (which corresponds to replacing $\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n \times m}$ in ℓ-succinct SIS with $\mathbf{W} := \mathbf{VG}, \mathbf{V} \leftarrow \mathbb{Z}_q^{\ell n \times n}$) and show that it implies ℓ-succinct SIS.

Assumption 3 ($\mathsf{BASIS}_{\mathsf{struct}}$ [51]) *Fix SIS parameters n, m, q, β The $\mathsf{BASIS}_{\mathsf{struct}}$ assumption with parameters (ℓ, σ) asserts that SIS is hard w.r.t. $\mathbf{B}$ (i.e., it is hard to find a non-zero $\mathbf{v} \in \mathbb{Z}_q^m$ such that $\mathbf{Bv} = \mathbf{0} \bmod q$ and $|\mathbf{v}| \leq \beta$) given $\mathbf{V}, \mathbf{T}$ where*

$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{V} \leftarrow \mathbb{Z}_q^{\ell n \times n}, \mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{VG}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$$

Remark 5. Our statement above is slightly more compact and general than that in [51]. The latter considers q prime, samples $\mathbf{V}_1, \ldots, \mathbf{V}_\ell \leftarrow \mathbb{Z}_q^{n \times n}$ (which are invertible w.h.p. if q is prime; invertibility is necessary for the schemes in [51]), and defines

$$\tilde{\mathbf{B}} := \begin{pmatrix} \mathbf{V}_1\mathbf{B} & \cdots & -\mathbf{G} \\ & \mathbf{V}_2\mathbf{B} \cdots & -\mathbf{G} \\ & \ddots & \vdots \quad \vdots \\ & \cdots & \mathbf{V}_\ell\mathbf{B} \; -\mathbf{G} \end{pmatrix}$$

and then samples $\mathbf{T} \leftarrow \tilde{\mathbf{B}}^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$, i.e. $\mathbf{T}$ is a random gadget trapdoor for $\tilde{\mathbf{B}}$. As shown in [41], for any invertible square matrix $\mathbf{M}$, random gadget trapdoors for $\tilde{\mathbf{B}}$ and for $\mathbf{M} \cdot \tilde{\mathbf{B}}$ are equivalent (i.e., we can efficiently convert between the two, up to small polynomial losses in the quality of the trapdoor). This means that whenever $\mathbf{V}_1, \ldots, \mathbf{V}_\ell$ are invertible, we can instead give out a random gadget trapdoor for

$$\begin{pmatrix} \mathbf{B} & \cdots & -\mathbf{V}_1^{-1}\mathbf{G} \\ & \mathbf{B} \cdots & -\mathbf{V}_2^{-1}\mathbf{G} \\ & \ddots & \vdots \quad \vdots \\ & \cdots & \mathbf{B} \; -\mathbf{V}_\ell^{-1}\mathbf{G} \end{pmatrix}$$

Setting $\mathbf{V}$ to the vertical concatenation of $-\mathbf{V}_1^{-1}, \ldots, -\mathbf{V}_\ell^{-1}$ yields the above formulation.

Lemma 2. *Fix SIS parameters n, m, q, β. Then, $\mathsf{BASIS}_{\mathsf{struct}}$ with parameters (ℓ, σ) implies ℓ-succinct SIS with parameters $(\ell, \sigma \cdot \mathsf{poly}(\ell, m))$.*

Proof. The reduction is straight-forward: we use the techniques in [1,41] to "randomize" $\mathbf{VG}$ to obtain a uniformly random $\mathbf{W}$ while transforming a trapdoor for $[\mathbf{I} \otimes \mathbf{B} \mid \mathbf{VG}]$ to one for $[\mathbf{I} \otimes \mathbf{B} \mid \mathbf{W}]$. Given $\mathbf{V} \leftarrow \mathbb{Z}_q^{\ell n \times n}, \mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{VG}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma)$, we sample $\mathbf{R} \leftarrow \{0,1\}^{\ell m \times \ell}$, and program

$$\mathbf{W} := (\mathbf{I}_\ell \otimes \mathbf{B})\mathbf{R} + \mathbf{VG}$$

so that $\mathbf{W}$ is statistically random, by the left-over hash lemma. Next, observe that

$$[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}] \cdot \overbrace{\begin{pmatrix} \mathbf{I} & -\mathbf{R} \\ & \mathbf{I} \end{pmatrix}}^{\mathbf{T}'} \cdot \mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G}$$

We can then use $\mathbf{T}'$ to sample $\mathbf{T}'' \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]^{-1}(\mathbf{I}_\ell \otimes \mathbf{G}, \sigma \cdot \mathsf{poly}(\ell, m))$. Now, run the adversary that breaks ℓ-succinct SIS on input $\mathbf{B}, \mathbf{W}, \mathbf{T}''$. □

6.2 Relation to Evasive LWE in [49,45]

The evasive LWE assumption was recently introduced in [49,45] and has since been used in [46,48,2,3]. We show that ℓ-succinct LWE follows from evasive LWE and LWE, in the super-polynomial modulus q regime, which is the primary setting for our ABE. All known lattice attacks scale well with the modulus and we do not expect attacks that work in the polynomial modulus regime but not in the super-polynomial regime (allowing some loss in the latter).

Assumption 4 (Evasive LWE [49,45]) *Fix LWE parameters n, m, q, χ and an efficiently samplable matrix distribution $\mathbf{P}$ along with public-coin auxiliary input* aux. *The evasive LWE assumption asserts that*

$$\begin{aligned} &\textbf{\textit{if}} && (\mathbf{B}, \mathbf{P}, \mathsf{aux}, \mathbf{sB}+\mathbf{e}, \mathbf{sP}+\mathbf{e}'') \approx_c (\mathbf{B}, \mathbf{P}, \mathsf{aux}, \mathbf{c}, \mathbf{c}''), \\ &\textbf{\textit{then}} && (\mathbf{B}, \mathsf{aux}, \mathbf{sB}+\mathbf{e}, \mathbf{B}^{-1}(\mathbf{P})) \approx_c (\mathbf{B}, \mathsf{aux}, \mathbf{c}, \mathbf{B}^{-1}(\mathbf{P})) \end{aligned}$$

where $\mathbf{c}, \mathbf{c}''$ are uniformly random and $\mathbf{e}''$ is a fresh noise vector.

Essentially, it says that given $\mathbf{B}, \mathbf{sB}+\mathbf{e}$, getting the additional component $\mathbf{B}^{-1}(\mathbf{P})$ is no more useful than just getting the product $(\mathbf{sB}+\mathbf{e}) \cdot \mathbf{B}^{-1}(\mathbf{P}) \approx \mathbf{sP}+\mathbf{e}''$.

Lemma 3. *Fix LWE parameters n, m, q, χ. Then, evasive LWE (plus LWE) implies ℓ-succinct LWE, up to small polynomial losses in the quality of the trapdoor.*

The proof is similar to that in [51, Section 6.1], where they related evasive LWE to $\mathsf{BASIS}_{\mathsf{struct}}$ in the LWE setting.

Proof. First, we consider a variant of ℓ-succinct LWE where $\mathbf{T}$ is an Ajtai trapdoor for $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$, i.e.

$$\mathbf{T} \leftarrow [\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]^{-1}(\mathbf{0}^{\ell n \times 2\ell m}, \sigma) \tag{14}$$

Note that we can efficiently convert between gadget trapdoors and Ajtai trapdoors, up to small polynomial losses in the quality of the trapdoor. Therefore, it suffices to show the result for an Ajtai trapdoor $\mathbf{T}$. Next, we recast the distribution of $\mathbf{T}$ in (14) in the form $\mathbf{B}^{-1}(\mathbf{P})$ following [51, Theorem 3.15]:

- Next, observe that the distribution of $\mathbf{T}$ in (14) is statistically close to the following distribution:

$$\mathbf{T} = \begin{pmatrix} \overline{\mathbf{T}} \\ \mathbf{R} \end{pmatrix} : \mathbf{R} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{m \times 2\ell m}, \overline{\mathbf{T}} \leftarrow (\mathbf{I}_\ell \otimes \mathbf{B})^{-1}(-\mathbf{WR}, \sigma)$$

  This follows from basis delegation [21] (also [51, Lemma 2.7]).

- Let us write $\mathbf{W} = \begin{pmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_\ell \end{pmatrix}$ where $\mathbf{W}_1, \ldots, \mathbf{W}_\ell \in \mathbb{Z}_q^{n \times m}$. Then,

$$(\mathbf{I}_\ell \otimes \mathbf{B})^{-1}(-\mathbf{WR}) \approx_s \begin{pmatrix} -\mathbf{B}^{-1}(\mathbf{W}_1\mathbf{R}) \\ \vdots \\ -\mathbf{B}^{-1}(\mathbf{W}_\ell\mathbf{R}) \end{pmatrix}$$

- Putting the two together, this means that $(\mathbf{T}, \mathbf{W})$ is completely determined given $\mathbf{W}_1, \ldots, \mathbf{W}_\ell \leftarrow \mathbb{Z}_q^{n \times m}$ along with

$$\mathbf{B}^{-1}(\mathbf{W}_1\mathbf{R}), \ldots, \mathbf{B}^{-1}(\mathbf{W}_\ell\mathbf{R}), \mathbf{R}$$

  where $\mathbf{R} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{m \times 2\ell m}$.

That is, to show ℓ-succinct LWE, it suffices to show that $(\mathbf{B}, \mathbf{sB}+\mathbf{e})$ is pseudorandom given $\mathbf{B}^{-1}([\mathbf{W}_1\mathbf{R} \mid \cdots \mid \mathbf{W}_\ell\mathbf{R}]), \mathbf{W}_1, \ldots, \mathbf{W}_\ell, \mathbf{R}$.

Now, we apply evasive LWE to the distribution

$$\mathbf{P} = [\mathbf{W}_1\mathbf{R} \mid \cdots \mid \mathbf{W}_\ell\mathbf{R}], \mathsf{aux} = (\mathbf{W}_1, \ldots, \mathbf{W}_\ell, \mathbf{R})$$

For the pre-condition, we have:

$$\begin{aligned} &(\mathbf{B}, \mathbf{sB}+\mathbf{e}, \{\mathbf{sW}_i\mathbf{R}+\mathbf{e}_i'', \mathbf{W}_i\}, \mathbf{R}) \\ &\approx_s (\mathbf{B}, \mathbf{sB}+\mathbf{e}, \{(\mathbf{sW}_i+\mathbf{e}_i)\mathbf{R}+\mathbf{e}_i'', \mathbf{W}_i\}, \mathbf{R}) \\ &\approx_c (\mathbf{B}, \mathbf{c}, \{\mathbf{s}_i\mathbf{R}+\mathbf{e}_i'', \mathbf{W}_i\}, \mathbf{R}), && \mathbf{c}, \mathbf{s}_i \leftarrow \mathbb{Z}_q^m \\ &\approx_c (\mathbf{B}, \mathbf{c}, \{\mathbf{c}_i, \mathbf{W}_i\}, \mathbf{R}), && \mathbf{c}_i \leftarrow \mathbb{Z}_q^{2\ell m} \end{aligned}$$

where the first $\approx_s$ uses noise flooding, and both $\approx_c$ relies on LWE (the latter using [15]). We may then conclude that evasive LWE plus LWE implies ℓ-succinct LWE. □

6.3 Targets for crypt-analysis

We suggest the following intermediate targets as open problems for cryptanalysis corresponding to potential improvements over attacking SIS/LWE directly, as a step towards understanding hardness:

- An attack on ℓ-succinct SIS/LWE that exploits $\mathbf{T}$ to obtain $\mathsf{poly}(\ell)$ speed-up over the best exponential-time attacks on standard SIS/LWE (i.e., a lattice analogue to Cheon's attack on q-type assumptions in pairings [23])
- An attack with complexity exponential in $\hat{m}$, where $\hat{m}$ is the width of $\mathbf{W}$ (c.f. Section 3). Indeed, when $\hat{m} = 0$, we have a trapdoor for $\mathbf{I}_\ell \otimes \mathbf{B}$, which breaks the assumption.

These targets are consistent with the following facts: ℓ-succinct SIS/LWE assumption becomes stronger as $\hat{m}$ decreases and as ℓ increases, via standard lattice delegation [21]. We stress that these targets are far from contradicting polynomial hardness of ℓ-succinct LWE with sub-exponential modulus-to-noise ratio, which is what we need for our results.

Acknowledgments. We thank David Wu for our lovely collaboration in [51] as well as numerous insightful discussions. We also thank the reviewers for the thoughtful feedback. Part of this work was done while visiting Divesh Aggarwal at CQT.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Heidelberg, May / June 2010.
2. S. Agrawal, S. Kumari, A. Yadav, and S. Yamada. Broadcast, trace and revoke with optimal parameters from polynomial hardness. In C. Hazay and M. Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 605–636. Springer, Heidelberg, Apr. 2023.
3. S. Agrawal, M. Rossi, A. Yadav, and S. Yamada. Constant input attribute based (and predicate) encryption from evasive and tensor LWE. In H. Handschuh and A. Lysyanskaya, editors, *CRYPTO*, 2023.
4. S. Agrawal, D. Wichs, and S. Yamada. Optimal broadcast encryption from LWE and pairings in the standard model. In R. Pass and K. Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 149–178. Springer, Heidelberg, Nov. 2020.
5. S. Agrawal and S. Yamada. CP-ABE for circuits (and more) in the symmetric key setting. In R. Pass and K. Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 117–148. Springer, Heidelberg, Nov. 2020.
6. S. Agrawal and S. Yamada. Optimal broadcast encryption from pairings and LWE. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 13–43. Springer, Heidelberg, May 2020.
7. M. Albrecht. Sis with hints zoo, 2023. https://malb.io/sis-with-hints.html.
8. M. R. Albrecht, V. Cini, R. W. F. Lai, G. Malavolta, and S. A. K. Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Y. Dodis and T. Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Heidelberg, Aug. 2022.
9. N. Attrapadung, B. Libert, and E. de Panafieu. Expressive key-policy attribute-based encryption with constant-size ciphertexts. In D. Catalano, N. Fazio, R. Gennaro, and A. Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 90–108. Springer, Heidelberg, Mar. 2011.
10. D. Balbás, D. Catalano, D. Fiore, and R. W. F. Lai. Functional commitments for circuits from falsifiable assumptions. In *TCC*, 2023.
11. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 719–737. Springer, Heidelberg, Apr. 2012.
12. D. Boneh, X. Boyen, and E.-J. Goh. Hierarchical identity based encryption with constant size ciphertext. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 440–456. Springer, Heidelberg, May 2005.
13. D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Heidelberg, May 2014.
14. D. Boneh, C. Gentry, and B. Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In V. Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 258–275. Springer, Heidelberg, Aug. 2005.
15. D. Boneh, K. Lewi, H. W. Montgomery, and A. Raghunathan. Key homomorphic PRFs and their applications. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Heidelberg, Aug. 2013.

16. Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In M. Naor, editor, *ITCS 2014*, pages 1–12. ACM, Jan. 2014.
17. Z. Brakerski and V. Vaikuntanathan. Constrained key-homomorphic PRFs from standard lattice assumptions - or: How to secretly embed a circuit in your PRF. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 1–30. Springer, Heidelberg, Mar. 2015.
18. Z. Brakerski and V. Vaikuntanathan. Circuit-ABE from LWE: Unbounded attributes and semi-adaptive security. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 363–384. Springer, Heidelberg, Aug. 2016.
19. Z. Brakerski and V. Vaikuntanathan. Lattice-inspired broadcast encryption and succinct ciphertext-policy ABE. In *ITCS*, pages 28:1–28:20, 2022.
20. R. Canetti, Y. Chen, J. Holmgren, A. Lombardi, G. N. Rothblum, R. D. Rothblum, and D. Wichs. Fiat-Shamir: from practice to theory. In M. Charikar and E. Cohen, editors, *51st ACM STOC*, pages 1082–1090. ACM Press, June 2019.
21. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 523–552. Springer, Heidelberg, May / June 2010.
22. J. Champion and D. J. Wu. Distributed broadcast encryption from lattices. In *TCC*, 2024.
23. J. H. Cheon. Security analysis of the strong Diffie-Hellman problem. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 1–11. Springer, Heidelberg, May / June 2006.
24. C. Cho, N. Döttling, S. Garg, D. Gupta, P. Miao, and A. Polychroniadou. Laconic oblivious transfer and its applications. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 33–65. Springer, Heidelberg, Aug. 2017.
25. V. Cini, R. W. F. Lai, and G. Malavolta. Lattice-based succinct arguments from vanishing polynomials. In *CRYPTO*, 2023.
26. V. Cini and H. Wee. Abe for circuits with $\mathsf{poly}(\lambda)$-sized keys from LWE. In *FOCS*, 2023.
27. F. Dong, Z. Hao, E. Mook, and D. Wichs. Laconic function evaluation, functional encryption and obfuscation for RAMs with sublinear computation. In *EUROCRYPT*, 2024.
28. N. Döttling, P. Gajland, and G. Malavolta. Laconic function evaluation for turing machines. In A. Boldyreva and V. Kolesnikov, editors, *PKC 2023, Part II*, volume 13941 of *LNCS*, pages 606–634. Springer, Heidelberg, May 2023.
29. B. Fisch, Z. Liu, and P. Vesely. Orbweaver: Succinct linear functional commitments from lattices. In *CRYPTO*, 2023.
30. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In R. E. Ladner and C. Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
31. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, Aug. 2013.
32. S. Goldwasser, Y. T. Kalai, R. A. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013.
33. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013.
34. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Predicate encryption for circuits from LWE. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 503–523. Springer, Heidelberg, Aug. 2015.
35. S. Gorbunov, V. Vaikuntanathan, and D. Wichs. Leveled fully homomorphic signatures from standard lattices. In R. A. Servedio and R. Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015.
36. S. Gorbunov and D. Vinayagamurthy. Riding on asymmetry: Efficient ABE for branching programs. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 550–574. Springer, Heidelberg, Nov. / Dec. 2015.
37. V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In A. Juels, R. N. Wright, and S. De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, Oct. / Nov. 2006. Available as Cryptology ePrint Archive Report 2006/309.
38. Y.-C. Hsieh, H. Lin, and J. Luo. Attribute-based encryption for circuits of unbounded depth from lattices: Garbled circuits of optimal size, laconic functional evaluation, and more. In *FOCS*, 2023.
39. A. Jain, H. Lin, and J. Luo. On the optimal succinctness and efficiency of functional encryption and attribute-based encryption. In C. Hazay and M. Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 479–510. Springer, Heidelberg, Apr. 2023.
40. H. Li, H. Lin, and J. Luo. ABE for circuits with constant-size secret keys and adaptive security. In E. Kiltz and V. Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 680–710. Springer, Heidelberg, Nov. 2022.
41. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, Apr. 2012.
42. C. Peikert and S. Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Heidelberg, Aug. 2019.
43. W. Quach, H. Wee, and D. Wichs. Laconic function evaluation and applications. In M. Thorup, editor, *59th FOCS*, pages 859–870. IEEE Computer Society Press, Oct. 2018.

44. A. Sahai and B. R. Waters. Fuzzy identity-based encryption. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
45. R. Tsabary. Candidate witness encryption from lattice techniques. In Y. Dodis and T. Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 535–559. Springer, Heidelberg, Aug. 2022.
46. V. Vaikuntanathan, H. Wee, and D. Wichs. Witness encryption and null-IO from evasive LWE. In S. Agrawal and D. Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 195–221. Springer, Heidelberg, Dec. 2022.
47. B. Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 619–636. Springer, Heidelberg, Aug. 2009.
48. B. Waters, H. Wee, and D. J. Wu. Multi-authority ABE from lattices without random oracles. In E. Kiltz and V. Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 651–679. Springer, Heidelberg, Nov. 2022.
49. H. Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In O. Dunkelman and S. Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Heidelberg, May / June 2022.
50. H. Wee and D. J. Wu. Lattice-based functional commitments: Fast veriôcation and cryptanalysis. In *ASIACRYPT*, 2023.
51. H. Wee and D. J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In C. Hazay and M. Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 385–416. Springer, Heidelberg, Apr. 2023.

A Comparison with WW

In this section, we describe an alternative derivation of our scheme, starting from the WW commitment scheme.

The WW commitment scheme. We begin with the core WW commitment scheme in [51, Remark 4.12], adapted to the notation and setting in this work. The scheme achieves succinct commitments of size independent of the input length ℓ; this succinct commitment can in turn be expanded to a GVW commitment of the same input.

- The public parameters comprise
$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{V} \in \mathbb{Z}_q^{\ell n\times n}$$
along with a random Gaussian $\mathbf{T} = \begin{pmatrix}\overline{\mathbf{T}}\\ \underline{\mathbf{T}}\end{pmatrix} \in \mathbb{Z}^{(\ell+1)m\times \ell m}$ where $\overline{\mathbf{T}} \in \mathbb{Z}^{\ell m\times \ell m}, \underline{\mathbf{T}} \in \mathbb{Z}^{m\times \ell m}$ such that
$$\overbrace{[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{VG}]\cdot \mathbf{T}}^{=(\mathbf{I}_\ell\otimes\mathbf{B})\cdot\overline{\mathbf{T}}+\mathbf{VG}\cdot\underline{\mathbf{T}}} = \mathbf{I}_\ell \otimes \mathbf{G} \tag{15}$$
That is, $\mathbf{T}$ is a random gadget trapdoor [41] for $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{VG}]$.[9]
- Given $\mathbf{x} \in \{0,1\}^\ell$, we multiply both sides of (15) on the right by $\mathbf{x}^\top \otimes \mathbf{I}_m$ to obtain
$$(\mathbf{I}_\ell \otimes \mathbf{B})\cdot \overbrace{\overline{\mathbf{T}}(\mathbf{x}^\top \otimes \mathbf{I}_m)}^{\text{opening}} + \mathbf{V}\cdot \overbrace{\mathbf{G}\underline{\mathbf{T}}(\mathbf{x}^\top \otimes \mathbf{I}_m)}^{\text{commitment}} = \mathbf{x}^\top \otimes \mathbf{G} \tag{16}$$
The commitment $\mathbf{C}$ to $\mathbf{x} \in \{0,1\}^\ell$ is given by $\mathbf{G}\cdot\underline{\mathbf{T}}(\mathbf{x}^\top\otimes\mathbf{I}_m) \in \mathbb{Z}_q^{n\times m}$ and the opening by $\overline{\mathbf{T}}(\mathbf{x}^\top\otimes\mathbf{I}_m) \in \mathbb{Z}^{\ell m\times m}$. Verification checks that the opening has low norm and satisfies the above relation in (16).

Binding follows from the $\mathsf{BASIS}_{\mathsf{struct}}$ assumption, which states that SIS is hard with respect to $\mathbf{B}$, given $\mathbf{V}, \mathbf{T}$. Moreover, we can expand $\mathbf{C}$ into $\mathbf{V}\cdot\mathbf{C} \in \mathbb{Z}_q^{\ell n\times m}$, which is a GVW commitment to $\mathbf{x}$ with opening $\overline{\mathbf{T}}(\mathbf{x}^\top \otimes \mathbf{I}_m)$.

Compressing $\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})$. In the BGGHNSVV ABE, the public key specifies a uniformly random $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times \ell m}$ and the ciphertext for an attribute $\mathbf{x} \in \{0,1\}^\ell$ contains
$$\underbrace{\mathbf{s}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})}_{\sim} \in \mathbb{Z}_q^{\ell m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n$$
Our goal is to compress the above quantity into a vector in $\mathbb{Z}_q^{O(m)}$ using $\mathbf{B}, \mathbf{V}, \mathbf{T}$.

First idea. A natural strategy following GVW would be to use $\underbrace{\mathbf{sC}}_{\sim}$ as the compressed ciphertext, where $\mathbf{C}$ is a homomorphic commitment to $\mathbf{x}$ (looking ahead, we will rely on homomorphic opening in the security proof). Instantiating this idea with the WW commitment is problematic because multiplying $\mathbf{C}$ on the left by $\mathbf{V}$ as in (16) interacts poorly with both the error term $\mathbf{e}$ and the secret $\mathbf{s}$. Instead, we will modify the commitment scheme and (16) as follows. We start by multiplying both sides of (15) on the *left* by $\mathbf{x}\otimes\mathbf{I}_n$ and use the fact that $\mathbf{x}\otimes\mathbf{I}_n$ "commutes" with $\mathbf{I}_\ell\otimes\mathbf{B}$ —i.e., $(\mathbf{x}\otimes\mathbf{I}_n)(\mathbf{I}_\ell\otimes\mathbf{B}) = \mathbf{B}(\mathbf{x}\otimes\mathbf{I}_m)$— to obtain:
$$\mathbf{B}\cdot \overbrace{(\mathbf{x}\otimes\mathbf{I}_m)\overline{\mathbf{T}}}^{\text{opening}} + \overbrace{(\mathbf{x}\otimes\mathbf{I}_n)\mathbf{VG}}^{\text{commitment}}\cdot\underline{\mathbf{T}} = \mathbf{x}\otimes\mathbf{G} \tag{17}$$
Now, consider a commitment $\mathbf{C}$ to $\mathbf{x}$ is given by $(\mathbf{x}\otimes\mathbf{I}_n)\mathbf{VG} \in \mathbb{Z}_q^{n\times m}$. This fixes both of the issues above: multiplying $\mathbf{C}$ on the right by the low-norm matrix $\underline{\mathbf{T}}$ is compatible with both $\mathbf{e}$ and $\mathbf{s}$, but introduces a security issue – given $\mathbf{sC}+\mathbf{e} = \mathbf{s}(\mathbf{x}\otimes\mathbf{I}_n)\mathbf{VG}+\mathbf{e}$, we can efficiently recover $\mathbf{s}$ due to the gadget matrix $\mathbf{G}$ in $\mathbf{C}$.

[9] The scheme as stated in WW parses $\mathbf{V}$ as $\mathbf{V}_1, \ldots, \mathbf{V}_\ell \in \mathbb{Z}_q^{n\times n}$, and gives out a random gadget trapdoor for
$$\begin{pmatrix}\mathbf{V}_1^{-1}\mathbf{B} & \cdots & \mathbf{G}\\ \vdots & \ddots & \vdots\\ & \cdots\ \mathbf{V}_\ell^{-1}\mathbf{B} & \mathbf{G}\end{pmatrix}$$

Second idea. To solve the latter issue, we append to the public key a matrix $\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n\times m}$, and our compressed ciphertext is now given by:

$$\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes \mathbf{I}_n)\mathbf{V}\mathbf{G}] \in \mathbb{Z}_q^{2m} \tag{18}$$

Towards decompression, add $\mathbf{B}_1\underline{\mathbf{T}}$ to both sides of (17) and flip the signs to obtain:

$$[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes \mathbf{I}_n)\mathbf{V}\mathbf{G}] \cdot \overbrace{\begin{pmatrix} -(\mathbf{x}\otimes \mathbf{I}_m)\overline{\mathbf{T}} \\ -\underline{\mathbf{T}} \end{pmatrix}}^{\mathbf{T}_\mathbf{x}\text{ small}} = \overbrace{-\mathbf{B}_1\underline{\mathbf{T}}}^{\mathbf{A}} - \mathbf{x}\otimes\mathbf{G} \tag{19}$$

We can now define $\mathbf{A} := -\mathbf{B}_1\underline{\mathbf{T}}$ and $\mathbf{T}_\mathbf{x} := \left(\begin{smallmatrix} -(\mathbf{x}\otimes \mathbf{I}_m)\overline{\mathbf{T}} \\ -\underline{\mathbf{T}} \end{smallmatrix}\right)$. Multiplying both sides of (19) by $\mathbf{s}$ on the left yields the desired decompression:

$$\mathbf{s}(\mathbf{A} - \mathbf{x}\otimes\mathbf{G}) \approx \mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes \mathbf{I}_n)\mathbf{V}\mathbf{G}] \cdot \mathbf{T}_\mathbf{x} \tag{20}$$

Weakening the assumption. The security of our scheme so far would rely on $\mathsf{BALWE}_{\mathsf{struct}}$ (the LWE analogue of $\mathsf{BASIS}_{\mathsf{struct}}$ introduced in [51]), namely $(\mathbf{B}, \mathbf{s}\mathbf{B}+\mathbf{e})$ is pseudorandom, given $\mathbf{V}, \mathbf{T}$. As noted in [51, § 6.1], $\mathsf{BALWE}_{\mathsf{struct}}$ is implied by evasive LWE plus the following non-standard variant of LWE (related to building simpler PRFs from lattices, c.f., the discussion in [11, §1.2, 1.3]), namely:

$$(\mathbf{B}, \mathbf{V}_1, \ldots, \mathbf{V}_\ell, \mathbf{R}, \mathbf{s}\mathbf{B}+\mathbf{e}, \mathbf{s}\mathbf{V}_i\mathbf{R}+\mathbf{e}'_i) \tag{21}$$

is pseudorandom, where $\mathbf{V}_i \leftarrow \mathbb{Z}_q^{n\times n}, \mathbf{R} \leftarrow \mathbb{Z}_q^{n\times 2m}, \mathbf{e} \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^m, \mathbf{e}'_i \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^{2m}$.

In this work, we introduce ℓ-succinct LWE, where we replace $\mathbf{W}$ in $\mathsf{BALWE}_{\mathsf{struct}}$ with $\mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n\times m}$. That is, ℓ-succinct LWE states that $(\mathbf{B}, \mathbf{s}\mathbf{B}+\mathbf{e})$ is pseudorandom, given $\mathbf{W}, \mathbf{T}$, where $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}] \cdot \mathbf{T} = \mathbf{I}_\ell \otimes \mathbf{G}$. We would then also replace $\mathbf{W}$ in our compressed LWE sample in (18) with $\mathbf{W}$ to obtain:

$$\mathbf{s}[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes \mathbf{I}_n)\mathbf{W}] \in \mathbb{Z}_q^{2m}$$

Extending the analysis in [51, § 6.1], we have that ℓ-succinct LWE is implied by evasive LWE, plus pseudorandomness of the following distribution:

$$(\mathbf{B}, \mathbf{W}_1, \ldots, \mathbf{W}_\ell, \mathbf{R}, \mathbf{s}\mathbf{B}+\mathbf{e}, \mathbf{s}\mathbf{W}_i\mathbf{R}+\mathbf{e}'_i)$$

where $\mathbf{W}_i \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{R} \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^{m\times 2m}, \mathbf{e} \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^m, \mathbf{e}'_i \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^{2m}$. The key distinctions from (21) are that $\mathbf{W}_i$ are wider than $\mathbf{V}_i$, and that $\mathbf{R}$ has low-norm, which allow us to base pseudorandomness of the latter on LWE, following [15].

B ABE for Inner Product

For the special case of inner product, we sketch a way to reduce the size of mpk in Construction 1 from $\widetilde{O}(\ell^2)$ to $\widetilde{O}(\ell)$, while blowing up the size of sk to $\widetilde{O}(\ell)$.

Construction 5 (ABE for inner product) *We construct an ABE scheme for inner product specified by*

$$\mathbf{x} \in \{0,1\}^\ell \overset{\mathbf{y}\in\mathbb{Z}_q^\ell}{\longmapsto} \mathbf{x}\mathbf{y}^\top$$

as follows:

- $\mathsf{Setup}(1^n, 1^\ell)$*: Sample*

$$(\mathbf{B}, \mathbf{T}_\mathbf{B}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q),\ \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n\times m},\ \mathbf{W} \leftarrow \mathbb{Z}_q^{\ell n\times m},\ \mathbf{P} \leftarrow \mathbb{Z}_q^{n\times\lambda}$$

  Output

$$\mathsf{mpk} := \big(\mathbf{B}, \mathbf{B}_1, \mathbf{W}, \mathbf{P}\big), \quad \mathsf{msk} := (\mathbf{T}_\mathbf{B})$$

– $\mathsf{Enc}(\mathsf{mpk}, \mathbf{x}, \mathbf{m})$. *Sample*

$$\mathbf{s} \leftarrow \mathbb{Z}_q^n,\ \mathbf{e}_0 \leftarrow \mathscr{D}_{\mathbb{Z},\chi}^m, \mathbf{e}_1 \leftarrow \mathscr{D}_{\mathbb{Z},\chi'}^m, \mathbf{e}_2 \leftarrow \mathscr{D}_{\mathbb{Z},\chi'}^\lambda,$$

Output

$$\mathsf{ct} := \Big(\overbrace{\mathbf{sB}+\mathbf{e}_0}^{\mathbf{c}_0},\ \overbrace{\mathbf{s}(\mathbf{B}_1+(\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W})+\mathbf{e}_1}^{\mathbf{c}_1},\ \overbrace{\mathbf{sP}+\mathbf{e}_2+\mathbf{m}\cdot\lfloor\tfrac{q}{2}\rfloor}^{\mathbf{c}_2}\Big) \in \mathbb{Z}_q^m \times \mathbb{Z}_q^m \times \mathbb{Z}_q^\lambda$$

– $\mathsf{KeyGen}(\mathsf{msk}, \mathbf{y})$*: Sample*

$$\begin{pmatrix}\mathbf{K}\\ \mathbf{R}\end{pmatrix} \leftarrow \mathsf{SamplePre}([\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}], \mathbf{I}_\ell \otimes \mathbf{T}_\mathbf{B}, \mathbf{y}^\top \otimes \mathbf{G}, \sigma_0)$$

$$\mathbf{D} = \begin{pmatrix}\mathbf{D}_0\\ \mathbf{D}_1\end{pmatrix} \leftarrow \mathsf{SamplePre}([\mathbf{B} \mid \mathbf{B}_1\mathbf{R}], \mathbf{T}_\mathbf{B}, \mathbf{P}, \sigma_0)$$

Output

$$\mathsf{sk} := (\mathbf{K}, \mathbf{R}, \mathbf{D}) \in \mathbb{Z}^{\ell m \times m} \times \mathbb{Z}^{m\times m} \times \mathbb{Z}^{2m\times\lambda}$$

– $\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk} = (\mathbf{K}, \mathbf{R}, \left(\begin{smallmatrix}\mathbf{D}_0\\ \mathbf{D}_1\end{smallmatrix}\right)), \mathbf{y}, \mathsf{ct} = (\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2), \mathbf{x})$*: Output*

$$\left\lfloor \tfrac{2}{q} \cdot \Big(\mathbf{c}_2 - [\mathbf{c}_0 \mid \mathbf{c}_1] \cdot \begin{pmatrix}\mathbf{D}_0 + (\mathbf{x}\otimes\mathbf{I}_m)\mathbf{K}\mathbf{D}_1\\ \mathbf{R}\mathbf{D}_1\end{pmatrix} \bmod q\Big) \right\rceil$$

Parameters. We will set LWE parameters as follows:

$$\begin{aligned} n &= \mathsf{poly}(\lambda, \log \ell)\\ m &= O(n \log q)\\ q &= \mathsf{poly}(\ell) \cdot \lambda^{\omega(1)} \end{aligned}$$

This yields the following parameter sizes assuming ℓ-succinct LWE:

$$|\mathsf{mpk}| = \widetilde{O}(\ell), \quad |\mathsf{ct}| = \widetilde{O}(1), \quad |\mathsf{sk}| = \widetilde{O}(\ell)$$

where $\widetilde{O}(\cdot)$ hides factors polynomial in $\lambda, \log \ell$.

Overview of analysis. Here, we have:

$$(\mathbf{I}_\ell \otimes \mathbf{B}) \cdot \mathbf{K} + \mathbf{W}\mathbf{R} = \mathbf{y}^\top \otimes \mathbf{G} \tag{22}$$

$$\mathbf{B} \cdot \mathbf{D}_0 + \mathbf{B}_1\mathbf{R} \cdot \mathbf{D}_1 = \mathbf{P} \tag{23}$$

Multiplying (22) on the left by $\mathbf{x}\otimes\mathbf{I}_n$ and on the right by $\mathbf{D}_1$ and adding to (23), and observing $(\mathbf{x}\otimes\mathbf{I}_n)(\mathbf{I}_\ell\otimes\mathbf{B}) = \mathbf{B}(\mathbf{x}\otimes\mathbf{I}_m)$ we obtain

$$[\mathbf{B} \mid \mathbf{B}_1 + (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}] \cdot \overbrace{\begin{pmatrix}\mathbf{D}_0 + (\mathbf{x}\otimes\mathbf{I}_m)\mathbf{K}\mathbf{D}_1\\ \mathbf{R}\mathbf{D}_1\end{pmatrix}}^{\text{small}} = \mathbf{P} + (\mathbf{x}\mathbf{y}^\top)\cdot\mathbf{G}\mathbf{D}_1$$

Correctness for $\mathbf{x}\mathbf{y}^\top = 0$ follows readily.

In the security proof, we sample $\mathbf{U} \leftarrow \mathscr{D}_{\mathbb{Z},\sigma_0}^{m\times m}$, and program $\mathbf{B}_1 := \mathbf{B}\mathbf{U} - (\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}$ as before. Multiplying (22) on the left by $\mathbf{x}\otimes\mathbf{I}_n$ and substituting $(\mathbf{x}\otimes\mathbf{I}_n)\mathbf{W}$ with $\mathbf{B}\mathbf{U} - \mathbf{B}_1$ yields:

$$\mathbf{B}_1\mathbf{R} = \mathbf{B} \cdot \overbrace{((\mathbf{x}\otimes\mathbf{I}_m)\mathbf{K} + \mathbf{U}\mathbf{R})}^{\text{small}} - (\mathbf{x}\mathbf{y}^\top)\mathbf{G}$$

This yields a trapdoor for $[\mathbf{B} \mid \mathbf{B}_1\mathbf{R}]$ whenever $\mathbf{x}\mathbf{y}^\top \neq 0$ (more precisely, $\gcd(\mathbf{x}\mathbf{y}^\top, q) = 1$), which allows us to simulate $\mathbf{D}$ in the corresponding secret keys (where $\mathbf{K}, \mathbf{R}$ are sampled using the trapdoor for $[\mathbf{I}_\ell \otimes \mathbf{B} \mid \mathbf{W}]$).

Theorem 6. *Under the ℓ-succinct LWE assumption, Construction 5 is selectively secure.*

Proof. We define a series of games:

- H_0: This is the real ABE security game.
- H_1: Same as H_1, except the challenger samples $\mathbf{B}_1, \mathbf{P}$ as follows:
  1. samples $\mathbf{U} \leftarrow \mathscr{D}_{\mathbb{Z},\sigma_0}^{m\times m}$, and programs $\mathbf{B}_1 := \mathbf{BU} - (\mathbf{x}\otimes \mathbf{I}_n)\mathbf{W}$
  2. samples $\mathbf{U}_0 \leftarrow \mathscr{D}_{\mathbb{Z},\sigma_0}^{m\times \lambda}$, and programs $\mathbf{P} = \mathbf{BU}_0$.

  $\mathsf{H}_0 \approx_s \mathsf{H}_1$ follows readily from [30].
- H_2: Same as H_1, except the challenger in Enc samples $\mathbf{c}_1 := \mathbf{c}_0\mathbf{U} + \mathbf{e}_1, \mathbf{c}_2 := \mathbf{c}_0\mathbf{U}_0 + \mathbf{e}_2$.

  $\mathsf{H}_1 \approx_s \mathsf{H}_2$ follows readily from noise-flooding, along with $\mathbf{c}_0\mathbf{U} \approx \mathbf{sBU} = \mathbf{s}(\mathbf{B}_1 + \mathbf{W}_{\mathbf{x}^*})$ and $\mathbf{c}_0\mathbf{U}_0 \approx \mathbf{sBU}_0 = \mathbf{sP}$.
- H_3: Same as H_2, except the challenger in KeyGen samples $\mathbf{D}$ using $\mathsf{SamplePre}([\mathbf{B} \mid \mathbf{B}_1\mathbf{R}], \binom{(\mathbf{x}\otimes \mathbf{I}_m)\mathbf{K}+\mathbf{UR}}{-\mathbf{I}_m}, \mathbf{P}, \sigma_0)$ instead of $\mathsf{SamplePre}([\mathbf{B} \mid \mathbf{B}_1\mathbf{R}], \mathbf{T}_{\mathbf{B}}, \mathbf{P}, \sigma_0)$.

  $\mathsf{H}_2 \approx_s \mathsf{H}_3$ follows from trapdoor sampling
- H_4: Same as H_3, except the challenger samples $\mathbf{c}_0 \leftarrow \mathbb{Z}_q^m$.

  $\mathsf{H}_3 \approx_c \mathsf{H}_4$ follows from (ℓ, m, σ)-succinct LWE, where we use the trapdoor for $[\mathbf{I}\otimes\mathbf{B} \mid \mathbf{W}]$ to sample $\binom{\mathbf{K}}{\mathbf{R}}$.
- H_5: Same as H_4, except the challenger samples $\mathbf{c}_2 \leftarrow \mathbb{Z}_q^\lambda$.

  $\mathsf{H}_4 \approx_s \mathsf{H}_5$ follows from left-over hash lemma, which tells us $(\mathbf{B}, \mathbf{c}_0, \mathbf{BU}_0, \mathbf{c}_0\mathbf{U}_0)$ is statistically close to uniform.

In H_5, the challenge bit b is perfectly hidden, so the advantage is 0. □

Broadcast encryption, again. Our ABE for inner product yields a broadcast encryption scheme for N users with parameters

$$|\mathsf{mpk}| = \widetilde{O}(N), \quad |\mathsf{ct}| = \widetilde{O}(1), \quad |\mathsf{sk}| = \widetilde{O}(N)$$

where $\widetilde{O}(\cdot)$ hides factors polynomial in $\lambda, \log N$.