ABE for Circuits with $\mathsf{poly}(\lambda)$-sized Keys from LWE

Valerio Cini and Hoeteck Wee

NTT Research, Sunnyvale, CA, USA

Abstract. We present a key-policy attribute-based encryption (ABE) scheme for circuits based on the Learning With Errors (LWE) assumption whose key size is independent of the circuit depth. Our result constitutes the first improvement for ABE for circuits from LWE in almost a decade, given by Gorbunov, Vaikuntanathan, and Wee (STOC 2013) and Boneh, et al. (EUROCRYPT 2014) – we reduce the key size in the latter from $\mathsf{poly}(\text{depth}, \lambda)$ to $\mathsf{poly}(\lambda)$. The starting point of our construction is a recent ABE scheme of Li, Lin, and Luo (TCC 2022), which achieves $\mathsf{poly}(\lambda)$ key size but requires pairings and generic bilinear groups in addition to LWE; we introduce new lattice techniques to eliminate the additional requirements.

1 Introduction

In key-policy attribute-based encryption (ABE) [SW05,GPSW06], ciphertexts ct are associated with an attribute $\mathbf{x} \in \{0,1\}^\ell$ and a message μ and keys sk with a predicate f, and decryption returns μ when $\mathbf{x}$ satisfies f (i.e, $f(\mathbf{x}) = 0$). 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 $\mathbf{x}$ satisfies none of these predicates. The fundamental goals in attribute-based encryption are two-fold: (i) to build expressive schemes that support a large class of policies and functions; and (ii) to obtain efficient instantiations based on widely-believed intractability of basic computational problems.

In 2013, Gorbunov, Vaikuntanathan, and Wee gave the first construction of ABE for circuits [GVW13]. The scheme relies on the Learning with Errors (LWE) assumption, and for depth d, size s circuits over ℓ-bit inputs where ℓ and d are fixed at set-up, achieves public key and ciphertext size $\ell \cdot \mathsf{poly}(d, \lambda)$ and key size $s \cdot \mathsf{poly}(d, \lambda)$. This was followed shortly by an improvement of Boneh, Gentry, Gorbunov, Halevi, Nikolaenko, Segev, Vaikuntanathan, and Vinayagamurthy, henceforth BGGHNSVV14, reducing the key size to $\mathsf{poly}(d, \lambda)$ [BGG$^+$14]. These works demonstrated for the first time the power of lattices not just for expressive computation like in fully homomorphic encryption (FHE) [Gen09,BV11], but also in conjunction with strong security properties, taking lattice and LWE-based cryptography far beyond just post-quantum security and FHE and enabling a broad range of cryptographic feasibility results over the past decade: reusable garbled circuits [GKP$^+$13], fully homomorphic signatures [GVW15b], constrained pseudorandom functions [BV15], predicate encryption [GVW15a], laconic function evaluation [QWW18,CDG$^+$17], correlation-intractable hashing and thus SNARGs for P [PS19,**?**,CJJ22], and many more.

In spite of this flurry of activities extending the BGGHNSVV14 techniques in myriad and innovative ways, there has been virtually no efficiency improvement to the BGGHNSVV14 ABE for circuits in the past decade. That is, until a recent clever and surprising work of Li, Lin, and Luo, LLL22 for short, further reducing the key size down to $\mathsf{poly}(\lambda)$ [LLL22] independent of circuit depth, but at the cost of additionally requiring pairings and generic bilinear groups. LLL22 crucially relies on techniques introduced in an earlier work of Agrawal and Yamada on optimal broadcast encryption with poly-logarithmic parameters [AY20]. Motivated by follow-up works to the latter in [AWY20,Wee22], two natural questions are whether we can replace the use of generic bilinear groups in LLL22 with either a falsifiable assumption like k-Lin, or with a post-quantum analogue such as evasive LWE [Wee22,Tsa22].

1.1 Our Results

In this work, we take a leap forward and completely eliminate the use of pairings and generic bilinear group model in LLL22. That is, we obtain an ABE for circuits with $\mathsf{poly}(\lambda)$ key size under the LWE assumption:

Reference	**Assumption**	$\lvert\mathsf{ct}\rvert$	$\lvert\mathsf{sk}\rvert$	$\lvert\mathsf{mpk}\rvert$
GVW13 [GVW13]	LWE	$\ell \cdot d^{2+1/\delta}$	$s \cdot d^{2+2/\delta}$	$\ell \cdot d^{2+2/\delta}$
BGGHNSVV14 [BGG+14]	LWE	$\ell \cdot d^{2+1/\delta}$	$d^{2+1/\delta}$	$\ell \cdot d^{2+2/\delta}$
LLL22 [LLL22]	LWE + bilinear GGM	$\ell \cdot d^{2+1/\delta}$	1	$\ell \cdot d^{2+2/\delta}$
this work	LWE	$\ell \cdot d^{2+1/\delta}$	1	$\ell \cdot d^{2+2/\delta}$

Fig. 1. Comparison with prior KP-ABE for circuits of size s and depth d, ignoring factors polynomial in λ. The quantities $\lvert\mathsf{ct}\rvert$, $\lvert\mathsf{sk}\rvert$ refer to the cryptographic overhead beyond transmitting $\mathbf{x}$ and f in the clear. In all of these schemes, the running time for encryption is $\ell \cdot \mathsf{poly}(d, \lambda)$ and that for key generation is $s \cdot \mathsf{poly}(d, \lambda)$.

Theorem 1 (informal). *Assuming the hardness of n-dimensional LWE with sub-exponential modulus-to-noise ratio 2^{n^δ}, there exists an attribute-based encryption (ABE) scheme for circuits whose key size is independent of the circuit depth. For depth d circuits over ℓ-bit inputs, we have*

$$\lvert\mathsf{mpk}\rvert = O_\lambda(\ell \cdot d^{2+2/\delta}), \quad \lvert\mathsf{ct}\rvert = O_\lambda(\ell \cdot d^{2+1/\delta}), \quad \lvert\mathsf{sk}\rvert = O_\lambda(1).$$

where $O_\lambda(\cdot)$ hides factors at most λ^3. The scheme achieves selective security against unbounded collusion.

Our security proof uses a number of new ingredients, including (i) correlated LWE secrets, (ii) pseudorandom public keys, (iii) trapdoor sampling with an approximate trapdoor. While the latter two techniques are not new to the LWE literature, this is the first time they are used to improve asymptotic efficiency in ABE. We are optimistic that our techniques will find further applications in lattice-based ABE and beyond, as has been the case for the BGGHNSVV14 techniques.

As an immediate corollary, we obtain improved reusable garbled circuits. A garbling scheme [Yao86,BHR12] allows us to encode an arbitrary circuit C into a garbled circuit $\hat{C}$, and an input x into a garbled input $\hat{x}$, so that $\hat{C}, \hat{x}$ reveal $C(x)$ and nothing else about C or x; in a reusable garbling scheme [GKP+13], security should hold even given multiple garbled inputs $\hat{x}_1, \hat{x}_2, \ldots$, etc. Instantiating the prior framework for reusable garbling schemes in [GKP+13,BGG+14] with our ABE scheme, we obtain reusable garbling schemes where the garbled circuit incurs only an *additive* $\mathsf{poly}(\lambda)$ overhead, independent of the circuit depth.

Concurrent work. A concurrent and independent work of Hsieh, Lin, and Luo [HLL23] constructs ABE for circuits where both master public key, secret keys and ciphertext have size independent of circuit depth. The scheme is proven (very) selectively secure against unbounded collusion under LWE and additionally evasive circular LWE assumptions, a stronger variant of the recently proposed evasive LWE assumption [Wee22,Tsa22].

2 Technical Overview

We begin with an overview of the LLL22 construction, which follows the following two-step blue-print (first developed in the context of pairing-based ABE):

- First, a 1-key secure ABE for circuits with key size $O_\lambda(1)$ using LWE and pairings; here-in lies the main technical and conceptual novelty of the LLL22 construction.
- Next, they achieve many-key security (i.e., security against unbounded collusion) by randomizing the secret keys "in the exponent"; this step is essentially the same as in [AY20] and crucially relies on the generic bilinear group model.

Our construction follows the same high-level blue-print. First, we construct a 1-key secure ABE for circuits with key size $O_\lambda(1)$ using just LWE: we simply replace the pairing-based building block in LLL22 with an existing LWE-based one. The main technical novelty of this work lies in carrying out the second step without pairings, and then, more notably, proving many-key security using just LWE.

2.1 BGGHNSVV14 ABE

We begin with the BGGHNSVV14 ABE scheme, which achieves $O_\lambda(d^{2+1/\delta})$-sized keys from LWE. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell \cdot m}$ be a matrix where $m = O(n \log q)$. Given $\mathbf{A}$ and a circuit $f : \{0,1\}^\ell \to \{0,1\}$ of depth d, we can derive [BGG+14,GSW13] 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

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

where $\mathbf{G}_{n,q} \in \mathbb{Z}_q^{n \times m}$ is the gadget matrix [MP12] and $\|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\| \le m^{O(d)}$.

BGGHNSVV14 ABE. The scheme is as follows, omitting error terms in the ciphertext:

$$\begin{aligned}
\mathsf{mpk} &= \mathbf{A}_0 \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{b} \leftarrow \mathbb{Z}_p^n, \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times \ell \cdot m}.\\
\mathsf{ct} &= (\overbrace{\mathbf{s} \cdot \mathbf{A}_0}^{\mathbf{c}_0},\ \overbrace{\mathbf{s} \cdot \mathbf{b}^\top + \mu \cdot \lfloor q/2 \rfloor}^{c_2},\ \overbrace{\mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})}^{\mathbf{c}_3}),\ \mathbf{s} \leftarrow \mathbb{Z}_q^n.\\
\mathsf{sk} &= \mathbf{k}_f^\top \leftarrow \mathcal{D}_{\mathbb{Z}^{2m},\tau} \text{ s.t. } [\mathbf{A}_0 \mid \mathbf{A}_f] \cdot \mathbf{k}_f^\top = \mathbf{b}^\top.
\end{aligned}$$

Decryption computes an approximation to $\mu \cdot \lfloor q/2 \rfloor$ for $f(\mathbf{x}) = 0$ as follows:

$$\mathbf{c}_2 - \overbrace{[\mathbf{c}_0 \mid \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}}]}^{\approx\, \mathbf{s} \cdot [\mathbf{A}_0 | \mathbf{A}_f]} \cdot \mathbf{k}_f^\top .$$

Parameters and efficiency. For correctness, we need the modulus-to-noise ratio to be at least $m^{O(d)}$ to accommodate the blow-up from multiplication by $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$; this means that $q \ge m^{O(d)}$ (for correctness) and $2^{n^\delta} \ge m^{O(d)}$ (for LWE hardness). In order to simulate secret keys in the security reduction, we will also require $\|\mathbf{k}_f\| \approx \tau \ge m^{O(d)}$. This means that the key size is at least $2m \log \tau \ge d^{2+1/\delta}$ bits. Indeed, it suffices to take

$$n = O_\lambda(d^{1/\delta}),\ \log q = O_\lambda(d),$$

which yields key size $O(n \cdot (\log q)^2) = O_\lambda(d^{2+1/\delta})$.

2.2 1-Key Security with $O_\lambda(1)$-sized Keys from LWE

As a warm-up to our main result, we describe a 1-key secure ABE scheme with $O_\lambda(1)$-sized keys based on LWE. The idea is to start with the LLL22 ABE scheme[1] with $O_\lambda(1)$-sized keys and (i) replace the underlying pairing-based building block with a LWE-based one, and (ii) omit secret key randomization "in the exponent", which is not needed for 1-key security.

Overview. The construction uses modulus switching [BV11,BTVW17] along with the LWE assumption over two different moduli:

- a large modulus q along with a modulus-to-noise ratio at least $m^{O(d)}$, as in BGGHNSVV14;
- a small modulus $p \le 2^{O(\lambda)}$ such that $q/p \gtrsim m^{O(d)}$.

The ciphertext contains $\mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})$ as before, except $\mathbf{s}$ is gaussian with norm smaller than p. During decryption, we will first compute $\mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{d}^\top$ where $\mathbf{d} \in \{0,1\}^m$ is a fixed low-norm vector and round that to obtain

$$\lfloor \mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{d}^\top \bmod q \rceil_p \approx \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p \bmod p,$$

[1] This refers to the many-key secure scheme in LLL22. LLL22 does not explicitly describe a 1-key secure ABE scheme with $O_\lambda(1)$ key size from LWE and pairings.

as long as $\mathbf{s}$ has sufficiently small norm. Here, we use $\lfloor \cdot \rceil_p$ to denote entry-wise rounding from $\mathbb{Z}_q$ to $\mathbb{Z}_p$.

Next, we use an inner product functional encryption (IPFE) scheme to (approximately) compute the inner product

$$\left[\, \mathbf{s} \mid \mu \,\right] \cdot \begin{bmatrix} \lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p \\ \lfloor p/2 \rfloor \end{bmatrix} = \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p + \mu \cdot \lfloor p/2 \rfloor \bmod p.$$

The ABE secret key is then an IPFE secret key $\mathbf{k}_f$ for the above computation. Concretely, using the LWE-based IPFE in [ALS16] with "constant-sized keys", the key $\mathbf{k}_f$ has size $\mathsf{poly}(\lambda, \log p) = O_\lambda(1)$.

1-key secure ABE. The scheme is as follows[2], omitting error terms in the ciphertext:

$$\begin{aligned}
\mathsf{mpk} &= \mathbf{B}_0 \leftarrow \mathbb{Z}_p^{n_0 \times O(n_0 \cdot \log p)}, \mathbf{B}_1 \leftarrow \mathbb{Z}_p^{n_0 \times O(n \cdot \log p)}, \mathbf{b} \leftarrow \mathbb{Z}_p^{n_0}, \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times \ell \cdot m}, \mathbf{d} \leftarrow \{0,1\}^m. \\
\mathsf{ct} &= (\overbrace{\mathbf{s}_0 \cdot \mathbf{B}_0}^{\mathbf{c}_0}, \overbrace{\mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n,p}}^{\mathbf{c}_1}, \overbrace{\mathbf{s}_0 \cdot \mathbf{b}^\top + \mu \cdot \lfloor p/2 \rfloor}^{c_2}, \overbrace{\mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})}^{\mathbf{c}_3}), \\
&\quad \mathbf{s}_0 \leftarrow \mathbb{Z}_p^{n_0}, \mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \chi}. \\
\mathsf{sk} &= \mathbf{k}_f^\top \leftarrow \mathcal{D}_{\mathbb{Z}^{O(n_0 \cdot \log p)}, \tau} \text{ s.t. } \mathbf{B}_0 \cdot \mathbf{k}_f^\top = \mathbf{b}^\top + \overbrace{\mathbf{B}_1 \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p)}^{\mathbf{b}_f^\top :=}.
\end{aligned}$$

Decryption computes an approximation to $\mu \cdot \lfloor p/2 \rfloor$ as follows:

$$\underbrace{\mathbf{c}_2 - \overbrace{\mathbf{c}_0 \cdot \mathbf{k}_f^\top}^{\approx \mathbf{s}_0 \cdot (\mathbf{b}^\top + \mathbf{b}_f^\top)} + \overbrace{\mathbf{c}_1 \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p)}^{\approx \mathbf{s}_0 \cdot \mathbf{b}_f^\top + \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p}}_{\approx \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p + \mu \cdot \lfloor p/2 \rfloor} - \overbrace{\lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{d}^\top \rceil_p}^{\approx \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p}.$$

Proving 1-key security. We provide a very brief sketch of "weakly selective" 1-key security, where the adversary fixes the single key query f and the challenge attribute $\mathbf{x}$ before seeing mpk. Since the security proof for our main result uses very different ideas, it is fine for the reader to skip this proof sketch. The proof strategy here is similar to that in LLL22 (along with ideas from [Agr17,QWW18,ALS16]). The reduction samples a random gaussian $\mathbf{k}_f$ and programs $\mathbf{b}$ to be $\mathbf{B}_0 \cdot \mathbf{k}_f^\top - \mathbf{b}_f^\top$. We can then use the decryption equation (plus noise flooding) to rewrite $\mathbf{c}_2$ as a function of $\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_3$ and $\mathbf{k}_f$, that is

$$\mathbf{c}_2 \approx \mathbf{c}_0 \cdot \mathbf{k}_f + \mathbf{c}_1 \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{d}^\top \rceil_p) - \lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{d}^\top \rceil_p + \mu \cdot \lfloor p/2 \rfloor$$

Using the fact that

$$\mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{d}^\top \approx \mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{d}^\top = \mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{d}^\top - f(\mathbf{x}) \cdot \mathbf{s} \cdot \mathbf{G}_{n,q} \cdot \mathbf{d}^\top$$

and that $f(\mathbf{x}) = 1$, security then boils down to showing pseudorandomness of

$$(\underbrace{\mathbf{s}_0 \cdot \mathbf{B}_0, \ \mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n,p}}_{\text{mod } p}, \ \underbrace{\mathbf{s} \cdot \mathbf{G}_{n,q} \cdot \mathbf{d}^\top, \ \mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})}_{\text{mod } q}),$$

given mpk. This follows from invoking LWE twice: first over modulus p to replace $\mathbf{s}_0 \cdot [\mathbf{B}_0 \mid \mathbf{B}_1]$ with random, and then over modulus q to replace $\mathbf{s} \cdot [\mathbf{G}_{n,q} \cdot \mathbf{d}^\top \mid \mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}]$ with random.

[2] The secret key $\mathbf{k}_f$ has a similar structure to the IPFE secret key in [ALS16], but with a different distribution.

2.3 Our ABE Scheme

In order to achieve many-key security (i.e., security against unbounded collusion), we first modify the secret key distribution, replacing $\mathbf{B}_0 \cdot \mathbf{k}_f^\top = \mathbf{b}^\top + \mathbf{b}_f^\top$ with

$$[\mathbf{B}_0 \mid \mathbf{B}_f] \cdot \mathbf{k}_f^\top = \mathbf{b}^\top, \tag{2}$$

where $\mathbf{B}_f \in \mathbb{Z}_p^{n_0 \times n_0 \log p}$ is defined analogously to $\mathbf{b}_f^\top$, except with a wider matrix $\mathbf{D} \in \{0,1\}^{m \times n_0 \log p}$ in place of $\mathbf{d}^\top \in \{0,1\}^{m \times 1}$. As in BGGHNSVV14, we will embed the gadget matrix $\mathbf{G}_{n_0,p}$ into $\mathbf{B}_f$ in the security proof. In fact, we need to make two more modifications for the security proof:

(i) $\mathbf{s}_0$ is gaussian (with norm slightly smaller than that of $\mathbf{s}$), so that we can correlate $\mathbf{s}_0$ and $\mathbf{s}$ in the security proof;

(ii) we replace (2) with

$$[\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}] \cdot \mathbf{k}_f^\top = \mathbf{b}^\top,$$

where $\mathbf{k}_f \in \mathbb{Z}^{O(n_0 \log p)}$; this allows us to simulate secret keys given just an approximate trapdoor [BTVW17], namely a low-norm $\mathbf{R}_f$ such that $\mathbf{B}_f \approx \mathbf{B}_0 \cdot \mathbf{R}_f + \mathbf{G}_{n_0,p}$ (where $\approx$ captures an error term arising from the use of $\lfloor \cdot \rceil_p$ rounding).

Setting parameters as in our 1-key secure ABE with $n_0 = O_\lambda(1), \log p = O_\lambda(1)$ (plus some minor tweaks to account for additional noise flooding), we obtain key size $O(n_0 \cdot (\log p)^2) = O_\lambda(1)$.

Our final ABE scheme. Our ABE scheme is as follows:

$$\mathsf{mpk} = \mathbf{B}_0 \leftarrow \mathbb{Z}_p^{n_0 \times O(n_0 \cdot \log p)}, \mathbf{B}_1 \leftarrow \mathbb{Z}_p^{n_0 \times O(n \cdot \log p)}, \mathbf{b} \leftarrow \mathbb{Z}_p^{n_0}, \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times \ell \cdot m}, \mathbf{D} \in \{0,1\}^{m \times O(n_0 \log p)}.$$

$$\begin{aligned} \mathsf{ct} = (&\overbrace{\mathbf{s}_0 \cdot \mathbf{B}_0}^{\mathbf{c}_0},\ \overbrace{\mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n,p}}^{\mathbf{c}_1},\ \overbrace{\mathbf{s}_0 \cdot \mathbf{b}^\top + \mu \cdot \lfloor p/2 \rfloor}^{c_2},\ \overbrace{\mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})}^{\mathbf{c}_3}), \\ &\mathbf{s}_0 \leftarrow \mathcal{D}_{\mathbb{Z},\chi_0}^{n_0}, \mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{n}. \end{aligned}$$

$$\mathsf{sk} = \mathbf{k}_f^\top \leftarrow \mathcal{D}_{\mathbb{Z}^{O(n_0 \log p)},\tau} \text{ s.t. } \big[\, \mathbf{B}_0 \mid \overbrace{\mathbf{B}_1 \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)}^{\mathbf{B}_f :=} \mid \mathbf{I}_{n_0} \,\big] \cdot \mathbf{k}_f^\top = \mathbf{b}^\top.$$

Decryption computes an approximation to $\mu \cdot \lfloor p/2 \rfloor$ as follows:

$$\mathbf{c}_2 - \overbrace{\Big[\mathbf{c}_0 \mid \mathbf{c}_1 \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p \mid \mathbf{0}_{n_0 \times n_0}\Big]}^{\approx\, \mathbf{s}_0 \cdot [\mathbf{B}_0 | \mathbf{B}_f | \mathbf{I}_{n_0}]} \cdot \mathbf{k}_f^\top.$$

Proof overview. Our proof strategy is very different from that in LLL22.

Step 1. First, we correlate $\mathbf{s}_0, \mathbf{s}$ by setting $\mathbf{s} = \mathbf{s}_0 \cdot \mathbf{W} + \mathbf{t}$ where $\mathbf{t}$ is a random gaussians and $\mathbf{W}$ a low-norm matrix. This allows us to rewrite the ciphertext as:

$$\begin{aligned} \mathsf{ct} = (&\mathbf{s}_0 \cdot \mathbf{B}_0,\ \mathbf{s}_0 \cdot (\mathbf{B}_1 + \mathbf{W} \cdot \mathbf{G}_{n,p}) + \mathbf{t} \cdot \mathbf{G}_{n,p},\ \mathbf{s}_0 \cdot \mathbf{b}^\top + \mu \cdot \lfloor p/2 \rfloor, \\ &\mathbf{s}_0 \cdot \mathbf{W} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}) + \mathbf{t} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})). \end{aligned}$$

Looking ahead, our goal is to invoke LWE to replace $\mathbf{s}_0 \cdot [\mathbf{B}_0 \mid \mathbf{b}^\top]$ with random. Towards this goal, we need to:

- account for the leakage on $\mathbf{s}_0$ in $\mathbf{s}_0 \cdot (\mathbf{B}_1 + \mathbf{W} \cdot \mathbf{G}_{n,p})$ and in $\mathbf{s}_0 \cdot \mathbf{W} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$;
- simulate key generation queries without knowing a trapdoor for $\mathbf{B}_0$.

Step 2. We replace $\mathbf{A}$ with a pseudorandom matrix such that $\mathbf{W} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}) = \mathbf{E}$ where $\mathbf{E}$ is a random gaussian matrix. This essentially follows from the LWE assumption with a low-norm $\mathbf{W}$ as the secret and $\mathbf{E}$ as the error term: concretely, we sample $\mathbf{W} = [\mathbf{I}_{n_0} \mid \mathbf{W}_0]$ and replace $\mathbf{A}$ with

$$\begin{bmatrix} \mathbf{W}_0 \cdot \tilde{\mathbf{A}} + \mathbf{E} \\ -\tilde{\mathbf{A}} \end{bmatrix} + \mathbf{x} \otimes \mathbf{G}_{n_1,q}.$$

Note that $\mathbf{t}$ completely masks any leakage about $\mathbf{W}$ in $\mathbf{s}$. Combined with noise flooding, this allows us to eliminate the term $\mathbf{s}_0 \cdot \mathbf{W} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}) = \mathbf{s}_0 \cdot \mathbf{E}$ in the challenge ciphertext.

Step 3. We program $\mathbf{B}_1 = \mathbf{B}_0 \cdot \mathbf{R} - \mathbf{W} \cdot \mathbf{G}_{n,p}$ where $\mathbf{R}$ is random low-norm. Then, we can write

$$\begin{aligned}
\mathbf{B}_f &= \mathbf{B}_1 \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) \\
&= \mathbf{B}_0 \cdot \mathbf{R} \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \mathbf{W} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p \\
&\approx \mathbf{B}_0 \cdot \mathbf{R} \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{W} \cdot \mathbf{A}_f \cdot \mathbf{D} \rceil_p \\
&= \mathbf{B}_0 \cdot \mathbf{R} \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{W} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} + f(\mathbf{x}) \cdot \mathbf{W} \cdot \mathbf{G}_{n,q} \cdot \mathbf{D} \rceil_p \\
&\approx \mathbf{B}_0 \cdot \mathbf{R} \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{E} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p - f(\mathbf{x}) \cdot \lfloor \mathbf{W} \cdot \mathbf{G}_{n,q} \cdot \mathbf{D} \rceil_p.
\end{aligned}$$

We now set $\mathbf{D}$ so that $\lfloor \mathbf{W} \cdot \mathbf{G}_{n,q} \cdot \mathbf{D} \rceil_p = \mathbf{G}_{n_0,p}$ (by setting $\mathbf{G}_{n,q} \cdot \mathbf{D} = \begin{bmatrix} \frac{q}{p} \cdot \mathbf{G}_{n_0,p} \\ \mathbf{0}_{(n-n_0) \times m_0} \end{bmatrix}$). This means that we can write

$$\mathbf{B}_f = \mathbf{B}_0 \cdot \mathbf{R}_f - \mathbf{E}_f - f(\mathbf{x}) \cdot \mathbf{G}_{n_0,p},$$

where both $\mathbf{R}_f = \mathbf{R} \cdot \mathbf{G}_{n,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)$ and $\mathbf{E}_f \approx \lfloor \mathbf{E} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p$ have low norm. Now, following [BTVW17], observe that:

$$\left[\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0} \right] \cdot \begin{bmatrix} \mathbf{R}_f \\ -\mathbf{I} \\ -\mathbf{E}_f \end{bmatrix} = f(\mathbf{x}) \cdot \mathbf{G}_{n_0,p}.$$

Since $f(\mathbf{x}) = 1$, this yields a gadget trapdoor which we can use to answer secret key queries (instead of a trapdoor for $\mathbf{B}_0$).

Step 4. At this point, we can rewrite the ciphertext as

$$\begin{aligned}
\mathsf{ct} = (&\mathbf{s}_0 \cdot \mathbf{B}_0,\ \mathbf{s}_0 \cdot \mathbf{B}_0 \cdot \mathbf{R} + \mathbf{t} \cdot \mathbf{G}_{n,p},\ \mathbf{s}_0 \cdot \mathbf{b}^\top + \mu \cdot \lfloor p/2 \rfloor, \\
&\mathbf{t} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q})).
\end{aligned}$$

We can then use the LWE assumption to replace $\mathbf{s}_0 \cdot [\mathbf{B}_0 \mid \mathbf{b}^\top]$ with random, thereby perfectly hiding the message μ.

3 Preliminaries

Notations. We use boldface lower case for row vectors (e.g. $\mathbf{r}$) and boldface upper case for matrices (e.g. $\mathbf{R}$). For integral vectors and matrices (i.e., those over $\mathbb{Z}$), we use the notation $\|\mathbf{r}\|$, $\|\mathbf{R}\|$ to denote the maximum absolute value over all the entries. We use $v \leftarrow D$ to denote a random sample from a distribution 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. We denoted by $\mathcal{D}_{\mathbb{Z}^m,\chi}$ the (centered) discrete Gaussian distribution over $\mathbb{Z}^m$ with parameter χ, i.e., the distribution over $\mathbb{Z}^m$ where for all $\mathbf{x}$, $\Pr[\mathbf{x}] \propto e^{-\pi \cdot (x_1^2 + \cdots + x_m^2)/\chi^2}$.

3.1 Attribute-based encryption

Syntax. A key policy attribute-based encryption (KP-ABE) scheme for some class $\mathcal{F}$ consists of four algorithms:

- $\mathsf{Setup}(1^\lambda, \mathcal{F}) \to (\mathsf{mpk}, \mathsf{msk})$. The setup algorithm gets as input the security parameter 1^λ and class description $\mathcal{F}$. It outputs the master public key mpk and the master secret key msk.
- $\mathsf{Enc}(\mathsf{mpk}, \mathbf{x}, \boldsymbol{\mu}) \to \mathsf{ct}_\mathbf{x}$. The encryption algorithm gets as input mpk, an input $\mathbf{x}$ and a message $\boldsymbol{\mu} \in \{0,1\}^\lambda$. It outputs a ciphertext $\mathsf{ct}_\mathbf{x}$. Note that $\mathbf{x}$ is public given $\mathsf{ct}_\mathbf{x}$.
- $\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, f) \to \mathsf{sk}_f$. The key generation algorithm gets as input $\mathsf{mpk}, \mathsf{msk}$ and $f \in \mathcal{F}$. It outputs a secret key sk_f. Note that f is public given sk_f.
- $\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}_f, f, \mathsf{ct}_\mathbf{x}, \mathbf{x}) \to \boldsymbol{\mu}$. The decryption algorithm gets as input sk_f and $\mathsf{ct}_\mathbf{x}$ along with mpk. It outputs a message $\boldsymbol{\mu}$.

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

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

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

$$\mathsf{Adv}^{\mathsf{ABE}}_{\mathcal{A}}(\lambda) := \Pr\left[b = b' : \begin{array}{l} \mathbf{x}^* \leftarrow \mathcal{A}(1^\lambda) \\ (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathcal{F}) \\ (\boldsymbol{\mu}_0, \boldsymbol{\mu}_1) \leftarrow \mathcal{A}^{\mathsf{KeyGen}(\mathsf{mpk},\mathsf{msk},\cdot)}(\mathsf{mpk}) \\ b \leftarrow \{0,1\}; \mathsf{ct}_{\mathbf{x}^*} \leftarrow \mathsf{Enc}(\mathsf{mpk}, \mathbf{x}^*, \boldsymbol{\mu}_b) \\ b' \leftarrow \mathcal{A}^{\mathsf{KeyGen}(\mathsf{mpk},\mathsf{msk},\cdot)}(\mathsf{ct}_{\mathbf{x}^*}) \end{array}\right] - \frac{1}{2},$$

with the restriction that all queries f that $\mathcal{A}$ sent to $\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, \cdot)$ satisfy $f(\mathbf{x}^*) = 1$. An ABE scheme is *selectively secure* if for all PPT adversaries $\mathcal{A}$, the advantage $\mathsf{Adv}^{\mathsf{ABE}}_{\mathcal{A}}(\lambda)$ is a negligible function in λ.

3.2 Lattices background

Learning with Errors. Given $n, m, q, \chi_e \in \mathbb{N}$, the $\mathsf{LWE}_{n,m,q,\chi_e}$ assumption states that

$$(\mathbf{A}, \mathbf{s} \cdot \mathbf{A} + \mathbf{e}) \approx_c (\mathbf{A}, \mathbf{c}),$$

where

$$\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi_e}, \mathbf{c} \leftarrow \mathbb{Z}_q^m.$$

An analog assumption holds when the secret $\mathbf{s}$ itself is chosen from a Gaussian distribution [ACPS09]. When $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^n, \chi_e}$, we refer to such assumption as the $\mathsf{sLWE}_{n,m,q,\chi_e}$ assumption.

We recall existing results from the literature useful for our construction and security proof.

Lemma 1 (Noise Flooding [GKPV10]). *Let $n \in \mathbb{N}$. For any real $\chi > \omega(\sqrt{\log n})$, and any $\mathbf{c} \in \mathbb{Z}^n$, it holds $\mathsf{SD}(\mathcal{D}_{\mathbb{Z}^n,\chi}, \mathcal{D}_{\mathbb{Z}^n,\chi} + \mathbf{c}) \leq \|\mathbf{c}\|/\chi$. In particular, if $\chi \geq \lambda^{\omega(1)} \cdot \|\mathbf{c}\|$, one has $\mathcal{D}_{\mathbb{Z}^n,\chi} \approx_s \mathcal{D}_{\mathbb{Z}^n,\chi} + \mathbf{c}$.*

Lemma 2 (Leftover Hash Lemma [HILL99]). *For $m \geq (n+1) \cdot \log q + 2 \cdot \lambda$ the distribution of $(\mathbf{A}, \mathbf{u} = \mathbf{A} \cdot \mathbf{x})$ for uniform and independent $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{x} \leftarrow \{0,1\}^m$ is statistically indistinguishable from uniformly random.*

Lemma 3 (Norm of Gaussian [GPV08]). *For all $\chi \geq 1$ and $k \geq 2$,*

$$\Pr[|x| \geq \chi\sqrt{k} \mid x \leftarrow \mathcal{D}_{\mathbb{Z},\chi}] \leq 2^{-k}.$$

Therefore, if $\mathbf{H} \leftarrow \mathcal{D}_{\mathbb{Z}^{m \times m'},\chi}$ then $\|\mathbf{H}\| \leq \lambda\chi\sqrt{m \cdot m'}$ except that with probability negligible in λ.

Trapdoor and preimage sampling [MP12,GPV08]. Let $n, q \in \mathbb{Z}$,

$$\mathbf{g}_q = (1, 2, 4, \ldots, 2^{\lceil \log q \rceil - 1}) \in \mathbb{Z}^{\lceil \log q \rceil}.$$

The gadget matrix $\mathbf{G}_{n,q}$ is defined as the diagonal concatenation of $\mathbf{g}_q$ n times. Formally, $\mathbf{G}_{n,q} = \mathbf{g}_q \otimes \mathbf{I}_n \in \mathbb{Z}^{n \times n \cdot \lceil \log q \rceil}$. For any $t \in \mathbb{Z}$, the function $\mathbf{G}_{n,q}^{-1} : \mathbb{Z}_q^{n \times t} \to \{0,1\}^{n \cdot \lceil \log q \rceil \times t}$ expands each entry $a \in \mathbb{Z}_q$ of the input matrix into a column of size $\lceil \log q \rceil$ consisting of the bit-representation of a. For any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times t}$ it holds that $\mathbf{G}_{n,q} \cdot \mathbf{G}_{n,q}^{-1}(\mathbf{A}) = \mathbf{A} \bmod q$. We refer to the gadget matrix simply as $\mathbf{G}$ when parameters n and q are clear from the context.

Let $n, m, q \in \mathbb{N}$ and consider a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. For all $\mathbf{V} \in \mathbb{Z}_q^{n \times m'}$ we let $\mathbf{A}^{-1}(\mathbf{V}, \tau)$ denote the random variable whose distribution is the Discrete Gaussian $\mathcal{D}_{\mathbb{Z}^{m \times m'}, \tau}$ conditioned on $\mathbf{A} \cdot \mathbf{A}^{-1}(\mathbf{V}, \tau) = \mathbf{V} \bmod q$.

Lemma 4 (Trapdoor Generation and Sampling [Ajt96,GPV08,MP12]). *There exists a pair of probabilistic polynomial-time algorithms:*

- SamplePre$(\mathbf{A}, \mathbf{T}, \mathbf{V}, \tau)$ *that given* $\mathbf{A}$ *and any* $\mathbf{T}$ *such that* $\mathbf{A} \cdot \mathbf{T} = \mathbf{G}$, $\tau \geq 2 \cdot m \cdot \sqrt{n \cdot \log q} \cdot \|\mathbf{T}\|$ *and* $\mathbf{V} \in \mathbb{Z}_q^{n \times m'}$, *outputs a sample from* $\mathbf{A}^{-1}(\mathbf{V}, \tau)$. *We call* $\mathbf{T}$ *a* τ*-trapdoor for* $\mathbf{A}$.
- TrapGen$(1^n, q, m)$ *that for all* $m \geq m_0 = m_0(n, q) = O(n \log q)$, *outputs* $(\mathbf{A}, \mathbf{T}_{\mathbf{A}})$ *s.t.* $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ *is within statistical distance* 2^{-n} *from uniform and* $\mathbf{T}_{\mathbf{A}}$ *is a* τ*-trapdoor for* $\mathbf{A}$ *where* $\tau = O(\sqrt{n} \cdot \log q \cdot \log n)$.

Homomorphic Computation on Matrices. We recall basic homomorphic computation on matrices used in BGGHNSVV14 as captured in (1):

Theorem 2 ([BGG+14,GSW13]). *There exist efficient deterministic algorithms* EvalF *and* EvalFX *such that for all* $n, q, \ell \in \mathbb{N}$, *and for any sequence of matrices* $\mathbf{A} = (\mathbf{A}_1, \ldots, \mathbf{A}_\ell) \in (\mathbb{Z}^{n \times n \cdot \lceil \log q \rceil})^\ell$, *for any depth-*$d$ *Boolean circuit* $f : \{0,1\}^\ell \to \{0,1\}$ *and for every* $\mathbf{x} = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$,[3] *the following properties hold.*

- *The outputs* $\mathbf{A}_f =$ EvalF$(\mathbf{A}, f)$ *and* $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} =$ EvalFX$(\mathbf{A}, f, \mathbf{x})$ *are matrices in* $\mathbb{Z}_q^{n \times (n \cdot \lceil \log q \rceil)}$ *and* $\mathbb{Z}^{(\ell \cdot n \cdot \lceil \log q \rceil) \times (n \cdot \lceil \log q \rceil)}$,
- *It holds that* $\|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\| \leq (n \cdot \log q)^{O(d)}$,
- *It holds that*

$$(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(\mathbf{x}) \cdot \mathbf{G}_{n,q} \mod q.$$

We will call this the "key equation" for matrix evaluation.

For a proof of this theorem, we refer the reader to [BCTW16, Fact 3.4].

3.3 Rounding Function and Properties

We define a "rounding" function $\lfloor \cdot \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$, where $q \geq p \geq 2$, as

$$\lfloor x \rceil_p = \lfloor (p/q) \cdot \bar{x} \rceil \bmod p,$$

where $\bar{x} \in \mathbb{Z}$ denote an arbitrary integer congruent to x modulo q. We extend $\lfloor \cdot \rceil_p$ component-wise to vectors and matrices over $\mathbb{Z}_q$. It can be seen that the definition of $\lfloor \cdot \rceil_p$ is independent of the choice of $\bar{x}$. Indeed, if $\bar{x}' = \bar{x} + k \cdot q \in \mathbb{Z}$, for $k \in \mathbb{Z}$, i.e., $\bar{x}' = \bar{x} \bmod q$, then one has

$$\lfloor (p/q) \cdot \bar{x}' \rceil = \lfloor (p/q) \cdot (\bar{x} + k \cdot q) \rceil = \lfloor (p/q) \cdot \bar{x} + \underbrace{p \cdot k}_{\in \mathbb{Z}} \rceil = \lfloor (p/q) \cdot \bar{x} \rceil + p \cdot k = \lfloor (p/q) \cdot \bar{x} \rceil \mod p,$$

as claimed. Next, we state some properties of the rounding function used in [LLL22].

Lemma 5 (Properties of $\lfloor \cdot \rceil_p$).

1. *For any* $\mathbf{a}, \mathbf{b} \in \mathbb{Z}_q^m$, $\exists \mathbf{e} \in \mathbb{Z}^m$ *with* $\|\mathbf{e}\| \leq 2$ *such that* $\lfloor \mathbf{a} + \mathbf{b} \bmod q \rceil_p = \lfloor \mathbf{a} \rceil_p + \lfloor \mathbf{b} \rceil_p + \mathbf{e} \mod p$.

[3] In order to support NOT gates, an additional constant 1 should be actually concatenated to the attribute $\mathbf{x}$.

2. *For any* $\mathbf{u} \in \mathbb{Z}^m, \mathbf{V} \in \mathbb{Z}_q^{m \times k}$, $\exists \mathbf{e} \in \mathbb{Z}^k$ *with* $\|\mathbf{e}\| \leq m \cdot \|\mathbf{u}\|$ *such that* $\lfloor \mathbf{u} \cdot \mathbf{V} \bmod q \rceil_p = \mathbf{u} \cdot \lfloor \mathbf{V} \rceil_p + \mathbf{e} \mod p$.

Proof. Property (1) follows from (2) by taking $\mathbf{u} = [1, 1]$, and $\mathbf{V} = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}$. It remains to prove property (2). Let $\bar{\mathbf{V}}$ denote the integer vector whose entry equals that of $\mathbf{V}$. We have

$$\begin{aligned} \lfloor \mathbf{u} \cdot \mathbf{V} \bmod q \rceil_p &= \lfloor (p/q) \cdot \mathbf{u} \cdot \bar{\mathbf{V}} \rceil \\ &= (p/q) \cdot \mathbf{u} \cdot \bar{\mathbf{V}} + \mathbf{f} \mod p, \end{aligned}$$

where $\mathbf{f} \in \mathbb{Q}^k$ with $\|\mathbf{f}\| \leq 1/2$. On the other hand

$$\begin{aligned} \mathbf{u} \cdot \lfloor \mathbf{V} \rceil_p &= \mathbf{u} \cdot \lfloor (p/q) \cdot \bar{\mathbf{V}} \rceil \\ &= \mathbf{u} \cdot \big((p/q) \cdot \bar{\mathbf{V}} + \mathbf{F} \big) \\ &= (p/q) \cdot \mathbf{u} \cdot \bar{\mathbf{V}} + \mathbf{u} \cdot \mathbf{F} \mod p, \end{aligned}$$

where $\mathbf{F} \in \mathbb{Q}^{m \times k}$, with $\|\mathbf{F}\| \leq 1/2$. We obtain that

$$\begin{aligned} \lfloor \mathbf{u} \cdot \mathbf{V} \bmod q \rceil_p &= (p/q) \cdot \mathbf{u} \cdot \bar{\mathbf{V}} + \mathbf{f} \\ &= (p/q) \cdot \mathbf{u} \cdot \bar{\mathbf{V}} + \mathbf{f} + \Big(\mathbf{u} \cdot \lfloor \mathbf{V} \rceil_p - (p/q) \cdot \mathbf{u} \cdot \bar{\mathbf{V}} + \mathbf{u} \cdot \mathbf{F} \Big) \\ &= \mathbf{u} \cdot \lfloor \mathbf{V} \rceil_p + \underbrace{\mathbf{f} - \mathbf{u} \cdot \mathbf{F}}_{\mathbf{e}} \mod p. \end{aligned}$$

Observe that $\mathbf{e}$ is an integer vector since it is the difference of two integer vectors modulo p. Next, we bound $\|\mathbf{e}\|$:

$$\begin{aligned} \|\mathbf{e}\| &= \|\mathbf{f} - \mathbf{u} \cdot \mathbf{F}\| \\ &\leq \|\mathbf{f}\| + \|\mathbf{u} \cdot \mathbf{F}\| \\ &\leq 1/2 + (1/2) \cdot m \cdot \|\mathbf{u}\| \\ &\leq m \cdot \|\mathbf{u}\|, \end{aligned}$$

where in the last inequality we have used that $\|\mathbf{u}\| \geq 1$, as otherwise $\mathbf{u} = 0$, and the claim is trivial. □

4 ABE for Circuits with $O_\lambda(1)$-sized Keys

In this section, we prove Theorem 1. We refer to Section 2 for an overview of the scheme[4] and the security proof.

Construction. Let the ABE $\Pi = (\mathsf{Setup}, \mathsf{Enc}, \mathsf{KeyGen}, \mathsf{Dec})$ for the family $\mathcal{F}_{\ell,d,s}$ of circuits of depth d and size s over ℓ-bit inputs be defined as follows:

- $\mathsf{Setup}(1^\lambda, 1^d, 1^\ell)$: Sample

$$(\mathbf{B}_0, \mathbf{T}_{\mathbf{B}_0}) \leftarrow \mathsf{TrapGen}(1^{n_0}, 1^{m_0}, p), \mathbf{B}_1 \leftarrow \mathbb{Z}_p^{n_0 \times m_1}, \mathbf{B}_2 \leftarrow \mathbb{Z}_p^{n_0 \times \lambda}, \mathbf{A} \leftarrow \mathbb{Z}_q^{n_1 \times \ell \cdot m_2},$$

  and let $\mathbf{D} := \mathbf{G}_{n_1,q}^{-1} \left(\begin{bmatrix} \frac{q}{p} \cdot \mathbf{G}_{n_0,p} \\ \mathbf{0}_{(n_1 - n_0) \times m_0} \end{bmatrix} \right) \in \mathbb{Z}^{m_2 \times m_0}$.
  Set $\mathsf{mpk} = (\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2, \mathbf{A}, \mathbf{D})$, and $\mathsf{msk} = \mathbf{T}_{\mathbf{B}_0}$. Return $(\mathsf{mpk}, \mathsf{msk})$.

[4] The parameters (n_1, m_2) here correspond to (n, m) in the overview, and we replaced $\mathbf{b}_2^\top$ with a matrix $\mathbf{B}_2$ so that we can directly encrypt messages in $\{0, 1\}^\lambda$.

- $\mathsf{Enc}(\mathsf{mpk}, \mathbf{x}, \boldsymbol{\mu} \in \{0,1\}^\lambda)$: Sample

$$\mathbf{s}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0},\chi_0}, \mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_1},\chi_3}, \mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_0},\chi_0}, \mathbf{e}_1 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_1},\chi_2}, \mathbf{e}_2 \leftarrow \mathcal{D}_{\mathbb{Z}^{\lambda},\chi_0}, \mathbf{e}_3 \leftarrow \mathcal{D}_{\mathbb{Z}^{\ell \cdot m_2},\chi_3}.$$

  Compute

$$\begin{aligned} \mathbf{c}_0 &:= \mathbf{s}_0 \cdot \mathbf{B}_0 + \mathbf{e}_0 && \bmod p, \\ \mathbf{c}_1 &:= \mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1 && \bmod p, \\ \mathbf{c}_2 &:= \mathbf{s}_0 \cdot \mathbf{B}_2 + \boldsymbol{\mu} \cdot \lfloor p/2 \rfloor + \mathbf{e}_2 && \bmod p, \\ \mathbf{c}_3 &:= \mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}_3 && \bmod q. \end{aligned}$$

  Output $\mathsf{ct}_{\mathbf{x}} := (\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$.

- $\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, f)$: Compute $\mathbf{A}_f = \mathsf{EvalF}(\mathbf{A}, f)$ and

$$\mathbf{B}_f := \mathbf{B}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) \in \mathbb{Z}_p^{n_0 \times m_0}.$$

  Sample

$$\mathbf{K}_f \leftarrow \mathsf{SamplePre}([\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}], \begin{bmatrix} \mathbf{T}_{\mathbf{B}_0} \\ \mathbf{0}_{m_0 \times m_0} \\ \mathbf{0}_{n_0 \times m_0} \end{bmatrix}, \mathbf{B}_2, \tau),$$

  i.e., $[\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}] \cdot \mathbf{K}_f = \mathbf{B}_2$. Output $\mathsf{sk}_f = (\mathbf{K}_f)$.[5]

- $\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}_f, f, \mathsf{ct}_{\mathbf{x}}, \mathbf{x})$: Compute $\mathbf{A}_f = \mathsf{EvalF}(\mathbf{A}, f)$ and $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x})$. Parse $\mathsf{ct}_{\mathbf{x}} = (\mathbf{c}_0, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$. Output

$$\left\lfloor \tfrac{2}{p} \cdot \left(\mathbf{c}_2 - \left[\mathbf{c}_0 \mid \mathbf{c}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p \mid \mathbf{0}_{n_0 \times n_0} \right] \cdot \mathbf{K}_f \right) \right\rceil \in \{0,1\}^\lambda. \tag{3}$$

Parameters. We have 4 gaussian parameters:

$$\overbrace{\chi_0}^{\approx\|\mathbf{e}_0\|,\|\mathbf{s}_0\|,\|\mathbf{e}_2\|} \leq \overbrace{\chi_1}^{\approx\|\mathbf{W}\|,\|\mathbf{E}\|} \leq \overbrace{\chi_2}^{\approx\|\mathbf{e}_1\|} \leq \overbrace{\chi_3}^{\approx\|\mathbf{s}\|,\|\mathbf{e}_3\|}.$$

where $\mathbf{W}, \mathbf{E}$ are introduced in the security proof.

The parameters requirements can be compactly specified as:

$$\begin{aligned} m_0 &\geq O(n_0 \log p) && \text{trapdoor generation (Lemma 4)} \\ 2^{n_0^\delta} &\geq p/\chi_0, \qquad \chi_0 \geq O(n_0 + \lambda) && \mathsf{sLWE}_{n_0,p,\chi_0} \text{ hardness } (\mathsf{H}_6 \approx_c \mathsf{H}_7) \\ 2^{(n_1-n_0)^\delta} &\geq q/\chi_1, \qquad \chi_1 \geq O(n_1 + \lambda) && \mathsf{sLWE}_{n_1-n_0,q,\chi_1} \text{ hardness } (\mathsf{H}_0 \approx_c \mathsf{H}_1) \\ \chi_3 &\geq \chi_0 \cdot \chi_1 \cdot \mathsf{poly}(\ell, m_2, \lambda) \cdot \lambda^{\omega(1)} && \text{noise flooding } (\mathsf{H}_1 \approx_s \mathsf{H}_2, \mathsf{H}_5 \approx_s \mathsf{H}_6) \\ \chi_2 &\geq \chi_0 \cdot \mathsf{poly}(m_2, \lambda) \cdot \lambda^{\omega(1)} && \text{noise flooding } (\mathsf{H}_5 \approx_s \mathsf{H}_6) \\ m_0 &\geq (n_0 + 1) \cdot \log p + 2\lambda && \text{LHL } (\mathsf{H}_2 \approx_s \mathsf{H}_3) \\ \tau &\geq \mathsf{poly}(m_2, \ell, \lambda) \cdot p/q \cdot B \cdot \chi_1 && \text{trapdoor generation } (\mathsf{H}_3 \approx_s \mathsf{H}_4) \\ p &\geq \mathsf{poly}(\lambda) \cdot \chi_0 + (p/q \cdot B \cdot \chi_3 + \chi_2 + \chi_3) \cdot \mathsf{poly}(m_2, \ell, \lambda) \cdot \tau && \text{correctness} \end{aligned}$$

We bound the adversarially chosen parameters d, ℓ by $\lambda^{\omega(1)}$.[6] Taking $\lambda_1 = \lambda^{\omega(1)}$, and additionally bounding each of $p/q \cdot B$ and the $\mathsf{poly}(\ell, m_2, \lambda)$ terms by λ_1, we can set

[5] Note that it suffices for correctness to output $\mathsf{sk}_f = (\mathbf{K}_f^0)$, where $\mathbf{K}_f^0 \in \mathbb{Z}^{2m_0 \times \lambda}$ correspond to the top $2m_0$ rows of $\mathbf{K}_f$, so that $[\mathbf{B}_0 \mid \mathbf{B}_f] \cdot \mathbf{K}_f^0 \approx \mathbf{B}_2$. Decryption would then compute $\mathbf{c}_2 - \left[\mathbf{c}_0 \mid \mathbf{c}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p\right] \cdot \mathbf{K}_f^0$ in place of (3).

[6] This is also the case in [LLL22]. In particular, the key sizes require multiplicative factors in $\mathsf{poly}(\log d, \log \ell)$.

$$\begin{array}{lll}
m_0 = n_0 \cdot \lceil \log p \rceil, & m_1 = n_1 \cdot \lceil \log p \rceil, & m_2 = n_1 \cdot \lceil \log q \rceil, \\
\chi_0 = \chi_1 = \lambda_1, & \chi_2 = \lambda_1^3, & \chi_3 = \lambda_1^4, \\
\tau = \lambda_1^3, & & \\
p = \lambda_1^9, & n_0 = O(\log \lambda_1)^{1/\delta}, & \\
q = B \cdot \lambda_1^8 = \lambda_1^{O(d)}, & n_1 = O(\log B + \log \lambda_1)^{1/\delta} = O(d \cdot \log \lambda_1)^{1/\delta}, &
\end{array} \tag{4}$$

where in the last line, we use $B := m_2^{O(d)} \leq \lambda_1^{O(d)}$.

Efficiency. Our ABE scheme achieves

$$|\mathsf{mpk}| = O(\ell \cdot (n_1 \cdot \log q)^2), \quad |\mathsf{ct}| = O(\ell \cdot n_1 \cdot (\log q)^2 + \lambda \cdot \log p), \quad |\mathsf{sk}| = O(n_0 \cdot \lambda \cdot (\log p)^2).$$

This yields the following parameter sizes (in bits) for our ABE scheme:

$$|\mathsf{mpk}| = O_\lambda(\ell \cdot d^{2+2/\delta}), \quad |\mathsf{ct}| = O_\lambda(\ell \cdot d^{2+1/\delta}), \quad |\mathsf{sk}| = O_\lambda(1).$$

where $O_\lambda(\cdot)$ hides factors polynomial in λ (bounded by λ^3). Here, we use $n_0 = O(\lambda), n_1 = O(d^{1/\delta} \cdot \lambda), \log p = O(\lambda), \log q = O(d \cdot \lambda)$, where we do optimize on the dependency on d but not on λ. In comparison, the BGGHNSVV14 ABE achieves the same $|\mathsf{mpk}|, |\mathsf{ct}|$ but with $|\mathsf{sk}| = O_\lambda(d^{2+1/\delta})$.

The running time for encryption is $\ell \cdot \mathsf{poly}(d, \lambda)$ and that for key generation is $s \cdot \mathsf{poly}(d, \lambda)$ where s is the circuit size for f.

Theorem 3 (Correctness). *Let Π be the* KP-ABE *construction described above, with parameters set as in Eq. (4). Then, Π is correct.*

Proof. Fix $\mathbf{x}, f$ such that $f(\mathbf{x}) = 0$. The bulk of the proof lies in showing that

$$\begin{aligned}
&\left[\mathbf{c}_0 \;\middle|\; \mathbf{c}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p \;\middle|\; \mathbf{0}_{n_0 \times n_0}\right] \\
&\qquad = \mathbf{s}_0 \cdot \left[\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}\right] + \mathbf{e}'_{f,\mathbf{x}} \bmod p
\end{aligned} \tag{5}$$

where $\|\mathbf{e}'_{f,\mathbf{x}}\|$ is small. Correctness then follows readily from the fact that

$$\mathbf{c}_2 - (\mathbf{s}_0 \cdot [\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}] + \mathbf{e}'_{f,\mathbf{x}}) \cdot \mathbf{K}_f = \boldsymbol{\mu} \cdot \lfloor p/2 \rfloor + \mathbf{e}_2 - \mathbf{e}'_{f,\mathbf{x}} \cdot \mathbf{K}_f \bmod p$$

To prove Eq. (5):

- First, we have that

$$\begin{aligned}
\mathbf{c}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) &= (\mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1) \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) \\
&= \mathbf{s}_0 \cdot \mathbf{B}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) + \mathbf{s} \cdot \mathbf{G}_{n_1,p} \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) \\
&\qquad + \boxed{\mathbf{e}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)} \\
&\approx \mathbf{s}_0 \cdot \mathbf{B}_1 \cdot \mathbf{G}_{n_1,p}^{-1}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) + \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p \quad \bmod p.
\end{aligned}$$

- Second, using the key equation

$$(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n_1,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f \bmod q,$$

as $f(\mathbf{x}) = 0$, we have

$$\begin{aligned}
\lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p &= \lfloor (\mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}_3) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p \\
&= \lfloor \mathbf{s} \cdot (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n_1,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} + \mathbf{e}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p \\
&= \lfloor \mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{D} + \mathbf{e}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \bmod q \rceil_p \\
&= \lfloor \mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{D} \rceil_p + \boxed{\lfloor \mathbf{e}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p} + \boxed{\varepsilon} \\
&\approx \lfloor \mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{D} \rceil_p \\
&= \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p + \boxed{\mathbf{e}_\mathbf{s}} \\
&\approx \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p \bmod p,
\end{aligned}$$

where the fifth equality uses that $\lfloor \mathbf{s} \cdot \mathbf{A}_f \cdot \mathbf{D} \rceil_p = \mathbf{s} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p + \mathbf{e}_\mathbf{s}$, with $\|\mathbf{e}_\mathbf{s}\| \leq n_1 \cdot \|\mathbf{s}\|$.

- We deduce that

$$\mathbf{c}_1 \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{c}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p \approx \mathbf{s}_0 \cdot \overbrace{\mathbf{B}_1 \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)}^{=\mathbf{B}_f} \,\mathrm{mod} p.$$

- We conclude that Eq. (5) holds with

$$\mathbf{e}'_{f,\mathbf{x}} = [\mathbf{e}_0 \mid \mathbf{e}_1 \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) - \lfloor \mathbf{e}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p - \varepsilon - \mathbf{e}_\mathbf{s} \mid -\mathbf{s}_0],$$

where, using Lemma 3, we have that with overwhelming probability in λ

$$\begin{aligned}
\|\mathbf{e}'_{f,\mathbf{x}}\| \leq\ & \lambda \cdot \sqrt{m_0} \cdot \chi_0 && \|\mathbf{e}_0\| \\
& + m_1 \cdot \lambda \cdot \sqrt{m_0} \cdot \chi_2 && \|\mathbf{e}_1 \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)\| \\
& + \frac{p}{q} \cdot \ell \cdot m_2 \cdot m_2 \cdot \lambda \cdot \sqrt{\ell \cdot m_2} \cdot \chi_3 \cdot m_2^{O(d)} && \|\lfloor \mathbf{e}_3 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{D} \rceil_p\| \\
& + 2 && \|\varepsilon\| \\
& + n_1 \cdot \lambda \cdot \sqrt{n_1} \cdot \chi_3 && \|\mathbf{e}_\mathbf{s}\| \\
& + \lambda \cdot \sqrt{n_0} \cdot \chi_0 && \|\mathbf{s}_0\|
\end{aligned}$$

It follows that the norm of the final error term is, with overwhelming probability in λ, bounded by

$$\begin{aligned}
\|\mathbf{e}_2\| + \|\mathbf{e}'_{f,\mathbf{x}} \cdot \mathbf{K}_f\| \leq\ & \lambda \cdot \sqrt{\lambda} \cdot \chi_0 \\
& + \lambda \cdot (2 \cdot m_0 + n_0)^2 \cdot \tau \cdot \Big(\lambda \cdot \sqrt{m_0} \cdot \chi_0 \\
& \qquad + m_1 \cdot \lambda \cdot \sqrt{m_0} \cdot \chi_2 \\
& \qquad + \frac{p}{q} \cdot \ell \cdot m_2 \cdot m_2 \cdot \lambda \cdot \sqrt{\ell \cdot m_2} \cdot \chi_3 \cdot m_2^{O(d)} \\
& \qquad + 2 \\
& \qquad + n_1 \cdot \lambda \cdot \sqrt{n_1} \cdot \chi_3 \\
& \qquad + \lambda \cdot \sqrt{n_0} \cdot \chi_0 \Big),
\end{aligned}$$

where we have used that $\|\mathbf{K}_f\| \leq \lambda \cdot \sqrt{(2 \cdot m_0 + n_0) \cdot n_0} \cdot \tau$ and that $\mathbf{e}'_{f,\mathbf{x}}$ is a vector of length $2 \cdot m_0 + n_0$. Since

$$p \geq \mathsf{poly}(\lambda) \cdot \chi_0 + (p/q \cdot B \cdot \chi_3 + \chi_2 + \chi_3) \cdot \mathsf{poly}(m_2, \ell, \lambda) \cdot \tau,$$

the theorem follows. □

Hybrid	mpk	ct	trapdoor for sk_f	justification
H_0	$(\mathbf{B}_0, \mathbf{T}_{\mathbf{B}_0}) \leftarrow \mathsf{TrapGen}(1^{n_0}, 1^{m_0}, p)$ $\mathbf{B}_1 \leftarrow \mathbb{Z}_p^{n_0 \times m_1}$ $\mathbf{B}_2 \leftarrow \mathbb{Z}_p^{n_0 \times \lambda}$ $\mathbf{A} \leftarrow \mathbb{Z}_q^{n_1 \times (\ell \cdot m_2)}$	$\mathbf{c}_0 \approx \mathbf{s}_0 \cdot \mathbf{B}_0$ $\mathbf{c}_1 \approx \mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n_1,p}$ $\mathbf{c}_2 \approx \mathbf{s}_0 \cdot \mathbf{B}_2 + \boldsymbol{\mu} \cdot \lfloor p/2 \rfloor$ $\mathbf{c}_3 \approx \mathbf{s} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q})$	$\mathbf{T}_{\mathbf{B}_0}$	
H_1	$\mathbf{A} = \begin{bmatrix} \mathbf{W}_0 \cdot \tilde{\mathbf{A}} + \mathbf{E} \\ -\tilde{\mathbf{A}} \end{bmatrix} + \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}$	$\downarrow$	$\downarrow$	$\mathsf{sLWE}_{n_1-n_0,\ell\cdot m_2,\chi_1,q}$
H_2	$\downarrow$	$\mathbf{s} = \mathbf{s}_0 \cdot \mathbf{W} + \mathbf{t}$	$\downarrow$	noise flooding (Lemma 1)
H_3	$\mathbf{B}_1 = \mathbf{B}_0 \cdot \mathbf{R} - \mathbf{W} \cdot \mathbf{G}_{n_1,p}$	$\downarrow$	$\downarrow$	LHL (Lemma 2)
H_4	$\downarrow$	$\downarrow$	$\begin{bmatrix} \mathbf{R}_f \\ -\mathbf{I}_{m_0} \\ -\mathbf{E}_f \end{bmatrix}$	SamplePre (Lemma 4)
H_5	$\mathbf{B}_0 \leftarrow \mathbb{Z}_p^{n_0 \times m_0}$	$\downarrow$	$\downarrow$	TrapGen (Lemma 4)
H_6	$\downarrow$	$\mathbf{c}_1 = \mathbf{c}_0 \cdot \mathbf{R} + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1'$ $\mathbf{c}_3 = \mathbf{t} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}_3'$	$\downarrow$	noise flooding (Lemma 1)
H_7	$\downarrow$	$\mathbf{c}_0 \leftarrow \mathbb{Z}_p^{m_0}$, $\mathbf{c}_2 \leftarrow \mathbb{Z}_p^{\lambda}$	$\downarrow$	$\mathsf{sLWE}_{n_0,m_0+\lambda,\chi_0,p}$

Fig. 2. Summary of our security proof. $\downarrow$ denotes the same as the previous hybrid. We omit the noise terms in H_0. In H_1, we introduce $\mathbf{W} = [\mathbf{I}_{n_0} \mid \mathbf{W}_0]$.

Theorem 4 (Security). *Let Π be the* KP-ABE *construction described above, with parameters set as in Eq. (4). Then, Π is selectively secure assuming* $\mathsf{sLWE}_{n_0,m_0+\lambda,\chi_0,p}$ *and* $\mathsf{sLWE}_{n_1-n_0,\ell\cdot m_2,\chi_1,q}$.

Proof. Consider the following sequence of hybrids, summarized in Fig. 2:

- H_0: This is identical to the real security game.
- H_1: This is the same as H_0, except for the following modification to $\mathbf{A}$ in mpk:
  - let $\mathbf{x}^*$ the challenge attribute provided by the adversary,
  - sample $\mathbf{W}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0 \times (n_1-n_0)},\chi_1}$ and set $\mathbf{W} = [\mathbf{I}_{n_0} \mid \mathbf{W}_0]$,
  - sample $\tilde{\mathbf{A}} \leftarrow \mathbb{Z}_q^{(n_1-n_0)\times \ell \cdot m_2}$ and $\mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0 \times \ell \cdot m_2},\chi_1}$,
  - set
  $$\mathbf{A} := \begin{bmatrix} \mathbf{W}_0 \cdot \tilde{\mathbf{A}} + \mathbf{E} \\ -\tilde{\mathbf{A}} \end{bmatrix} + \mathbf{x}^* \otimes \mathbf{G}_{n_1,q} \mod q.$$

  By $\mathsf{sLWE}_{n_1-n_0,\ell\cdot m_2,\chi_1,q}$, the distribution of $\mathbf{A}$ in H_1 is computationally indistinguishable to the one in H_0. The reduction works as follows:
  - it receives the sLWE instance $(\mathbf{A}_{\mathsf{bot}}, \mathbf{A}_{\mathsf{top}}) \in \mathbb{Z}_q^{(n_1-n_0)\times \ell \cdot m_2} \times \mathbb{Z}_q^{n_0 \times \ell \cdot m_2}$,
  - it obtains $\mathbf{x}^*$ from the adversary $\mathcal{A}$,
  - computes $\mathbf{B}_0$, $\mathbf{B}_1$, and $\mathbf{B}_2$ as in H_0, and sets
  $$\mathbf{A} := \begin{bmatrix} \mathbf{A}_{\mathsf{top}} \\ -\mathbf{A}_{\mathsf{bot}} \end{bmatrix} + \mathbf{x}^* \otimes \mathbf{G}_{n_1,q} \mod q.$$
  - it answers KeyGen queries using $\mathbf{T}_{\mathbf{B}_0}$,
  - whenever the adversary $\mathcal{A}$ produces $(\boldsymbol{\mu}_0, \boldsymbol{\mu}_1)$, it samples $b \leftarrow \{0,1\}$, $\mathbf{s}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0},\chi_0}$, $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_1},\chi_3}$, $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_0},\chi_0}$, $\mathbf{e}_1 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_1},\chi_2}$, $\mathbf{e}_2 \leftarrow \mathcal{D}_{\mathbb{Z}^{\lambda},\chi_0}$, $\mathbf{e}_3 \leftarrow \mathcal{D}_{\mathbb{Z}^{\ell\cdot m_2},\chi_3}$, and outputs
  $$\mathsf{ct} = (\mathbf{s}_0 \cdot \mathbf{B}_0 + \mathbf{e}_0, \mathbf{s}_0 \cdot \mathbf{B}_1 + \mathbf{s} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1, \mathbf{s}_0 \cdot \mathbf{B}_2 + \boldsymbol{\mu}_b \cdot \lfloor p/2 \rfloor + \mathbf{e}_2, \mathbf{s} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}_3).$$

  Observe that

- if $(\mathbf{A}_{\mathsf{bot}}, \mathbf{A}_{\mathsf{top}})$ is a uniform random instance, the view of $\mathcal{A}$ is identical to H_0;
- if $(\mathbf{A}_{\mathsf{bot}}, \mathbf{A}_{\mathsf{top}})$ is a structured $\mathsf{sLWE}_{n_1-n_0,\ell\cdot m_2,\chi_1,q}$ instance, the view of the adversary $\mathcal{A}$ is identical to H_1.

We conclude that $\mathsf{H}_0 \approx_c \mathsf{H}_1$.

– H_2: This is the same as H_1, except for the following modification to $\mathbf{s}$ in the challenge ciphertext:

- sample $\mathbf{t} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_1},\chi_3}$ and set

$$\mathbf{s} := \mathbf{s}_0 \cdot \mathbf{W} + \mathbf{t}.$$

Recall that in H_1, we have $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_1},\chi_3}$. By noise flooding (Lemma 1), we have

$$(\mathbf{s}_0, \mathbf{W}, \mathbf{s}) \approx_s (\mathbf{s}_0, \mathbf{W}, \mathbf{s}_0 \cdot \mathbf{W} + \mathbf{t}),$$

as long as

$$\chi_3 \geq n_0 \cdot \sqrt{n_0} \cdot \chi_0 \cdot \sqrt{n_0 \cdot n_1} \cdot \chi_1 \cdot \lambda^{\omega(1)}.$$

We conclude that $\mathsf{H}_1 \approx_s \mathsf{H}_2$. Observe that we can now rewrite $(\mathbf{c}_1, \mathbf{c}_3)$ in the challenge ciphertext in H_2 as follows:

$$\begin{aligned}\mathbf{c}_1 &= \mathbf{s}_0 \cdot \mathbf{B}_1 + (\mathbf{s}_0 \cdot \mathbf{W} + \mathbf{t}) \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1 \\ &= \mathbf{s}_0 \cdot (\mathbf{B}_1 + \mathbf{W} \cdot \mathbf{G}_{n_1,p}) + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1 \bmod p,\end{aligned}$$

and

$$\begin{aligned}\mathbf{c}_3 &= (\mathbf{s}_0 \cdot \mathbf{W} + \mathbf{t}) \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}_3 \\ &= \mathbf{t} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \mathbf{s}_0 \cdot \mathbf{E} + \mathbf{e}_3 \mod q,\end{aligned}$$

where we used that

$$\mathbf{W} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) = [\mathbf{I}_{n_0} \mid \mathbf{W}_0] \cdot \begin{bmatrix} \mathbf{W}_0 \cdot \tilde{\mathbf{A}} + \mathbf{E} \\ -\tilde{\mathbf{A}} \end{bmatrix} = \mathbf{E} \mod q.$$

– H_3: This is the same as H_2, except for the following modification to $\mathbf{B}_1$ in mpk:

- sample $\mathbf{R} \leftarrow \{0,1\}^{m_0 \times m_1}$,
- set $\mathbf{B}_1 := \mathbf{B}_0 \cdot \mathbf{R} - \mathbf{W} \cdot \mathbf{G}_{n_1,p} \bmod p$.

Since $\mathbf{R}$ is sampled uniformly and $m_0 \geq (n_0 + 1) \cdot \log p + 2 \cdot \lambda$, indistinguishability ($\mathsf{H}_2 \approx_s \mathsf{H}_3$) follows from the leftover hash lemma (Lemma 2). Notice that we can now rewrite $\mathbf{c}_1$ in the challenge ciphertext in H_3 as

$$\begin{aligned}\mathbf{c}_1 &= \mathbf{s}_0 \cdot (\mathbf{B}_1 + \mathbf{W} \cdot \mathbf{G}_{n_1,p}) + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1 \\ &= \mathbf{s}_0 \cdot \mathbf{B}_0 \cdot \mathbf{R} + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}_1 \mod p.\end{aligned}$$

– H_4: This is the same as H_3, except for the following modification to KeyGen queries:

- recall the key equation

$$(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}_{n_1,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(\mathbf{x}) \cdot \mathbf{G}_{n_1,q} \mod q,$$

and that a valid adversary can only make KeyGen queries for functions f for which $f(\mathbf{x}^*) = 1$. Using these facts, one has that

$$
\begin{aligned}
\mathbf{B}_f &= \mathbf{B}_1 \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) \\
&= (\mathbf{B}_0 \cdot \mathbf{R} - \mathbf{W} \cdot \mathbf{G}_{n_1,p}) \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p) \\
&= \mathbf{B}_0 \cdot \underbrace{\mathbf{R} \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)}_{\mathbf{R}_f} - \mathbf{W} \cdot \lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \mathbf{W} \cdot \lfloor ((\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} + f(\mathbf{x}^*) \cdot \mathbf{G}_{n_1,q}) \cdot \mathbf{D} \rceil_p \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \mathbf{W} \cdot \lfloor ((\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} + \mathbf{G}_{n_1,q}) \cdot \mathbf{D} \rceil_p \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \lfloor \mathbf{W} \cdot ((\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} + \mathbf{G}_{n_1,q}) \cdot \mathbf{D} \rceil_p - \mathbf{E}_{\mathbf{W}} \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \left\lfloor \underbrace{\mathbf{W} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q})}_{\mathbf{E}} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} \cdot \mathbf{D} + \mathbf{W} \cdot \mathbf{G}_{n_1,q} \cdot \mathbf{D} \right\rceil_p - \mathbf{E}_{\mathbf{W}} \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \left\lfloor \underbrace{\mathbf{E} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} \cdot \mathbf{D}}_{\mathbf{E}'} + \mathbf{W} \cdot \mathbf{G}_{n_1,q} \cdot \mathbf{D} \right\rceil_p - \mathbf{E}_{\mathbf{W}} \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \lfloor \mathbf{W} \cdot \mathbf{G}_{n_1,q} \cdot \mathbf{D} \rceil_p - \underbrace{(\lfloor \mathbf{E}' \rceil_p + \mathbf{E}_{\mathbf{W}} + \mathbf{E}_+)}_{\mathbf{E}_f} \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \left\lfloor \frac{q}{p} \cdot \mathbf{G}_{n_0,p} \right\rceil_p - \mathbf{E}_f \\
&= \mathbf{B}_0 \cdot \mathbf{R}_f - \mathbf{G}_{n_0,p} - \mathbf{E}_f \mod p,
\end{aligned}
$$

where in the second-to-last equality we have used the definition of $\mathbf{W}$ and $\mathbf{D}$.

- compute $\mathbf{T}_f := \begin{bmatrix} \mathbf{R}_f \\ -\mathbf{I}_{m_0} \\ -\mathbf{E}_f \end{bmatrix}$ and observe that $[\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}] \cdot \mathbf{T}_f = \mathbf{G}_{n_0,p}$.
- compute

$$\mathbf{K}_f \leftarrow \mathsf{SamplePre}([\mathbf{B}_0 \mid \mathbf{B}_f \mid \mathbf{I}_{n_0}], \mathbf{T}_f, \mathbf{B}_2, \tau),$$

to answer KeyGen queries. By the properties of $\mathsf{SamplePre}$ (Lemma 4), this works as long as

$$
\begin{aligned}
\tau &\geq 2 \cdot (2 \cdot m_0 + n_0) \cdot \sqrt{n_0 \cdot \log p} \cdot \|\mathbf{T}_f\| \\
&= O(m_0^2 \cdot (\|\mathbf{R}_f\| + \|\mathbf{E}_f\|)),
\end{aligned}
$$

where

$$
\begin{aligned}
\|\mathbf{R}_f\| &= \|\mathbf{R} \cdot \mathbf{G}^{-1}_{n_1,p}(\lfloor \mathbf{A}_f \cdot \mathbf{D} \rceil_p)\| \\
&\leq m_1,
\end{aligned}
$$

and

$$
\begin{aligned}
\|\mathbf{E}_f\| &= \|\lfloor \mathbf{E}' \rceil_p + \mathbf{E}_{\mathbf{W}} + \mathbf{E}_+\| \\
&\leq \|\lfloor \mathbf{E}' \rceil_p\| + \|\mathbf{E}_{\mathbf{W}}\| + \|\mathbf{E}_+\| \\
&\leq \|\lfloor \mathbf{E} \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}^*} \cdot \mathbf{D} \rceil_p\| + n_1 \cdot \sqrt{n_0 \cdot n_1} \cdot \lambda \cdot \chi_1 + 2 \\
&\leq \frac{p}{q} \cdot \lambda \cdot \sqrt{n_0 \cdot \ell \cdot m_2} \cdot \chi_1 \cdot \ell \cdot m_2 \cdot m_2^{O(d)} \cdot m_2 + n_1 \cdot \sqrt{n_0 \cdot n_1} \cdot \lambda \cdot \chi_1 + 2.
\end{aligned}
$$

Therefore, since $\tau \geq \mathsf{poly}(m_2, \ell, \lambda) \cdot p/q \cdot B \cdot \chi_1$ satisfies such constraint, we have that $\mathsf{H}_3 \approx_s \mathsf{H}_4$. Notice that the reduction does not use $\mathbf{T}_{\mathbf{B}_0}$ anymore.

- H_5: This is the same as H_4, except for the following modification to $\mathbf{B}_0$ in mpk:
  - sample $\mathbf{B}_0 \leftarrow \mathbb{Z}_p^{n_0 \times m_0}$ instead of $(\mathbf{B}_0, \mathbf{T}_{\mathbf{B}_0}) \leftarrow \mathsf{TrapGen}(1^{n_0}, 1^{m_0}, p)$.

  By the properties of the TrapGen algorithm (Lemma 4), the distribution of $\mathbf{B}_0$ is statistically indistinguishable between H_4 and H_5. Therefore, $\mathsf{H}_4 \approx_s \mathsf{H}_5$.
- H_6: This is the same as H_5, except for the following modification to $\mathbf{c}_1, \mathbf{c}_3$ in the challenge ciphertext:
  - sets
  $$\mathbf{c}_1 \coloneqq \mathbf{c}_0 \cdot \mathbf{R} + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}'_1 \mod p,$$
  $$\mathbf{c}_3 \coloneqq \mathbf{t} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}'_3 \mod q,$$
  for $\mathbf{t} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_1},\chi_3}$, $\mathbf{e}'_1 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_1},\chi_2}$, and $\mathbf{e}_3 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_2},\chi_3}$.

  First, recall that in H_5, we have
  $$\mathbf{c}_1 = (\overbrace{\mathbf{s}_0 \cdot \mathbf{B}_0 + \mathbf{e}_0}^{\mathbf{c}_0}) \cdot \mathbf{R} + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \boxed{\mathbf{e}_1 - \mathbf{e}_0 \cdot \mathbf{R}} \mod p,$$
  $$\mathbf{c}_3 = \mathbf{t} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \boxed{\mathbf{s}_0 \cdot \mathbf{E} + \mathbf{e}_3} \mod q,$$
  where the boxed terms are the terms in H_5 that will be modified in H_6. By noise flooding (Lemma 1), we have
  $$(\mathbf{e}_0, \mathbf{R}, \boxed{\mathbf{e}_1 - \mathbf{e}_0 \cdot \mathbf{R}}) \approx_s (\mathbf{e}_0, \mathbf{R}, \mathbf{e}'_1),$$
  $$(\mathbf{s}_0, \mathbf{E}, \boxed{\mathbf{s}_0 \cdot \mathbf{E} + \mathbf{e}_3}) \approx_s (\mathbf{s}_0, \mathbf{E}, \mathbf{e}'_3).$$
  as long as
  $$\chi_2 \geq m_0 \cdot \sqrt{m_0} \cdot \chi_0 \cdot \lambda^{\omega(1)}, \quad \text{and}$$
  $$\chi_3 \geq n_0 \cdot \sqrt{n_0} \cdot \chi_0 \cdot \sqrt{n_0 \cdot \ell \cdot m_2} \cdot \chi_1 \cdot \lambda^{\omega(1)}.$$
  We conclude that $\mathsf{H}_5 \approx_s \mathsf{H}_6$.
- H_7: This is the same as H_6, except for the following modification to $\mathbf{c}_0, \mathbf{c}_2$ in the challenge ciphertext:
  - sample
  $$\mathbf{c}_0 \leftarrow \mathbb{Z}_p^{m_0}, \quad \mathbf{c}_2 \leftarrow \mathbb{Z}_p^{\lambda}.$$

  Recall that in H_6, we have
  $$[\mathbf{c}_0 \mid \mathbf{c}_2] = \mathbf{s}_0 \cdot [\mathbf{B}_0 \mid \mathbf{B}_2] + [\mathbf{e}_0 \mid \mathbf{e}_2] + [\mathbf{0} \mid \boldsymbol{\mu}_b \cdot \lfloor p/2 \rfloor].$$
  where $\mathbf{s}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0},\chi_0}, \mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_0},\chi_0}, \mathbf{e}_2 \leftarrow \mathcal{D}_{\mathbb{Z}^{\lambda},\chi_0}$. To show that $\mathsf{H}_6 \approx_c \mathsf{H}_7$, we rely on $\mathsf{sLWE}_{n_0,m_0+\lambda,\chi_0,p}$. The reduction works as follows:
  - it parses $\mathbf{B} = [\mathbf{B}_0 \mid \mathbf{B}_2] \in \mathbb{Z}_p^{n_0 \times (m_0+\lambda)}$ and $\tilde{\mathbf{c}} = [\tilde{\mathbf{c}}_0 \mid \tilde{\mathbf{c}}_2] \in \mathbb{Z}_p^{m_0+\lambda}$ obtained from the $\mathsf{sLWE}_{n_0,m_0+\lambda,\chi_0,p}$ instance,
  - it obtains $\mathbf{x}^*$ from the adversary $\mathcal{A}$,
  - it samples $\mathbf{R} \leftarrow \{0,1\}^{m_0 \times m_1}$, $\mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0 \times \ell \cdot m_2},\chi_1}$, $\mathbf{W}_0 \leftarrow \mathcal{D}_{\mathbb{Z}^{n_0 \times (n_1-n_0)},\chi_1}$, and computes $\mathbf{A}, \mathbf{B}_1$ in mpk as in H_6, while using $\mathbf{B}_0, \mathbf{B}_2$ obtained from the sLWE instance,
  - it answers KeyGen queries using $\mathbf{T}_f$ (which can be computed starting from $\mathbf{R}, \mathbf{E}$) as in H_6,
  - whenever the adversary $\mathcal{A}$ produces $(\boldsymbol{\mu}_0, \boldsymbol{\mu}_1)$, it samples $b \leftarrow \{0,1\}$, $\mathbf{t} \leftarrow \mathcal{D}_{\mathbb{Z}^{n_1},\chi_3}$, $\mathbf{e}'_1 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_1},\chi_2}$, $\mathbf{e}'_3 \leftarrow \mathcal{D}_{\mathbb{Z}^{m_2},\chi_3}$, and outputs
  $$\mathsf{ct} = (\tilde{\mathbf{c}}_0, \tilde{\mathbf{c}}_0 \cdot \mathbf{R} + \mathbf{t} \cdot \mathbf{G}_{n_1,p} + \mathbf{e}'_1, \tilde{\mathbf{c}}_2 + \boldsymbol{\mu}_b \cdot \lfloor p/2 \rfloor, \mathbf{t} \cdot (\mathbf{A} - \mathbf{x}^* \otimes \mathbf{G}_{n_1,q}) + \mathbf{e}'_3).$$

  Observe that
  - if $(\mathbf{B}, \tilde{\mathbf{c}})$ is a structured $\mathsf{sLWE}_{n_0,m_0+1,\chi_0,p}$ instance, the view of the adversary $\mathcal{A}$ is identical to H_6;
  - if $(\mathbf{B}, \tilde{\mathbf{c}})$ is a uniform random instance, the view of $\mathcal{A}$ is identical to H_7.

We conclude that $\mathsf{H}_6 \approx_c \mathsf{H}_7$.

Since the message $\boldsymbol{\mu}_b$ and the challenge bit b are perfectly hidden in H_7, this concludes the proof. □

Remark 1 (ciphertext pseudorandomness). In fact, we can prove a stronger property, namely that the challenge ciphertext is pseudorandom; this is also the case for the BGGHNSVV14 ABE. To do this, we need to show that $(\mathbf{c}_1, \mathbf{c}_3)$ in H_7 are pseudorandom. This requires introducing the following additional hybrids:

- H_8: use $\mathbf{T}_{\mathbf{B}_0}$ to simulate secret keys.
- H_9: replace $\mathbf{c}_1$ with a random $\mathbf{c}_1 \leftarrow \mathbb{Z}_p^{m_1}$. We have $\mathsf{H}_8 \approx_s \mathsf{H}_9$ via the leftover hash lemma (Lemma 2).
- H_{10}: replace $\mathbf{A}$ with a random $\mathbf{A} \leftarrow \mathbb{Z}_q^{n_1 \times (\ell \cdot m_2)}$. We have $\mathsf{H}_9 \approx_c \mathsf{H}_{10}$ from sLWE as in $\mathsf{H}_0 \approx_c \mathsf{H}_1$.
- H_{11}: replace $\mathbf{c}_3$ with a random $\mathbf{c}_3 \leftarrow \mathbb{Z}_q^{\ell m_2}$. We have $\mathsf{H}_{10} \approx_c \mathsf{H}_{11}$ from sLWE (with $\mathbf{t}$ as the LWE secret).

4.1 Reusable Garbled Circuits

Goldwassser et al. [GKP+13], with improvements from Boneh et al. [BGG+14], showed that starting from (i) an ABE scheme for $\mathcal{F}_{\ell,d,s}$ 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), we can construct a reusable garbling scheme for $\mathcal{F}_{\ell,d,s}$ 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 $s + \mathsf{poly}(\lambda) \cdot K(\ell', d', s')$;

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, we have the following corollary:

Corollary 1 (Reusable garbling scheme). *Assuming the hardness of LWE with 2^{n^δ} modulus-to-noise ratio, we have a reusable garbling scheme for $\mathcal{F}_{\ell,d,s}$ in the CRS model where*

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

Acknowledgements. We would like to thank Huijia (Rachel) Lin for helpful discussions about LLL22, as well as the FOCS reviewers for meticulous and constructive feedback. This work was completed while V.C. was at AIT. He was funded by the Austrian Science Fund (FWF) and netidee SCIENCE grant P31621-N38 (PROFET).

References

ACPS09. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Berlin, Heidelberg, August 2009. 7

Agr17. Shweta Agrawal. Stronger security for reusable garbled circuits, general definitions and attacks. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 3–35. Springer, Cham, August 2017. 4

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. 8

ALS16. Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 333–362. Springer, Berlin, Heidelberg, August 2016. 4

AWY20. Shweta Agrawal, Daniel Wichs, and Shota Yamada. Optimal broadcast encryption from LWE and pairings in the standard model. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 149–178. Springer, Cham, November 2020. 1

AY20. Shweta Agrawal and Shota Yamada. Optimal broadcast encryption from pairings and LWE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 13–43. Springer, Cham, May 2020. 1, 2

BCTW16. Zvika Brakerski, David Cash, Rotem Tsabary, and Hoeteck Wee. Targeted homomorphic attribute-based encryption. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 330–360. Springer, Berlin, Heidelberg, October / November 2016. 8

BGG$^+$14. Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Berlin, Heidelberg, May 2014. 1, 2, 3, 8, 17

BHR12. Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012. 2

BTVW17. Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained PRFs (and more) from LWE. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 264–302. Springer, Cham, November 2017. 3, 5, 6

BV11. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011. 1, 3

BV15. Zvika Brakerski and Vinod Vaikuntanathan. Constrained key-homomorphic PRFs from standard lattice assumptions - or: How to secretly embed a circuit in your PRF. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 1–30. Springer, Berlin, Heidelberg, March 2015. 1

CDG$^+$17. Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic oblivious transfer and its applications. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 33–65. Springer, Cham, August 2017. 1

CJJ22. Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. SNARGs for $\mathcal{P}$ from LWE. In *62nd FOCS*, pages 68–79. IEEE Computer Society Press, February 2022. 1

Gen09. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. 1

GKP$^+$13. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013. 1, 2, 17

GKPV10. Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In Andrew Chi-Chih Yao, editor, *ICS 2010*, pages 230–240. Tsinghua University Press, January 2010. 7

GPSW06. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309. 1

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 7, 8

GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Berlin, Heidelberg, August 2013. 3, 8

GVW13. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013. 1, 2

GVW15a. Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits from LWE. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 503–523. Springer, Berlin, Heidelberg, August 2015. 1

GVW15b. Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015. 1

HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. 7

HLL23. Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices: Garbled circuits of optimal size, laconic functional evaluation, and more. In *FOCS*, 2023. 2

LLL22. Hanjun Li, Huijia Lin, and Ji Luo. ABE for circuits with constant-size secret keys and adaptive security. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 680–710. Springer, Cham, November 2022. 1, 2, 8, 10

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. 3, 8

PS19. Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Cham, August 2019. 1

QWW18. Willy Quach, Hoeteck Wee, and Daniel Wichs. Laconic function evaluation and applications. In Mikkel Thorup, editor, *59th FOCS*, pages 859–870. IEEE Computer Society Press, October 2018. 1, 4

SW05. Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Berlin, Heidelberg, May 2005. 1

Tsa22. Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 535–559. Springer, Cham, August 2022. 1, 2

Wee22. Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Cham, May / June 2022. 1, 2

Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986. 2