Optimal Broadcast Encryption and CP-ABE from Evasive Lattice Assumptions

Hoeteck Wee

NTT Research and ENS, Paris

Abstract. We present a new, simple candidate broadcast encryption scheme for N users with parameter size $\mathsf{poly}(\log N)$. We prove security of our scheme under a non-standard variant of the LWE assumption where the distinguisher additionally receives short Gaussian pre-images, while avoiding zeroizing attacks. This yields the first candidate optimal broadcast encryption that is plausibly post-quantum secure, and enjoys a security reduction to a simple assumption. As a secondary contribution, we present a candidate ciphertext-policy attribute-based encryption (CP-ABE) scheme for circuits of a-priori bounded polynomial depth where the parameter size is independent of the circuit size, and prove security under an additional non-standard assumption.

1 Introduction

In this work, we study broadcast encryption [28] as well as attribute-based encryption schemes [42,35,11]. In ciphertext-policy attribute-based encryption (CP-ABE), ciphertexts ct are associated with a predicate f and a message m and keys sk with an attribute x, and decryption returns m when x satisfies f. Broadcast encryption is a special case of CP-ABE where the predicate is specified by a set $S \subseteq [N]$, and decryption returns m when $x \in S$. In both cases, we require security against unbounded collusions, so that an adversary that sees a ciphertext along with secret keys for an arbitrary number of attributes $x_1, x_2, \ldots$ learns nothing about m as long as none of these attributes satisfies f.

Broadcast encryption has been an active area of research since their introduction in the 1990s, where a major goal is to obtain schemes with short parameters, that is, short ciphertexts ct, public keys mpk and secret keys sk. In a celebrated work from 2005, Boneh, Gentry and Waters [14] presented the first broadcast encryption scheme with sublinear-sized parameters from bilinear groups where $|\mathsf{ct}| + |\mathsf{mpk}| + |\mathsf{sk}| = O(N^{1/2})$, [16,34,23], recently improved to $O(N^{1/3})$ [45]. On the other hand, in spite of the tremendous advances in lattice-based cryptography over the past decade, we do not know a LWE-based broadcast encryption scheme achieving $|\mathsf{ct}| = o(N)$.

A more recent line of works focuses on *optimal* broadcast encryption with parameter size $\mathsf{poly}(\log N)$, where the first feasibility results relied on either multi-linear maps [17] or indistinguishability obfuscation [18].[1]. In a recent remarkable break-through, Agrawal and Yamada [7] –along with a follow-up with Wichs [5]– constructed an optimal broadcast encryption scheme from bilinear groups *and* LWE. Independently, Brakerski and Vaikuntathan [21] presented a candidate "lattice-inspired' optimal broadcast encryption scheme that is plausibly post-quantum secure, but they were unable to provide a reduction to LWE or any simple lattice assumption.

Our Contributions. Our main contribution is a new, simple candidate optimal broadcast encryption scheme with $\mathsf{poly}(\log N)$-sized parameters. We prove selective security of our scheme assuming *evasive LWE*, a non-standard variant of the LWE assumption where the distinguisher additionally receives short Gaussian pre-images while avoiding zeroizing attacks. This yields the first candidate optimal broadcast encryption that is plausibly post-quantum secure, and enjoys a security reduction to a simple assumption. As a secondary contribution, we present a candidate CP-ABE scheme for circuits of a-priori bounded polynomial depth where the parameter size is independent of the circuit size, and prove security under an additional non-standard assumption. We refer to Fig 1 for a comparison with prior works, and proceed with a brief overview of our constructions.

[1] For simplicity of exposition and due to the sheer complexity and impracticality of the ensuing schemes, we ignore obfuscation-based broadcast in the rest of the introduction, deferring a comparison to Section 2.3.

Reference	Assumption	Post-Quantum	CP-ABE
AY20 [7]	LWE + bilinear GGM		NC^1, $\|\mathsf{ct}\| = \mathsf{poly}(\ell, d, \log s)$
AWY20 [5]	LWE + bilinear KOALA		NC^1, $\|\mathsf{ct}\| = \mathsf{poly}(\ell, d, \log s)$
BV22 [21]	×	✓	circuits, $\|\mathsf{ct}\| = \mathsf{poly}(\ell, d, \log s)$
Section 5.3	evasive LWE	✓	NC^1, $\|\mathsf{ct}\| = \mathsf{poly}(2^d, \log s)$
Section 5.4	evasive LWE + tensor LWE	✓	circuits, $\|\mathsf{ct}\| = \mathsf{poly}(d, \log s)$

Fig. 1. Comparison with prior optimal broadcast encryption schemes (sans obfuscation), all of which also yield CP-ABE schemes for either NC^1 or circuits of (a-prior bounded) polynomial depth d. Broadcast encryption for N users correspond to circuits of size $O(N \log N)$ and depth $O(\log\log N)$. CP-ABE decryption time in AY20, AMY20 grows with 2^d, hence the limitation to NC^1 circuits. As in [5], our broadcast encryption schemes achieve selective security, and our CP-ABE schemes achieve very selective security. BV22 only shows LWE-hardness against a subclass of attacks on a specific component of their scheme and does not provide any reduction for their full scheme. Both evasive LWE and bilinear KOALA are non-falsifiable assumptions. Finally, bilinear GGM ⇒ bilinear KOALA.

2 Technical Overview

Our optimal broadcast encryption scheme follows the Agrawal-Yamada-Wichs, henceforth AYW, blue-print laid out in [7,5] (and partially in [21]): (i) we start with a one-key secure CP-ABE for circuits based on LWE and randomize the secret keys to achieve security against collusions, and (ii) we show that for an appropriate family of circuits, our CP-ABE scheme implies optimal broadcast encryption. The AYW schemes achieve randomization via exponentiation with random scalars in a bilinear group. Security relies on LWE in addition to a hardness assumption about the bilinear group, either the generic group model (GGM) [7], or non-standard knowledge assumption (KOALA) [5,12]. We proceed to sketch two new technical ideas in this work that allows us to eliminate the use of bilinear maps, thereby achieving plausible post-quantum security.

Randomization via tensors. We randomize secret keys by tensoring with random Gaussian (row) vectors $\mathbf{r} \leftarrow \mathcal{D}^m_{\mathbb{Z},\chi}$, which satisfies the following correctness and security properties:

- Following prior ABE schemes based on LWE [13], given $\mathbf{x} \in \{0,1\}^\ell, \mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, we can homomorphically evaluate a circuit f on $\mathbf{A} - \mathbf{x} \otimes \mathbf{G}$ to obtain a quantity of the form $\mathbf{A}_f - f(x)\mathbf{G}$ via right-multiplication by some low-norm matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$. This property is preserved under tensoring with random Gaussian vectors $\mathbf{r}$: we can homomorphically evaluate f on $(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}^\top$ to obtain $(\mathbf{A}_f - f(x)\mathbf{G}) \otimes \mathbf{r}^\top$ via right multiplication by $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \otimes \mathbf{I}$. Note that homomorphic evaluation is not possible if we replace tensor product with vector multiplication (on the right).
- Tensoring "amplifies" a single LWE secret $\mathbf{s}$ into Q independent LWE secrets $\mathbf{s}_1, \ldots, \mathbf{s}_Q$. More formally, under the LWE assumption, we have

$$\left\{ \left(\boxed{\mathbf{s}(\mathbf{I}_n \otimes \mathbf{r}_i^\top) + \mathbf{e}_i}, \mathbf{r}_i^\top \right) \right\}_{i \in [Q]} \approx_c \left\{ \left(\mathbf{s}_i, \mathbf{r}_i^\top \right) \right\}_{i \in [Q]} \tag{1}$$

  where $\mathbf{s} \leftarrow \mathbb{Z}_q^{nm}, \mathbf{s}_i \leftarrow \mathbb{Z}_q^n, \mathbf{e}_i \leftarrow \mathcal{D}^n_{\mathbb{Z},\chi}, \mathbf{r}_i \leftarrow \mathcal{D}^m_{\mathbb{Z},\chi}$ [15,22]. In our analysis, Q corresponds to the number of key queries, and having Q independent secrets enables a hybrid argument over the key queries.

We describe an alternative derivation of these ideas at the end of Section 2.1.

An evasive lattice assumption. We describe a simple variant of the *evasive LWE* assumption we put forth in this work. Fix an efficiently samplable distribution $\mathbf{P}$ over $\mathbb{Z}_q^{n \times t}$. The evasive LWE assumption allows us to assert statements of the form

$$(\mathbf{B}, \boxed{\mathbf{sB} + \mathbf{e}}, \mathbf{B}^{-1}(\mathbf{P})) \approx_c (\mathbf{B}, \mathbf{c}, \mathbf{B}^{-1}(\mathbf{P}))$$

where $\mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{c} \in \mathbb{Z}_q^{m'}$ are uniformly random, $m = O(n \log q) \le t$ (so that $\mathbf{P}$ is wider than $\mathbf{B}$). We have two distinguishing strategies in the literature:

– ignore $\mathbf{B}^{-1}(\mathbf{P})$ and distinguish $(\mathbf{B}, \mathbf{sB}+\mathbf{e})$ from $(\mathbf{B}, \mathbf{c})$ – this covers lattice attacks on LWE;
– compute $\mathbf{c}^* = (\mathbf{sB}+\mathbf{e}') \cdot \mathbf{B}^{-1}(\mathbf{P}) \approx \mathbf{sP}$ and distinguish the latter from uniform – this includes zeroizing attacks on multi-linear map and obfuscation candidates [25,41,24,37].

The evasive LWE assumption essentially asserts that these are the only distinguishing attacks. Namely,

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

where $\mathbf{e}''$ is a fresh noise vector. Note that $\mathbf{sP}+\mathbf{e}'' \approx_c \mathbf{c}''$ implies that the high-order bits of $(\mathbf{sB}+\mathbf{e}') \cdot \mathbf{B}^{-1}(\mathbf{P}) \approx \mathbf{sP}$ are pseudorandom, thereby defeating the second distinguishing strategy.[2] Overall, we note that the statement of evasive LWE is fairly simple and general, and does not refer to tensor products, circuits, or structured distributions like $\mathbf{A}-\mathbf{x}\otimes\mathbf{G}$ or $\mathbf{A}_f$. That is, the assumption encapsulates a principled approach towards (conjectured) computational hardness, rather than one that is tailored to our scheme.

Proof strategy. Our security proof proceeds in two steps: first, we rely on evasive LWE to reduce security of our scheme to a simpler statement with no short Gaussians, and then we prove this latter statement from LWE, using (6) along the way. For the second step, we need to modify the scheme to perform homomorphic evaluation on $\mathbf{A}-\mathbf{x}\otimes\mathbf{I}$ where $\mathbf{A}$ is a low-norm matrix, and we replaced the gadget matrix $\mathbf{G}$ with the identity matrix $\mathbf{I}$; in the security proof, we will use the fact that if $\mathbf{A}-\mathbf{x}\otimes\mathbf{I}$ is low-norm, then

$$\underbrace{\mathbf{s}((\mathbf{A}-\mathbf{x}\otimes\mathbf{I})\otimes\mathbf{r}^\top)}_{\approx} \approx \underbrace{\mathbf{s}(\mathbf{I}\otimes\mathbf{r}^\top)}_{\approx} \cdot (\mathbf{A}-\mathbf{x}\otimes\mathbf{I})$$

upon which we can invoke (6) to replace $\mathbf{s}(\mathbf{I}\otimes\mathbf{r}^\top)$ on the RHS with random.

Homomorphic evaluation on $\mathbf{A}-\mathbf{x}\otimes\mathbf{I}$ works as before with $\mathbf{G}$, except the noise growth is now doubly (instead of singly) exponential in circuit depth. This yields a CP-ABE scheme with $|\mathsf{ct}| = \mathsf{poly}(2^d, \log s)$ for NC^1 circuits of multiplicative depth d and size s, and we show that this is sufficient for optimal broadcast encryption. In particular, broadcast encryption for N users correspond to circuits of multiplicative depth $O(\log\log N)$ and size $O(N\log N)$. To obtain a CP-ABE for a-prior bounded depth circuits with $|\mathsf{ct}| = \mathsf{poly}(d, \log s)$, we keep $\mathbf{A}-\mathbf{x}\otimes\mathbf{G}$ as before, and instead prove security based a new (falsifiable) "*tensor LWE*" assumption in the second step.

2.1 Our CP-ABE Schemes

We describe our CP-ABE schemes in more detail. The schemes rely on the following strengthening of our earlier statement of evasive LWE: we consider distributions over pairs of matrices $(\mathbf{A}', \mathbf{P})$ together with auxiliary input aux (instead of just $\mathbf{P}$) and require that

$$\begin{aligned} &\textbf{if} && (\mathbf{A}', \mathbf{B}, \mathbf{P}, \boxed{\mathbf{sA}'+\mathbf{e}'}, \boxed{\mathbf{sB}+\mathbf{e}}, \boxed{\mathbf{sP}+\mathbf{e}''}, \mathsf{aux}) \approx_c (\mathbf{A}', \mathbf{B}, \mathbf{P}, \mathbf{c}', \mathbf{c}, \mathbf{c}'', \mathsf{aux}), && (2) \\ &\textbf{then} && (\mathbf{A}', \mathbf{B}, \boxed{\mathbf{sA}'+\mathbf{e}'}, \boxed{\mathbf{sB}+\mathbf{e}}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{A}', \mathbf{B}, \mathbf{c}', \mathbf{c}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) && (3) \end{aligned}$$

In our applications, the auxiliary input includes the coin tosses used to sample $\mathbf{A}', \mathbf{P}$, which rules out obfuscation-based counter-examples.

A one-key secure CP-ABE. We consider CP-ABE for circuits $f : \{0,1\}^\ell \to \{0,1\}$ of depth d and size s. Following [7,21], we begin with a one-key secure CP-ABE (where we use curly underlines in place of noise terms):

$$\begin{aligned} \mathsf{mpk} &:= \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{A} \leftarrow \mathbb{Z}_q^{n\times \ell m}, \mathbf{u}^\top \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^m \\ \mathsf{ct}_f &:= \underbrace{\mathbf{sA}_f\mathbf{u}^\top + \mu\cdot\mathbf{g}}, \underbrace{\mathbf{sB}_1}, \text{ where } \mathbf{s} \leftarrow \mathbb{Z}_q^n \\ \mathsf{sk}_{\mathbf{x}} &:= \mathbf{B}_1^{-1}(\mathbf{A}-\mathbf{x}\otimes\mathbf{G}) \end{aligned}$$

[2] Note that the error distribution $\mathbf{e}\cdot\mathbf{B}^{-1}(\mathbf{P})$ in $\mathbf{c}^*$ is different from the fresh Gaussian error $\mathbf{e}''$. Differences in error distributions can make or break a scheme if $\mathbf{c}^*$ has small norm, but we do not know attacks exploiting these differences when $\mathbf{c}^*$ has large norm, as is the case here.

Note that the ciphertext size is independent of ℓ. Decryption for $f(\mathbf{x}) = 0$ uses $(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(\mathbf{x})\mathbf{G}$, which implies $\underbrace{\mathbf{s}\mathbf{B}_1} \cdot \mathbf{B}_1^{-1}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} \cdot \mathbf{u}^\top \approx \mathbf{s}\mathbf{A}_f\mathbf{u}^\top$.

Next, we show that the scheme is one-key secure assuming LWE and evasive LWE. Intuitively, evasive LWE says that we can replace the terms $\underbrace{\mathbf{s}\mathbf{B}_1}, \mathbf{B}_1^{-1}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$ with their product $\underbrace{\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$. Then, it suffices to show that μ is hidden given

$$\mathbf{B}_1, \mathbf{A}, \underbrace{\mathbf{s}\mathbf{A}_f\mathbf{u}^\top} + \mu \cdot \mathbf{g}, \underbrace{\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$$

Next, we can write $\mathbf{s}\mathbf{A}_f\mathbf{u}^\top$ in terms of $\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$ and $f(\mathbf{x}) \cdot \mathbf{s}\mathbf{G}\mathbf{u}^\top$ using homomorphic computation. Since $f(\mathbf{x}) = 1$, it suffices to show that μ is hidden given

$$\mathbf{B}_1, \mathbf{A}, \underbrace{\mathbf{s}\mathbf{G}\mathbf{u}^\top} + \mu \cdot \mathbf{g}, \underbrace{\mathbf{s}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$$

which follows quite readily from LWE.

Note that this scheme is insecure if the adversary is allowed to make two key queries: given secret keys for 0^ℓ and 1^ℓ, an adversary can compute $\underbrace{\mathbf{s}\mathbf{A}}, \underbrace{\mathbf{s}(\mathbf{A} - 1^\ell \otimes \mathbf{G})}$, substract the two to obtain $\underbrace{\mathbf{s}(1^\ell \otimes \mathbf{G})}$ and solve for $\mathbf{s}$ and thus μ. To defeat this attack, we randomize the secret keys by tensoring with random Gaussian vectors.

First modification. We replace $\mathbf{A} - \mathbf{x} \otimes \mathbf{G}$ in sk with $(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}^\top$ and $\mathbf{s}\mathbf{A}_f\mathbf{u}^\top$ in ct with $\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I})$, so that

$$\mathsf{ct}_f := \underbrace{\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I})} + \mu \cdot \mathbf{g}, \underbrace{\mathbf{s}\mathbf{B}_1}, \text{ where } \mathbf{s} \leftarrow \mathbb{Z}_q^{nm}$$

$$\mathsf{sk}_{\mathbf{x}} := \mathbf{B}_1^{-1}((\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}^\top), \mathbf{r}^\top$$

Decryption computes the following quantities:

$$(\underbrace{\mathbf{s}(\mathbf{A}_f \otimes \mathbf{I})} + \mu \cdot \mathbf{g}) \cdot (\mathbf{I} \otimes \mathbf{r}^\top) \approx \mathbf{s}(\mathbf{A}_f \otimes \mathbf{r}^\top) + \mu \cdot \mathbf{g} \cdot (\mathbf{I} \otimes \mathbf{r}^\top)$$

$$\underbrace{\mathbf{s}\mathbf{B}_1} \cdot \mathbf{B}_1^{-1}((\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}^\top) \cdot (\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \otimes \mathbf{I}) \approx \mathbf{s}(\mathbf{A}_f \otimes \mathbf{r}^\top)$$

and subtracts the two to recover μ. The attacker from before now learns $\underbrace{\mathbf{s}(\mathbf{A} \otimes \mathbf{r}_1^\top)}, \underbrace{\mathbf{s}((\mathbf{A} - 1^\ell \otimes \mathbf{G}) \otimes \mathbf{r}_2^\top}$ and since $\mathbf{r}_1 \neq \mathbf{r}_2$ w.h.p., we can no longer carry out the attack from before.

We do not know an attack on the preceding scheme. However, adapting the security proof for the one-key setting to the many-key setting runs into two difficulties. Upon applying evasive LWE as before, we want to argue that μ is hidden given

$$\mathbf{B}_1, \mathbf{A}, \underbrace{\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I})} + \mu \cdot \mathbf{g}, \left\{\underbrace{\mathbf{s}((\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}) \otimes \mathbf{r}_i^\top)}, \mathbf{r}_i^\top\right\}_{i \in [Q]}$$

- The first difficulty lies in handling $\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I})$: using homomorphic computation as before allows us to write $\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{r}_i^\top)$ in terms of $\mathbf{s}((\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}) \otimes \mathbf{r}_i^\top)$ and $f(\mathbf{x}_i) \cdot \mathbf{s}(\mathbf{G}\mathbf{u}^\top \otimes \mathbf{r}_i^\top)$. We then need to bridge the gap between $\left\{\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{r}_i^\top)\right\}_{i \in [Q]}$ (what we know how to simulate) and $\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I})$ (what appears in the ciphertext). The next modification addresses this difficulty while relying only on the LWE assumption.
- This leaves us with arguing pseudorandomness of $\left\{\underbrace{\mathbf{s}((\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}) \otimes \mathbf{r}_i^\top)}, \mathbf{r}_i^\top\right\}_{i \in [Q]}$, for which we present two solutions.

  The first (and less satisfactory) is to simply assert pseudorandomness via a new assumption, which we refer to as *tensor LWE*. This assumption is qualitatively different from evasive LWE in that there are no Gaussian pre-images. The second solution relies only on the LWE assumption, but incurs a 2^d blow-up, which is nonetheless sufficient for optimal broadcast encryption.

Second modification. We mask $\mathbf{s}(\mathbf{A}_f \otimes \mathbf{I})$ in the cipertext with a fresh LWE sample $\mathbf{s}_0\mathbf{A}_0 + \mathbf{e}_0$ and during decryption, compute

$$\mathbf{s}(\mathbf{A}_f \otimes \mathbf{r}^\top) \approx \overbrace{(\mathbf{s}(\mathbf{A}_f \otimes \mathbf{I}_m) + \mathbf{s}_0\mathbf{A}_0 + \mathbf{e})}^{\mathsf{ct}} \cdot (1 \otimes \overbrace{\mathbf{r}^\top}^{\mathsf{sk}}) - \overbrace{(\mathbf{s}_0\mathbf{B}_0 + \mathbf{e}_0)}^{\mathsf{ct}} \cdot \overbrace{\mathbf{B}_0^{-1}(\mathbf{A}_0\mathbf{r}^\top)}^{\mathsf{sk}} \tag{4}$$

where $\mathbf{s}_0\mathbf{B}_0 + \mathbf{e}_0$ appears in ct_f and $\mathbf{B}_0^{-1}(\mathbf{A}_0\mathbf{r}^\top)$ in $\mathsf{sk}_\mathbf{x}$.[3] This yields the following CP-ABE scheme for bounded depth circuits:

$$\begin{aligned}
\mathsf{mpk} &:= \mathbf{A}_0, \mathbf{B}_0 \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{mn\times m^2}, \mathbf{A} \leftarrow \mathbb{Z}_q^{n\times \ell m} \\
\mathsf{ct}_f &:= \underbrace{\mathbf{s}_0\mathbf{B}_0}, \underbrace{\mathbf{s}(\mathbf{A}_f \otimes \mathbf{I}_m) + \mathbf{s}_0\mathbf{A}_0 + \mu\cdot\mathbf{g}}, \underbrace{\mathbf{s}\mathbf{B}_1}, \text{ where } \mathbf{s} \leftarrow \mathbb{Z}_q^{mn}, \mathbf{s}_0 \leftarrow \mathbb{Z}_q^n \\
\mathsf{sk}_\mathbf{x} &:= \mathbf{B}_0^{-1}(\mathbf{A}_0\mathbf{r}^\top), \mathbf{B}_1^{-1}((\mathbf{A} - \mathbf{x}\otimes\mathbf{G})\otimes\mathbf{r}^\top), \mathbf{r}^\top, \text{ where } \mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^m
\end{aligned}$$

Decryption for $f(\mathbf{x}) = 0$ computes (approximately)

$$\begin{aligned}
\mu\cdot\mathbf{g}\cdot(1\otimes\mathbf{r}^\top) \approx\ & (\underbrace{\mathbf{s}(\mathbf{A}_f\otimes\mathbf{I}_m) + \mathbf{s}_0\mathbf{A}_0} + \mu\cdot\mathbf{g})\cdot(1\otimes\mathbf{r}^\top) - \underbrace{\mathbf{s}_0\mathbf{B}_0}\cdot\mathbf{B}_0^{-1}(\mathbf{A}_0\mathbf{r}^\top) \\
& + \underbrace{\mathbf{s}\mathbf{B}_1}\cdot\mathbf{B}_1^{-1}((\mathbf{A}-\mathbf{x}\otimes\mathbf{G})\otimes\mathbf{r}^\top)\cdot(\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\otimes\mathbf{I})
\end{aligned}$$

Again, via the evasive LWE assumption (upon additionally combining $\mathbf{B}_0, \mathbf{B}_1$ into a single matrix $\mathbf{B}$), ABE security reduces to proving pseudorandomness of

$$\mathbf{A}_0, \mathbf{A}, \mathbf{u}^\top, \overbrace{\underbrace{\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}) + \mathbf{s}_0\mathbf{A}_0}}^{\mathbf{c}'}, \left\{\underbrace{\mathbf{s}((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{G})\otimes\mathbf{r}_i^\top)}, \underbrace{\mathbf{s}_0\mathbf{A}_0\mathbf{r}_i^\top}, \mathbf{r}_i^\top\right\}_{i\in[Q]}$$

Observe that

$$\underbrace{\mathbf{s}_0\mathbf{A}_0\mathbf{r}_i^\top} \approx \mathbf{c}'\cdot\mathbf{r}_i^\top - \underbrace{\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{r}_i^\top)}$$

We can then use the LWE assumption with secret $\mathbf{s}_0$ to replace $\mathbf{c}'$ with random. This leaves us with proving pseudorandomness of

$$\mathbf{A}_0, \mathbf{A}, \mathbf{u}^\top, \left\{\underbrace{\mathbf{s}((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{G})\otimes\mathbf{r}_i^\top)}, \underbrace{\mathbf{s}(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{r}_i^\top)}, \mathbf{r}_i^\top\right\}_{i\in[Q]}$$

At this point, we can apply homomorphic computation to $\mathbf{s}((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{G})\otimes\mathbf{r}_i^\top)$ as before in the one-key scheme, upon which we are left with proving pseudorandomness of

$$\mathbf{A}, \mathbf{u}^\top, \left\{\underbrace{\mathbf{s}((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{G})\otimes\mathbf{r}_i^\top)}, \underbrace{\mathbf{s}(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)}, \mathbf{r}_i^\top\right\}_{i\in[Q]} \tag{5}$$

The tensor LWE assumption essentially states that the above distribution is pseudorandom.

Third modification. The third and final modification allows us to handle the second difficulty without introducing the additional tensor LWE assumption but with a 2^d blow-up. The idea is to replace $\mathbf{G}$ in $\mathsf{sk}_\mathbf{x}$ with $\mathbf{I}_m$ and sample $\mathbf{A} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{m\times\ell m}$ so that $\mathbf{A} - \mathbf{x}\otimes\mathbf{I}_m$ has low-norm:

$$\begin{aligned}
\mathsf{mpk} &:= \mathbf{A}_0, \mathbf{B}_0 \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{m^2\times O(m^2\log q)}, \mathbf{A} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{m\times\ell m} \\
\mathsf{ct}_f &:= \underbrace{\mathbf{s}_0\mathbf{B}_0}, \underbrace{\mathbf{s}(\mathbf{A}_f\otimes\mathbf{I}_m) + \mathbf{s}_0\mathbf{A}_0 + \mu\cdot\mathbf{g}}, \underbrace{\mathbf{s}\mathbf{B}_1}, \text{ where } \mathbf{s} \leftarrow \mathbb{Z}_q^{m^2}, \mathbf{s}_0 \leftarrow \mathbb{Z}_q^n \\
\mathsf{sk}_\mathbf{x} &:= \mathbf{B}_0^{-1}(\mathbf{A}_0\mathbf{r}^\top), \mathbf{B}_1^{-1}((\mathbf{A}-\mathbf{x}\otimes\mathbf{I})\otimes\mathbf{r}^\top), \mathbf{r}^\top, \text{ where } \mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^m
\end{aligned}$$

In the security proof, instead of (5), we need to prove pseudorandomness of

$$\mathbf{A}, \mathbf{u}^\top, \left\{\underbrace{\mathbf{s}((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)}, \underbrace{\mathbf{s}(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)}, \mathbf{r}_i^\top\right\}_{i\in[Q]}$$

Both $\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m$ and $\mathbf{u}^\top$ have low-norm, so

$$\begin{aligned}
\underbrace{\mathbf{s}((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)} &\approx \underbrace{\mathbf{s}(\mathbf{I}\otimes\mathbf{r}_i^\top)}\cdot(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}) \\
\underbrace{\mathbf{s}(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)} &\approx \underbrace{\mathbf{s}(\mathbf{I}\otimes\mathbf{r}_i^\top)}\cdot\mathbf{u}^\top
\end{aligned}$$

[3] This modification plays a role similar to that of inner product functional encryption in [5], instantiated using ideas from the LWE-based scheme in [3]. The latter does not support inner product modulo q, which is what we need here.

We may then invoke (6) to replace $\mathbf{s}(\mathbf{I}\otimes\mathbf{r}_i^\top)$ with $\mathbf{s}_i \leftarrow \mathbb{Z}_q^m$, upon which it suffices to prove pseudorandomness of

$$\mathbf{A}, \mathbf{u}^\top, \left\{\mathbf{s}_i(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m), \mathbf{s}_i\mathbf{u}^\top\right\}_{i\in[Q]}$$

This in turn follows from LWE via a straight-forward hybrid argument over $i \in [Q]$.

An alternative derivation. We present an alternative derivation of randomization via tensors.[4] As before, we want to "amplify" a single LWE secret $\mathbf{s}$ into Q independent LWE secrets $\mathbf{s}_1, \ldots, \mathbf{s}_Q$. Following [21], we replace $\mathbf{s}$ with a matrix $\mathbf{S}$ and observe that by LWE, we have

$$\left\{(\mathbf{r}_i\mathbf{S}+\mathbf{e}_i, \mathbf{r}_i^\top)\right\}_{i\in[Q]} \approx_c \left\{(\mathbf{s}_i, \mathbf{r}_i^\top)\right\}_{i\in[Q]} \tag{6}$$

where $\mathbf{S} \leftarrow \mathbb{Z}_q^{m\times n}, \mathbf{s}_i \leftarrow \mathbb{Z}_q^n, \mathbf{e}_i \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^n, \mathbf{r}_i \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^m$ [15,22]. That is, during decryption, we want to (approximately) compute the product

$$\mathbf{r}\cdot\mathbf{S}\cdot(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})$$

Observe that the term $\mathbf{S}$ which depends on the ciphertext is sandwiched between two terms $\mathbf{r}, \mathbf{x}$ that depend on the key. The key observation is that we can rewrite the above product as

$$\mathsf{flat}(\mathbf{S})\cdot((\mathbf{A}-\mathbf{x}\otimes\mathbf{G})\otimes\mathbf{r}^\top)$$

where flat "flattens" a matrix into a row vector by concatenating the rows of the input matrix, which we can in turn write as the product of $\mathsf{flat}(\mathbf{S})\cdot\mathbf{B}_1$ and $\mathbf{B}_1^{-1}((\mathbf{A}-\mathbf{x}\otimes\mathbf{G})\otimes\mathbf{r}^\top)$. This yields the following variant of the scheme described in our first modification:

$$\mathsf{ct}_f := \mathbf{S}\mathbf{A}_f\mathbf{u}^\top + \mu\cdot\mathbf{g}^\top, \mathsf{flat}(\mathbf{S})\cdot\mathbf{B}_1, \text{ where } \mathbf{S} \leftarrow \mathbb{Z}_q^{m\times n}$$
$$\mathsf{sk}_\mathbf{x} := \mathbf{B}_1^{-1}((\mathbf{A}-\mathbf{x}\otimes\mathbf{G})\otimes\mathbf{r}^\top), \mathbf{r}^\top$$

Decryption first computes

$$\mathsf{flat}(\mathbf{S})\cdot\mathbf{B}_1\cdot\mathbf{B}_1^{-1}((\mathbf{A}-\mathbf{x}\otimes\mathbf{G})\otimes\mathbf{r}^\top) \approx \mathbf{r}\cdot\mathbf{S}\cdot(\mathbf{A}-\mathbf{x}\otimes\mathbf{G})$$

We can then carry out homomorphic evaluation whenever $f(\mathbf{x}) = 0$ to recover $\mathbf{r}\cdot\mathbf{S}\cdot\mathbf{A}_f\mathbf{u}^\top$, which we can combine with the first term of ct_f to recover μ.

2.2 On Evasive Lattice Assumptions

In the past decade, we have witnessed a large number of "lattice-inspired" schemes, on which weaknesses and attacks were subsequently discovered. A partial list includes:

- multi-linear maps and key exchange [29,33] and attacks in [25,26]
- obfuscation for branching programs [30,33,36] and attacks in [24]
- noisy inner product functional encryption [1] with attacks and fixes [4]
- obfuscation from circular security [19,32,46,20] with attacks on [32,46] in [37]

In fact, our evasive LWE assumption shares some structural similarities to the GGH15-based multi-linear maps [33] corresponding to the first two items on the list above. There is however a key conceptual distinction which we briefly alluded to earlier and shall expand on next.

[4] We arrived at this alternative derivation while preparing the conference talk, after submitting the camera-ready version of this work.

The zeroizing regime. All of the afore-mentioned attacks have one thing in common: they pertain to the *zeroizing regime* where an attacker can easily obtain sufficiently many equations in low-norm secret values —low-norm LWE secrets, error vectors, or both— over the integers that information-theoretically determine these secret values.[5] These equations arise naturally from the interaction of the correctness constraints and the security requirements. Such attacks are referred to in the literature as *zeroizing attacks*. Prior zeroizing attacks basically proceed in two steps: (i) collect many of these equations, and (ii) using these equations to recover some secret value and break security. The first step is typically fairly straight-forward; most of the technical and creative work lies in the second step, which varies from computing a linear-algebraic quantity (e.g., kernel [25] or rank [24,4]) of a carefully crafted matrix over the integers/reals, to more sophisticated sum-of-squares attacks [9,40].

Our evasive LWE assumption falls outside of this zeroizing regime in that we do not see any straight-forward way to collect even a single equation of the underlying LWE error vectors over the integers. As explained earlier in the introduction, the straight-forward adaptation of prior attacks would be to compute $\mathbf{c}^* = (\mathbf{sB} + \mathbf{e}_0) \cdot \mathbf{B}^{-1}(\mathbf{P}) \approx \mathbf{sP}$, but the pre-condition for evasive LWE implies that $\mathbf{c}^*$ has large norm and does not yield an equation over the integers. The setting for our assumption is closer to that for prior witness encryption candidates, specifically, the GGH15-based witness encryption candidate in [24], which also fall outside the zeroizing regime. Indeed, there are no known attacks on any witness encryption candidates in the literature, giving us additional confidence in our evasive LWE assumption. For the crypt-analysts who believe that existing witness encryption candidates are broken but haven't found an attack, our evasive LWE assumption provides a much simpler target for crypt-analysis.

Perspective. To the best of our knowledge, our evasive LWE assumption is the first simple lattice assumption that falls outside of the zeroizing regime. We firmly believe that the study of such *evasive* lattice assumptions —hardness, attacks, and constructions— constitutes an important and promising research direction, as well as a rich source of open problems. More broadly, non-standard variants of LWE and evasive lattice assumptions are conceptually similar to q-type assumptions, knowledge assumptions, and generic/algebraic group model assumptions that have played an essential role in our study of group and pairing-based cryptography.[6] This analogy provides additional impetus for the study of evasive lattice assumptions.

Looking ahead. Looking ahead, we see 4 possible scenarios, starting with the most optimistic:

1. This work ultimately leads to optimal broadcast encryption based on LWE, as has been the case for several lattice-based schemes where the initial candidates were based on non-standard assumptions (outside the zeroizing regime), such as fully homomorphic encryption and its multi-key variant and the Fiat-Shamir heuristic.
2. The evasive LWE assumption surives cryptanalysis: this could enable other advanced encryption primitives such as witness encryption [43,44].
3. The evasive LWE assumption is broken but the broadcast encryption scheme is not. This would require new and valuable crypt-analytic advances beyond the state-of-the-art zeroizing attacks. The current statement of evasive LWE is fairly general, and an attack could guide us towards identify more secure variants of the assumption that would suffice for our broadcast encryption scheme.
4. Both the evasive LWE assumption and the broadcast encryption are broken. Could these new attacks be extended to current GGH15-based witness encryption candidates?

We believe any of these scenarios would advance our current scientific understanding of lattice-based cryptography and assumptions (hardness and/or attacks).

[5] As a point of comparison, we have examples such as k-LWE [39] and inner product functional encryption [3] based on LWE where it is easy to obtain a few such equations, but the equations do *not* information-theoretically determine the secret values.

[6] Security based on evasive LWE can be viewed as ruling out restricted adversaries that replaces $\mathbf{sB} + \mathbf{e}, \mathbf{B}^{-1}(\mathbf{P})$ with their product $\mathbf{sP} + \mathbf{e}''$ (with fresh noise) and ignoring $\mathbf{B}^{-1}(\mathbf{P})$ thereafter. Viewed this way, evasive LWE can be seen as a partial analogue of the generic/algebraic group model used in group and pairing-based cryptography. Several works studied analogues of the generic group model for multi-linear maps [31,10], but they were in the zeroizing regime.

2.3 Additional Related Work

We describe additional related work.

Relation to GGH15 multi-linear maps. Our work draws upon several insights in the study of GGH15 multi-linear maps [33,22]. First, randomization via tensors and the statement in (6) both appeared in [22], but in a different context. Second, the intuition for evasive LWE in terms of an "optimal" distinguishing strategy also underlies earlier GGH15-based schemes, with the crucial distinction that evasive LWE falls outside the zeroizing regime. Our evasive LWE assumption also provides a concise statement of this intuition in a setting that falls outside the zeroizing regime.

Obfuscation-based broadcast. We can obtain optimal broadcast encryption schemes by combining the obfuscation-based scheme in [18] with the state-of-the-art obfuscaton schemes/candidates. The ensuing schemes would be extremely complex and impractical, inherited from the current obfuscation schemes/candidates, compounded with the use of non-black-box techniques. Nonetheless, there is value in understanding the ensuing schemes from the perspective of assumptions. In particular, if we rely on the Jain-Lin-Sahai obfuscation scheme [38], we would require both bilinear groups and LWE similar to the AYW schemes, and would not achieve post-quantum security. If we turn to the post-quantum obfuscation candidates, e.g. [4,32,46,20,24], then we would require hardness or assumptions in the zeroizing regime.

CP-ABE from LWE. The state of the art for CP-ABE from LWE is that of Agrawal and Yamada [6] supporting circuits of depth d and size s over $\{0,1\}^\ell$ with $|\mathsf{ct}| = \mathsf{poly}(d,s)$ and key generation running in time $\mathsf{poly}(\ell,d,\log s)$; this improves upon the "trivial" CP-ABE from LWE based on the KP-ABE for circuits from LWE in [13], where key generation runs in time $\mathsf{poly}(d,s)$. Both of these schemes achieve $|\mathsf{sk}| = \mathsf{poly}(d,\log s)$. We note that the recent CP-ABE for NC1 from LWE in [27] achieves $|\mathsf{ct}|, |\mathsf{sk}| = \mathsf{poly}(s)$. In contrast, the CP-ABE schemes described in Fig 1 achieve $|\mathsf{ct}| = \mathsf{poly}(\ell,d,\log s)$ or $|\mathsf{ct}| = \mathsf{poly}(\ell, 2^d, \log s)$ (i.e., almost independent of circuit size s).

Variant of evasive LWE. In an independent work, Tsabary [43] (also [44]) introduced a variant of evasive LWE, where the matrix $\mathbf{B}$ is secret, and the vector $\mathbf{s}$ is replaced by a matrix $\mathbf{S}$ that could be drawn from an arbitrary distribution. Concretely, fix some efficiently samplable distributions $(\mathbf{S},\mathbf{P},\mathsf{aux})$ over $\mathbb{Z}_q^{n'\times n} \times \mathbb{Z}_q^{n\times t} \times \{0,1\}^*$. The assumption (following the exposition in [44]) asserts that

$$\textbf{if} \qquad (\boxed{\mathbf{SB}+\mathbf{E}}, \boxed{\mathbf{SP}+\mathbf{E}'}, \mathsf{aux}) \approx_c (\mathbf{C}, \mathbf{C}', \mathsf{aux}), \tag{7}$$

$$\textbf{then} \qquad (\boxed{\mathbf{SB}+\mathbf{E}}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{C}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \tag{8}$$

where $\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{C} \leftarrow \mathbb{Z}_q^{n'\times m}$ are uniformly random and $\mathbf{E}'$ is a fresh noise matrix of sufficiently larger magnitude than $\mathbf{E}$. Think of parameters $O(n\log q) \le m \le t$, so that $\mathbf{P}$ is wider than $\mathbf{B}$.

- The formulation of evasive LWE in [43] additionally allows $(\mathbf{P},\mathsf{aux})$ to depend on $\mathbf{B}$, whereas ours and that in [44] does not. In particular, our formulation of evasive LWE is more conservative.
- The use of evasive LWE in [43,44] for witness encryption requires that the assumption holds for general, private-coin auxiliary input, which is unlikely to hold in its completely full generality, e.g. for highly contrived aux that contains a carefully crafted obfuscated program (containing a trapdoor for $\mathbf{P}$). In contrast, the application in this work only requires evasive LWE to hold for public-coin auxiliary input.

SIS analogue of evasive LWE. Finally, we can consider *evasive SIS* —the SIS analogue of evasive LWE— which asserts that if SIS is hard for $[\mathbf{B} \mid \mathbf{P}]$ (given aux), then SIS is hard for $\mathbf{B}$, even given $\mathbf{B}^{-1}(\mathbf{P})$ (and aux). We believe evasive SIS constitutes a useful heuristic for understanding non-standard variants of SIS such as those put forth in recent follow-up works [8,47].

3 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 $v \leftarrow \mathcal{D}$ to denote a random sample from a distribution $\mathcal{D}$, 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

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

A useful corollary of the mixed-product property says that for any pair of row vectors $\mathbf{u}, \mathbf{v} \in \mathbb{Z}^n$,

$$\begin{aligned} \mathbf{u} \otimes \mathbf{v} &= (\mathbf{u} \otimes 1)(\mathbf{I}_n \otimes \mathbf{v}) = (1 \otimes \mathbf{v})(\mathbf{u} \otimes \mathbf{I}_n) \\ &= \mathbf{u}(\mathbf{I}_n \otimes \mathbf{v}) = \mathbf{v}(\mathbf{u} \otimes \mathbf{I}_n) \end{aligned}$$

We adopt the convention that matrix multiplication takes precedence over tensor product, so that we can write $\mathbf{A} \otimes \mathbf{BC}$ to mean $\mathbf{A} \otimes (\mathbf{BC})$.

3.1 Lattices background

We use $\mathcal{D}_{\mathbb{Z},\chi}$ 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

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

where

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

Trapdoor and preimage sampling. 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 $\mathcal{D}_{\mathbb{Z}^{m \times n'},\sigma}$ conditioned on $\mathbf{BY} = \mathbf{Z} \pmod q$. We sometimes suppress σ when the context is clear.

There is a p.p.t. algorithm $\mathsf{TrapGen}(1^n, q)$ that, given the modulus $q \geq 2$ and dimension n, outputs $\mathbf{B} \approx_s U(\mathbb{Z}_q^{n \times 2n\log q})$ with a trapdoor τ. Moreover, there is a p.p.t. algorithm that given $(\mathbf{B}, \tau) \leftarrow \mathsf{TrapGen}(1^n, q)$, $\mathbf{Z} \in \mathbb{Z}_q^{n \times n'}$, and $\sigma \geq 2\sqrt{n \log q}$, outputs a sample from $\mathbf{B}^{-1}(\mathbf{Z}, \sigma)$.

3.2 Attribute-based encryption

Syntax. A ciphertext-policy attribute-based encryption (CP-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.

$\mathsf{Enc}(\mathsf{mpk}, f, \mu) \to \mathsf{ct}_f$. The encryption algorithm gets as input mpk, $f \in \mathcal{F}$ and a message $\mu \in \{0,1\}$. It outputs a ciphertext ct_f.

$\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, x) \to \mathsf{sk}_x$. The key generation algorithm gets as input mpk, msk and $x \in \{0,1\}^\ell$. It outputs a secret key sk_x.

$\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}_x, \mathsf{ct}_f) \to m$. The decryption algorithm gets as input sk_x and ct_f such that $f(x) = 0$ along with mpk. It outputs a message μ.

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

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

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

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

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

$$\mathsf{Adv}^{\mathrm{ABE}}_{\mathcal{A}}(\lambda) := \Pr\left[b = b' : \begin{array}{c} (f, x_1, \ldots, x_Q) \leftarrow \mathcal{A}(1^\lambda) \\ (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathcal{F}) \\ \mathsf{sk}_i \leftarrow \mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, x_i), i = 1, \ldots, Q \\ (\mu_0, \mu_1) \leftarrow \mathcal{A}(\mathsf{mpk}, \mathsf{sk}_1, \ldots, \mathsf{sk}_Q) \\ b \leftarrow \{0,1\};\ \mathsf{ct}_f \leftarrow \mathsf{Enc}(\mathsf{mpk}, f, \mu_b) \\ b' \leftarrow \mathcal{A}(\mathsf{ct}_f)\end{array}\right] - \frac{1}{2}$$

Broadcast encryption. Here,

$$\mathcal{X} = \{0,1\}^N, \mathcal{Y} = [N]$$

where we think of $\{0,1\}^N$ as the power set of $[N]$ (i.e., set of all subsets of $[N]$), and

$$\mathsf{P}(S, y) = 1 \iff y \in S$$

As noted in [7,5], very selective security for broadcast encryption implies selective security since an adversary can simply ask for all keys outside S.

4 Evasive LWE

We proceed to provide a formal statement of our evasive LWE assumption, stated informally in Section 1.

Evasive LWE. Let Samp be a PPT algorithm that on input 1^λ, outputs

$$\mathbf{A}' \in \mathbb{Z}_q^{n \times m'}, \mathbf{P} \in \mathbb{Z}_q^{n \times t}, \mathsf{aux} \in \{0,1\}^*$$

We define the following advantage functions:

$$\begin{aligned}\mathsf{Adv}^{\mathrm{PRE}}_{\mathcal{A}_0}(\lambda) := \Pr[\mathcal{A}_0(\boxed{\mathbf{sA}'+\mathbf{e}'},\boxed{\mathbf{sB}+\mathbf{e}},\boxed{\mathbf{sP}+\mathbf{e}''},\mathbf{A}',\mathbf{B},\mathsf{aux})=1]\\ -\Pr[\mathcal{A}_0(\mathbf{c},\mathbf{c}_0,\mathbf{c}',\mathbf{A}',\mathbf{B},\mathsf{aux})=1], \qquad (9)\\ \mathsf{Adv}^{\mathrm{POST}}_{\mathcal{A}_1}(\lambda) := \Pr[\mathcal{A}_1(\boxed{\mathbf{sA}'+\mathbf{e}'},\boxed{\mathbf{sB}+\mathbf{e}},\mathbf{K},\mathbf{A}',\mathbf{B},\mathsf{aux})=1]\\ -\Pr[\mathcal{A}_1(\mathbf{c},\mathbf{c}_0,\mathbf{K},\mathbf{A}',\mathbf{B},\mathsf{aux})=1] \qquad (10)\end{aligned}$$

where

$$\begin{aligned}&(\mathbf{A}',\mathbf{P},\mathsf{aux})\leftarrow\mathsf{Samp}(1^\lambda)\\ &\mathbf{B}\leftarrow\mathbb{Z}_q^{n\times m},\mathbf{s}\leftarrow\mathbb{Z}_q^n,\\ &\mathbf{c}\leftarrow\mathbb{Z}_q^{m'},\mathbf{c}_0\leftarrow\mathbb{Z}_q^m,\mathbf{c}'\leftarrow\mathbb{Z}_q^t,\\ &\mathbf{e}\leftarrow\mathcal{D}^m_{\mathbb{Z},\chi},\mathbf{e}'\leftarrow\mathcal{D}^{m'}_{\mathbb{Z},\chi},\mathbf{e}''\leftarrow\mathcal{D}^t_{\mathbb{Z},\chi}\\ &\mathbf{K}\leftarrow\mathbf{B}^{-1}(\mathbf{P})\text{ with standard deviation }O(\sqrt{m\log q})\end{aligned}$$

We say that the *evasive LWE* assumption holds if for every PPT Samp, $\mathcal{A}_1$, there exists another PPT $\mathcal{A}_0$ and a polynomial $Q(\cdot)$ such that

$$\mathsf{Adv}^{\mathrm{PRE}}_{\mathcal{A}_0}(\lambda)\geq\mathsf{Adv}^{\mathrm{POST}}_{\mathcal{A}_1}(\lambda)/Q(\lambda)-\mathsf{negl}(\lambda)$$

We consider parameter settings for which $\mathsf{LWE}_{n,q,\chi}$ holds.

Remark 1 (restricted samplers). As in [5], we only require that the assumption holds for samplers where aux additionally contains all of the coin tosses used by Samp. This avoids obfuscation-based counter-examples where aux contains an obfuscation of a program related to a trapdoor for matrix $\mathbf{P}$.

Remark 2 (noise magnitudes). For simplicity, we stated the assumption with all the LWE error terms $\mathbf{e},\mathbf{e}',\mathbf{e}''$ having the same Gaussian parameter χ. It is straight-forward to adapt the assumption and the scheme to a quantitatively weaker variant where the error terms in the post-condition (10) have a larger Gaussian parameter than those in the pre-condition.

Remark 3 (weaker pseudorandomness). For the security of our scheme, it suffices to consider a weaker variant of the assumption where only $\mathbf{sA}'+\mathbf{e}'$ is required to be pseudorandom in the post-condition.

We refer to Section 6 for further discussion on the assumption.

5 Main Constructions

In this section, we present our main constructions:

- a CP-ABE scheme for NC^1 achieving $|\mathsf{ct}|=\mathsf{poly}(2^d,\log s,\lambda)$;
- an "optimal" broadcast encryption scheme for N users with $|\mathsf{mpk}|+|\mathsf{ct}|+|\mathsf{sk}|=\mathsf{poly}(\log N,\lambda)$;
- a CP-ABE scheme for circuits achieving $|\mathsf{ct}|=\mathsf{poly}(d,\log s,\lambda)$;

The first scheme serves as the basis for the second and the third scheme. The first two schemes rely on evasive LWE whereas the third requires an additional "tensor LWE" assumption. We prove very selective security for all three schemes, which implies selective security for broadcast encryption.

5.1 Homomorphic Computation on Matrices

We recall basic homomorphic computation on matrices used in prior LWE-based ABE [13].

Lemma 1 ($\mathsf{EvalF}_{\mathbf{G}}, \mathsf{EvalFX}_{\mathbf{G}}$)**.** *Fix parameters n, q, ℓ and $m = O(n \log q)$. Given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$ and a circuit $f : \{0,1\}^\ell \to \{0,1\}$ of depth d and size s, we can efficiently compute a matrix $\mathbf{A}_f \in \mathbb{Z}_q^{n \times m}$ such that for all $\mathbf{x} \in \{0,1\}^\ell$, there exists a matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \in \mathbb{Z}^{\ell m \times m}$ with $|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}| = m^{O(d)} \cdot s$ such that*

$$(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(\mathbf{x})\mathbf{G} \tag{11}$$

where $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ is the gadget matrix. Moreover, $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ is efficiently computable given $\mathbf{A}, f, \mathbf{x}$. We use $\mathsf{EvalF}_{\mathbf{G}}(\mathbf{A}, f), \mathsf{EvalFX}_{\mathbf{G}}(\mathbf{A}, f, \mathbf{x})$ to denote the algorithms computing $\mathbf{A}_f, \mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ respectively.

Low-norm variant. We also consider a variant where $\mathbf{A}$ has low-norm and we replace $\mathbf{G}$ with $\mathbf{I}$: when deriving $\mathbf{A}_f$, addition gates correspond to matrix addition and multiplication gates correspond to matrix multiplication.[7] The magnitude of the noise squares with each multiplication gate, leading to noise growth that is doubly exponential in d.

Lemma 2 ($\mathsf{EvalF}, \mathsf{EvalFX}$)**.** *Fix parameters m, ℓ. Given a matrix $\mathbf{A} \in \mathbb{Z}^{m \times \ell m}$ and a circuit $f : \{0,1\}^\ell \to \{0,1\}$ of depth d and size s, we can efficiently compute a matrix $\mathbf{A}_f \in \mathbb{Z}^{m \times m}$ such that $|\mathbf{A}_f| = (|\mathbf{A}|m)^{O(2^d)} \cdot s$ and for all $\mathbf{x} \in \{0,1\}^\ell$, there exists a matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \in \mathbb{Z}^{\ell m \times m}$ with $|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}| = (|\mathbf{A}|m)^{O(2^d)} \cdot s$ such that*

$$(\mathbf{A} - \mathbf{x} \otimes \mathbf{I}_m) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f - f(\mathbf{x})\mathbf{I}_m \tag{12}$$

Moreover, $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ is efficiently computable given $\mathbf{A}, f, \mathbf{x}$. We use $\mathsf{EvalF}(\mathbf{A}, f), \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x})$ to denote the algorithms computing $\mathbf{A}_f, \mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ respectively.

5.2 CP-ABE for $\mathbf{NC^1}$ Circuits

We present our CP-ABE scheme for NC^1 circuits.

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

$$(\mathbf{B}, \tau) \leftarrow \mathsf{TrapGen}(1^{n+m^2}, q),\ \mathbf{A}_0 \leftarrow \mathbb{Z}_q^{n \times m},\ \mathbf{A} \leftarrow \mathcal{D}_{\mathbb{Z},\chi''}^{m \times \ell m}, \mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z},\chi''}^{m}$$

  Output

$$\mathsf{mpk} := \left(\mathbf{B}, \mathbf{A}_0, \mathbf{A}, \mathbf{u}^\top\right), \quad \mathsf{msk} := \tau$$

- $\mathsf{Enc}(\mathsf{mpk}, f, \mu \in \{0,1\})$. Compute $\mathbf{A}_f = \mathsf{EvalF}(\mathbf{A}, f)$. Sample

$$\mathbf{s}_0 \leftarrow \mathbb{Z}_q^n, \mathbf{s}_1 \leftarrow \mathbb{Z}_q^{m^2},\ \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{m}, \mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{O((n+m^2)\log q)},$$

  Output

$$\mathsf{ct}_f := \Big(\overbrace{(\mathbf{s}_0 \mid \mathbf{s}_1)\mathbf{B} + \mathbf{e}_0}^{\mathbf{c}_0}, \overbrace{\mathbf{s}_0\mathbf{A}_0 + \mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I}_m) + \mu \cdot \mathbf{g} + \mathbf{e}}^{\mathbf{c}}\Big)$$

- $\mathsf{KeyGen}(\mathsf{msk}, \mathbf{x})$: Sample

$$\mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z},\chi''}^{m},\ \mathbf{K} \leftarrow \mathbf{B}^{-1}\begin{pmatrix} \mathbf{A}_0\mathbf{r}^\top \\ (\mathbf{A} - \mathbf{x} \otimes \mathbf{I}_m) \otimes \mathbf{r}^\top \end{pmatrix},$$

  using τ with standard deviation $O(\sqrt{(n+m^2)\log q})$. Output

$$\mathsf{sk}_{\mathbf{x}} := \left(\mathbf{K}, \mathbf{r}^\top\right)$$

- $\mathsf{Dec}(\mathsf{sk}, \mathbf{x}, \mathsf{ct}, f)$: Compute $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathsf{EvalFX}(\mathbf{A}, f, \mathbf{x})$. Output

$$\mathsf{round}_{\beta_0}\left(\mathbf{c} \cdot \mathbf{r}^\top - \mathbf{c}_0 \cdot \mathbf{K} \cdot \begin{pmatrix} 1 \\ \mathbf{K}_1 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top \end{pmatrix}\right)$$

  where $\mathsf{round}_{\beta_0}(x)$ outputs 0 if $|x| < \beta_0$ and 1 otherwise.

[7] That is, $x_i + x_j$ corresponds to $\mathbf{A}_i + \mathbf{A}_j$ and $x_i \cdot x_j$ corresponds to $\mathbf{A}_i \cdot \mathbf{A}_j$ instead of $\mathbf{A}_i \cdot \mathbf{G}^{-1}(\mathbf{A}_j)$. More generally, we can represent a circuit f of depth d and size s as a polynomial comprising the sum of s monomials, each of total degree at most 2^d. Then, $\mathbf{A}_f = f(\mathbf{A}_1, \ldots, \mathbf{A}_\ell)$.

	$\mathbf{c}'$	$\mathbf{c}'_{1,i}$	$c'_{0,i}$
H_0	$\mathbf{s}_0\mathbf{A}_0 + \mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top \otimes \mathbf{I}_m)$	$\mathbf{s}_1((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)$	$\mathbf{s}_0\mathbf{A}_0\mathbf{r}_i^\top$
H_1	$\downarrow$	$\downarrow$	$\mathbf{c}'\cdot\mathbf{r}_i^\top - \mathbf{c}'_1\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top - \mathbf{s}_1(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)$
H_2	$\mathbf{c}' \leftarrow \mathbb{Z}_q^m$	$\downarrow$	$\downarrow$
H_3	$\downarrow$	$(\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)+\mathbf{e}'_i)\cdot(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)$	$\mathbf{c}'\cdot\mathbf{r}_i^\top - \mathbf{c}'_1\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top - (\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)+\mathbf{e}'_i)\cdot\mathbf{u}^\top$
H_4	$\downarrow$	$\mathbf{s}_i(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)$	$\mathbf{c}'\cdot\mathbf{r}_i^\top - \mathbf{c}'_1\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top - \mathbf{s}_i\mathbf{u}^\top$
H_5	$\downarrow$	$\mathbf{c}'_{1,i} \leftarrow \mathbb{Z}_q^{\ell m}$	$c'_{0,i} \leftarrow \mathbb{Z}_q$

Fig. 2. Summary of the hybrid sequence, with $\mathsf{H}_0 \approx_s \mathsf{H}_1 \approx_c \mathsf{H}_2 \approx_s \mathsf{H}_3 \approx_c \mathsf{H}_4 \approx_c \mathsf{H}_5$. We suppress the additive noise terms in $\mathbf{c}', \mathbf{c}'_{1,i}, \mathbf{c}'_{0,i}$; $\downarrow$ denotes same as previous hybrid; we sample $\mathbf{e}'_i \leftarrow \mathcal{D}^m_{\mathbb{Z},\chi''}$ in H_3 and $\mathbf{s}_i \leftarrow \mathbb{Z}_q^m$ in H_4.

Parameters. Recall $|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}|$ is bounded by $\beta = (\chi'' m)^{O(2^d)} \cdot s$. We set

$$n = \mathsf{poly}(\lambda, \log\beta), \quad m = O(n\log q), \quad \chi'', \chi' = \lambda^{\omega(1)}, \quad \chi = \beta\cdot\lambda^{\omega(1)},$$
$$\beta_0 = \chi^2\cdot\chi''\cdot\beta\cdot\mathsf{poly}(m), \quad q = \beta_0\cdot\lambda^{\omega(1)}$$

In particular, this means

$$|\mathsf{mpk}| = \ell\cdot\mathsf{poly}(2^d, s, \lambda), \quad |\mathsf{ct}| = \mathsf{poly}(2^d, s, \lambda), \quad |\mathsf{sk}| = \ell\cdot\mathsf{poly}(2^d, s, \lambda)$$

Correctness. Fix $\mathbf{x}, f$ such that $f(\mathbf{x}) = 0$. First, we have

$$\begin{aligned}
\mathbf{c}_0\cdot\mathbf{K}\cdot\begin{pmatrix}1\\ \mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top\end{pmatrix} &\approx (\mathbf{s}_0 \mid \mathbf{s}_1)\mathbf{B}\cdot\mathbf{B}^{-1}\begin{pmatrix}\mathbf{A}_0\mathbf{r}^\top & \\ & (\mathbf{A}-\mathbf{x}\otimes\mathbf{I}_m)\otimes\mathbf{r}^\top\end{pmatrix}\cdot\begin{pmatrix}1\\ \mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top\end{pmatrix}\\
&= \mathbf{s}_0\mathbf{A}_0\mathbf{r}^\top + \mathbf{s}_1((\mathbf{A}-\mathbf{x}\otimes\mathbf{I}_m)\otimes\mathbf{r}^\top)\cdot(\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top\otimes 1)\\
&= \mathbf{s}_0\mathbf{A}_0\mathbf{r}^\top + \mathbf{s}_1((\mathbf{A}-\mathbf{x}\otimes\mathbf{I}_m)\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top\otimes\mathbf{I}_m)\cdot(1\otimes\mathbf{r}^\top)\\
&= \mathbf{s}_0\mathbf{A}_0\mathbf{r}^\top + \mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m)\cdot\mathbf{r}^\top
\end{aligned}$$

where the final equality uses $(\mathbf{A}-\mathbf{x}\otimes\mathbf{I}_m)\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}} = \mathbf{A}_f$. This means

$$\begin{aligned}
&\mathbf{c}\cdot\mathbf{r}^\top - \mathbf{c}_0\cdot\mathbf{K}\cdot\begin{pmatrix}1\\ \mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top\end{pmatrix}\\
&\approx (\mathbf{s}_0\mathbf{A}_0 + \mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m) + \mu\cdot\mathbf{g})\cdot\mathbf{r}^\top - \mathbf{s}_0\mathbf{A}_0\mathbf{r}^\top - \mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m)\cdot\mathbf{r}^\top\\
&= \mu\cdot\mathbf{g}\cdot\mathbf{r}^\top
\end{aligned}$$

In particular, the error term is bounded by

$$|\mathbf{e}\cdot\mathbf{r}^\top| + |\mathbf{e}_0\cdot\mathbf{K}\cdot\begin{pmatrix}1\\ \mathbf{H}_{\mathbf{A},f,\mathbf{x}}\mathbf{u}^\top\end{pmatrix}| \le \chi^2\cdot\chi''\cdot\beta\cdot\mathsf{poly}(m) \le \beta_0$$

Now, $\mathbf{g}\cdot\mathbf{r}^\top$ is statistically close to uniform over $\mathbb{Z}_q$, and correctness follows as long as $q \ge \beta_0\cdot\lambda^{\omega(1)}$.

Security. Suppose the (very selective) ABE adversary $\mathcal{A}$ with randomness $\mathsf{coins}_{\mathcal{A}}$ queries f and $\mathbf{x}_1,\ldots,\mathbf{x}_Q$ such that $f(\mathbf{x}_1)=\cdots=f(\mathbf{x}_Q)=1$. We invoke our evasive LWE hardness assumption with Gaussian parameter χ and the following sampler Samp:

$$\begin{aligned}
\mathsf{aux} &= (\overbrace{\mathbf{x}_1,\ldots,\mathbf{x}_Q,f,\mathsf{coins}_{\mathcal{A}}}^{\mathsf{aux}_0},\mathbf{r}_1^\top,\ldots,\mathbf{r}_Q^\top,\mathbf{A}_0,\mathbf{A},\mathbf{u}^\top)\\
\mathbf{P}_0 &= \mathbf{A}_0[\mathbf{r}_1^\top \mid \cdots \mid \mathbf{r}_Q^\top]\\
\mathbf{P}_1 &= [(\mathbf{A}-\mathbf{x}_1\otimes\mathbf{I}_m)\otimes\mathbf{r}_1^\top \mid \cdots \mid (\mathbf{A}-\mathbf{x}_Q\otimes\mathbf{I}_m)\otimes\mathbf{r}_Q^\top]\\
\mathbf{P} &= \begin{pmatrix}\mathbf{P}_0 & \\ & \mathbf{P}_1\end{pmatrix}\\
\mathbf{A}' &= \begin{pmatrix}\mathbf{A}_0\\ \mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m\end{pmatrix}
\end{aligned}$$

where $\mathbf{r}_1^\top,\ldots,\mathbf{r}_Q^\top,\mathbf{A}_0,\mathbf{A},\mathbf{u}^\top$ are sampled as in our CP-ABE scheme. Note that Samp satisfies the restriction in Remark 1. At this point, it suffices to show pseudorandomness of

$$\mathsf{aux}_0,\mathbf{A}_0,\mathbf{A},\mathbf{u}^\top,\mathbf{B},\overbrace{\mathbf{s}_0\mathbf{A}_0+\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m)+\mathbf{e}'}^{\approx(\mathbf{s}_0|\mathbf{s}_1)\mathbf{A}'},(\mathbf{s}_0\mid\mathbf{s}_1)\mathbf{B}+\mathbf{e}_0,\{\overbrace{\mathbf{s}_1((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)+\mathbf{e}'_{1,i}}^{\approx(\mathbf{s}_0|\mathbf{s}_1)\mathbf{P}}\mathbf{s}_0\mathbf{A}_0\mathbf{r}_i^\top+e'_{0,i},\mathbf{r}_i^\top\}_{i\in[Q]} \tag{13}$$

where $\mathbf{A}_0,\mathbf{u}^\top,\mathbf{s}_0,\mathbf{s}_1,\mathbf{r}_i^\top$ are also sampled as in the CP-ABE scheme and $\mathbf{e}'\leftarrow\mathcal{D}^m_{\mathbb{Z},\chi''},\mathbf{e}'_{1,i}\leftarrow\mathcal{D}^{\ell m}_{\mathbb{Z},\chi''},e'_{0,i}\leftarrow\mathcal{D}_{\mathbb{Z},\chi}$. By noise flooding, this implies pseudorandomness of (13) when $\mathbf{e}'\leftarrow\mathcal{D}^m_{\mathbb{Z},\chi},\mathbf{e}'_{1,i}\leftarrow\mathcal{D}^{\ell m}_{\mathbb{Z},\chi}$. Combined with our hardness assumption, the latter would imply:

$$\begin{aligned}
&\overbrace{\mathbf{A}_0,\mathbf{A},\mathbf{u}^\top,\mathbf{B}}^{\mathsf{mpk}},\ \overbrace{(\mathbf{s}_0\mid\mathbf{s}_1)\mathbf{B}+\mathbf{e}_0,\underbrace{\boxed{\mathbf{s}_0\mathbf{A}_0+\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m)+\mathbf{e}}}_{\mathsf{sk}_i}+\mu\cdot\mathbf{g}}^{\mathsf{ct}},\\
&\qquad\qquad\{\overbrace{\mathbf{B}^{-1}\begin{pmatrix}\mathbf{A}_0\mathbf{r}_i^\top\\ (\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top\end{pmatrix},\mathbf{r}_i^\top)}^{}\}_{i\in[Q]}\\
\approx_c\ &\mathbf{A}_0,\mathbf{A},\mathbf{u}^\top,\mathbf{B},\ \mathbf{c}_0,\boxed{\mathbf{c}}+\mu\cdot\mathbf{g},\\
&\qquad\qquad\{\mathbf{B}^{-1}\begin{pmatrix}\mathbf{A}_0\mathbf{r}_i^\top\\ (\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top\end{pmatrix},\mathbf{r}_i^\top)\}_{i\in[Q]}
\end{aligned}$$

(where the distinguisher additionally gets aux_0) from which ABE security follows readily. Next, we prove pseudorandomness of (13) from LWE via a hybrid sequence summarized in Fig 2:

- H_0: the distribution in (13)

$$\mathsf{aux}_0,\mathbf{A}_0,\mathbf{A},\mathbf{u}^\top,\mathbf{B},\overbrace{\mathbf{s}_0\mathbf{A}_0+\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m)+\mathbf{e}'}^{\mathbf{c}'},\overbrace{(\mathbf{s}_0\mid\mathbf{s}_1)\mathbf{B}+\mathbf{e}_0}^{\mathbf{c}_0}\{\overbrace{\mathbf{s}_1((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)+\mathbf{e}'_{1,i}}^{\mathbf{c}'_{1,i}},\overbrace{\mathbf{s}_0\mathbf{A}_0\mathbf{r}_i^\top+e'_{0,i}}^{c'_{0,i}},\mathbf{r}_i^\top\}_{i\in[Q]}$$

- H_1: same as H_0, except we compute

$$c'_{0,i}:=\boxed{\mathbf{c}'\cdot\mathbf{r}_i^\top-\mathbf{c}'_1\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top-\mathbf{s}_1(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)}+e'_{0,i}.$$

We claim that $\mathsf{H}_0\approx_s\mathsf{H}_1$. First, observe that

$$\begin{aligned}
\mathbf{s}_0\mathbf{A}_0\mathbf{r}_i^\top &= \overbrace{(\mathbf{s}_0\mathbf{A}_0+\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m))}^{\approx\mathbf{c}'}\cdot\mathbf{r}_i^\top-\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{r}_i^\top)\\
&= (\mathbf{s}_0\mathbf{A}_0+\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m))\cdot\mathbf{r}_i^\top-\mathbf{s}_1((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top+\mathbf{u}^\top)\otimes\mathbf{r}_i^\top)\\
&= (\mathbf{s}_0\mathbf{A}_0+\mathbf{s}_1(\mathbf{A}_f\mathbf{u}^\top\otimes\mathbf{I}_m))\cdot\mathbf{r}_i^\top-\overbrace{\mathbf{s}_1((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)}^{\approx\mathbf{c}'_{1,i}}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top+\mathbf{s}_1(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)
\end{aligned}$$

where the first and third equalities uses the mixed-product property, and the second equality uses $(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}=\mathbf{A}_f-f(\mathbf{x}_i)\mathbf{I}_m$ and $f(\mathbf{x}_i)=1$. Then, $\mathsf{H}_0\approx_s\mathsf{H}_1$ follows from noise flooding using $e'_{0,i}$, namely

$$e'_{0,i}\approx_s e'_{0,i}+\mathbf{e}'\cdot\mathbf{r}_i^\top-\mathbf{e}'_{1,i}\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top$$

which in turn follows from $\chi\geq\chi'\cdot\beta\cdot\lambda^{\omega(1)}$.

– H_2: same as H_1, except we sample $\mathbf{c}'\leftarrow\mathbb{Z}_q^m,\mathbf{c}_0\leftarrow\mathbb{Z}_q^{O((n+m^2)\log q)}$. We have $\mathsf{H}_1\approx_c\mathsf{H}_2$, since

$$(\mathbf{A}_0,\mathbf{B}_0,\mathbf{s}_0\mathbf{A}_0+\mathbf{e}',\mathbf{s}_0\mathbf{B}_0+\mathbf{e}_0,)\approx_c(\mathbf{A}_0,\mathbf{c}',\mathbf{c}_0),\ \mathbf{c}'\leftarrow\mathbb{Z}_q^m,\mathbf{c}_0\leftarrow\mathbb{Z}_q^{O((n+m^2)\log q)},\mathbf{B}_0\leftarrow\mathbb{Z}_q^{n\times O((n+m^2)\log q)}$$

via the LWE assumption. (In the reduction, $\mathbf{B}_0$ corresponds to the top n rows of $\mathbf{B}$.)

– H_3: same as H_2, except we compute

$$\begin{aligned}\mathbf{s}_i&:=\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)+\mathbf{e}'_i,\quad \mathbf{e}'_i\leftarrow\mathcal{D}^m_{\mathbb{Z},\chi''}\\ \mathbf{c}'_{1,i}&:=\mathbf{s}_i(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)+\mathbf{e}'_{1,i}\\ c'_{0,i}&:=\mathbf{c}'\cdot\mathbf{r}_i^\top-\mathbf{c}'_1\cdot\mathbf{H}_{\mathbf{A},f,\mathbf{x}_i}\mathbf{u}^\top-\mathbf{s}_i\mathbf{u}^\top+e'_{0,i}\end{aligned}$$

We claim that $\mathsf{H}_2\approx_s\mathsf{H}_3$. First, observe that

$$\overbrace{\mathbf{s}_1((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes\mathbf{r}_i^\top)}^{\approx\mathbf{c}'_{1,i}}=\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)((\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\otimes1)=\overbrace{\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)}^{\approx\mathbf{s}_i}(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)$$

$$\mathbf{s}_1(\mathbf{u}^\top\otimes\mathbf{r}_i^\top)=\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)(\mathbf{u}^\top\otimes1)=\overbrace{\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)}^{\approx\mathbf{s}_i}\mathbf{u}^\top$$

where the first equality in each line uses the mixed-product property. Then, $\mathsf{H}_2\approx_s\mathsf{H}_3$ follows from noise flooding using $\mathbf{e}'_{1,i}$ and $e'_{0,i}$, namely

$$\begin{aligned}\mathbf{e}'_{1,i}&\approx_s\mathbf{e}'_{1,i}+\mathbf{e}'_i\cdot(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)\\ e'_{0,i}&\approx_s e'_{0,i}+\mathbf{e}'_i\cdot\mathbf{u}^\top\end{aligned}$$

which in turn follows from $\chi',\chi''\geq\lambda^{\omega(1)}$.

– H_4: same as H_3, except we sample $\mathbf{s}_i\leftarrow\mathbb{Z}_q^m$. We have $\mathsf{H}_3\approx_c\mathsf{H}_4$, since

$$\left\{\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top),\mathbf{r}_i^\top\right\}_{i\in[Q]}\approx_c\left\{\mathbf{s}_i,\mathbf{r}_i^\top\right\}_{i\in[Q]},\ \mathbf{s}_1\leftarrow\mathbb{Z}_q^{m^2},\mathbf{r}_i\leftarrow\mathcal{D}^m_{\mathbb{Z},\chi},\mathbf{s}_i\leftarrow\mathbb{Z}_q^m$$

via the LWE assumption [15,22]. In particular, if we write $\mathbf{s}_1=(\mathbf{s}_{1,1},\ldots,\mathbf{s}_{1,m})$ where $\mathbf{s}_{1,1},\ldots,\mathbf{s}_{1,m}\in\mathbb{Z}_q^m$, then $\mathbf{s}_1(\mathbf{I}_m\otimes\mathbf{r}_i^\top)=(\mathbf{s}_{1,1}\mathbf{r}_i^\top,\ldots,\mathbf{s}_{1,m}\mathbf{r}_i^\top)$.

– H_5 : same as H_4, except we sample $c'_{0,i}\leftarrow\mathbb{Z}_q^{\ell m},\mathbf{c}'_{1,i}\leftarrow\mathbb{Z}_q^{\ell m}$. We have $\mathsf{H}_4\approx_c\mathsf{H}_5$. This follows from a hybrid argument over $i=1,\ldots,Q$, where in the i'th step, we switch the distribution of $c'_{0,i},\mathbf{c}'_{1,i}$ to random via:

$$\begin{aligned}&(\mathbf{A}, && \mathbf{u}^\top, && \mathbf{s}_i\mathbf{u}^\top+e'_{0,i}, && \mathbf{s}_i(\mathbf{A}-\mathbf{x}_i\otimes\mathbf{I}_m)+\mathbf{e}'_{1,i})\\ \approx_s\ &(\mathbf{A}+\mathbf{x}_i\otimes\mathbf{I}_m, && \mathbf{u}^\top, && \mathbf{s}_i\mathbf{u}^\top+e'_{0,i}, && \mathbf{s}_i\mathbf{A}+\mathbf{e}'_{1,i})\\ \approx_c\ &(\mathbf{A}+\mathbf{x}_i\otimes\mathbf{I}_m, && \mathbf{u}^\top, && c'_{0,i}, && \mathbf{c}'_{1,i})\\ \approx_s\ &(\mathbf{A}, && \mathbf{u}^\top, && c'_{0,i}, && \mathbf{c}'_{1,i})\end{aligned}$$

In particular,

- the first and last $\approx_s$ rely on noise flooding with $\mathbf{A}\approx_s\mathbf{A}+\mathbf{x}_i\otimes\mathbf{I}_m$, which in turn follows from $\chi''\geq\lambda^{\omega(1)}$;
- the $\approx_c$ relies on LWE [15] which tells us

$$(\mathbf{A},\mathbf{u}^\top,\mathbf{s}_i\mathbf{A}+\mathbf{e}'_{1,i},\mathbf{s}_i\mathbf{u}^\top+e'_{0,i})\approx_c(\mathbf{A},\mathbf{u}^\top,\text{random})$$

and where the reduction samples $\mathbf{s}_{i+1},\ldots,\mathbf{s}_Q$ and computes $c'_{0,i+1},\ldots,c'_{0,Q},\mathbf{c}'_{1,i+1},\ldots,\mathbf{c}'_{1,Q}$ as in H_5.

5.3 Optimal Broadcast Encryption

To handle broadcast encryption with N users, we identify a user $x \in [N]$ with a bit string $\mathbf{x} \in \{0,1\}^{\lceil \log N \rceil}$. Let $I_{\mathbf{y}}(\cdot)$ be the point function wrt $\mathbf{y}$, that is,

$$I_{\mathbf{y}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} = \mathbf{y} \\ 0 & \text{otherwise} \end{cases}$$

We can then associate each set $S \subseteq [N]$ with the circuit $f_S : \{0,1\}^{\lceil \log N \rceil} \to \{0,1\}$ given by

$$\mathbf{x} \mapsto 1 - \sum_{\mathbf{y} \in S} I_{\mathbf{y}}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in S \\ 1 & \text{if } \mathbf{x} \notin S \end{cases}$$

It is easy to see that f_S can be computed by a circuit of depth $O(\log\log N)$ and size $O(N \log N)$:[8]

- each $I_{\mathbf{y}}(\cdot)$ can be computed by a circuit of depth $O(\log\log N)$ and size $O(\log N)$;
- followed by an addition gate with fan-in N.

We can then instantiate our CP-ABE scheme with $\beta = \lambda^{\mathsf{poly}(\log N)} \cdot N \log N$ (via the bound in Lemma 2) which yields a broadcast encryption scheme with

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

5.4 CP-ABE for Polynomial-Depth Circuits

Tensor LWE. We introduce an additional tensor LWE assumption which states that for all $\mathbf{x}_1, \ldots, \mathbf{x}_Q \in \{0,1\}^\ell$, we have

$$\mathbf{A}, \left\{ \mathbf{s}(\mathbf{I}_m \otimes \mathbf{r}_i^\top)(\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}) + \mathbf{e}_i, \mathbf{r}_i^\top \right\}_{i \in [Q]} \approx_c \mathbf{A}, \left\{ \mathbf{c}_i, \mathbf{r}_i^\top \right\}_{i \in [Q]}$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times \ell m}, \mathbf{s} \leftarrow \mathbb{Z}_q^{mn}, \mathbf{e}_i \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{\ell m}, \mathbf{r}_i^\top \leftarrow \mathcal{D}_{\mathbb{Z},\chi}^{m}, \mathbf{c}_i \leftarrow \mathbb{Z}_q^{\ell m}$. We consider the same parameter settings as $\mathsf{LWE}_{n,q,\chi}$, with $\ell, Q = \mathsf{poly}(\lambda)$. Our analysis in Section 5.2 shows that if we use a low-norm $\mathbf{A}$ and replace $\mathbf{G}$ with $\mathbf{I}$, then LWE implies tensor LWE.

CP-ABE scheme. We modify our CP-ABE scheme in Section 5.2 as follows:

- we sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times \ell m}, (\mathbf{B}_1, \tau_1) \leftarrow \mathsf{TrapGen}(1^{mn}, q), \mathbf{s}_1 \leftarrow \mathbb{Z}_q^{mn}$;
- we replace $\mathbf{I}_m$ in sk with the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ and we replace $\mathsf{EvalF}, \mathsf{EvalFX}$ with $\mathsf{EvalF}_{\mathbf{G}}, \mathsf{EvalFX}_{\mathbf{G}}$ respectively;
- we set $\chi'' = \mathsf{poly}(\lambda)$.

That is, we have:

$$\mathsf{ct}_f := \left((\mathbf{s}_0 \mid \mathbf{s}_1)\mathbf{B} + \mathbf{e}_0,\ \mathbf{s}_0\mathbf{A}_0 + \mathbf{s}_1(\mathbf{A}_f \mathbf{u}^\top \otimes \mathbf{I}_m) + \mu \cdot \mathbf{g} + \mathbf{e} \right)$$

$$\mathsf{sk}_{\mathbf{x}} := \left(\mathbf{B}^{-1} \begin{pmatrix} \mathbf{A}_0 \mathbf{r}^\top \\ (\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}^\top \end{pmatrix}, \mathbf{r}^\top \right)$$

As before, we have: $|\mathsf{mpk}| = \ell \cdot \mathsf{poly}(\log \beta, \lambda), |\mathsf{ct}| = \mathsf{poly}(\log \beta, \lambda), |\mathsf{sk}| = \ell \cdot \mathsf{poly}(\log \beta, \lambda)$. Now, for circuits of depth d and size s, we have $|\mathbf{H}_{\mathbf{A},f,\mathbf{x}}| = \lambda^{O(d)} \cdot s$ so that we can $\beta = \lambda^{O(d)} \cdot s$, which yields:

$$|\mathsf{mpk}| = \ell \cdot \mathsf{poly}(d, \log s, \lambda), \quad |\mathsf{ct}| = \mathsf{poly}(d, \log s, \lambda), \quad |\mathsf{sk}| = \ell \cdot \mathsf{poly}(d, \log s, \lambda)$$

[8] As explained in [13], "To support multiplication and addition of constants, we may assume that we have an extra 0-th input to the circuit that always carries the value 1." That is, we will set $\ell = \lceil \log N \rceil + 1$ in our CP-ABE scheme.

Security. The proof of security requires the following modifications:

- H_1:
$$c'_{0,i} := \boxed{\mathbf{c}' \cdot \mathbf{r}_i^\top - \mathbf{c}'_1 \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}_i} \mathbf{u}^\top - \mathbf{s}_1 (\mathbf{G}\mathbf{u}^\top \otimes \mathbf{r}_i^\top)} + e'_{0,i}.$$
- the proof of $\mathsf{H}_0 \approx_s \mathsf{H}_1$ uses $(\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}) \cdot \mathbf{H}_{\mathbf{A},f,\mathbf{x}_i} = \mathbf{A}_f - f(\mathbf{x}_i)\mathbf{G}$.
- we omit $\mathsf{H}_3, \mathsf{H}_4$ and directly argue that $\mathsf{H}_2 \approx_c \mathsf{H}_5$. Tensor LWE implies that for all $\mathbf{x}_1, \ldots, \mathbf{x}_Q \in \{0,1\}^\ell$,
$$\begin{aligned} &\mathbf{A}, \mathbf{u}^\top, \{\mathbf{r}_i^\top\}_{i\in[Q]}, \{\mathbf{s}_1 (\mathbf{I}_m \otimes \mathbf{r}_i^\top)(\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}) + \mathbf{e}'_{1,i}, \mathbf{s}_1 (\mathbf{I}_m \otimes \mathbf{r}_i^\top)\mathbf{G}\mathbf{u}^\top + e'_{0,i}\}_{i\in[Q]} \\ \approx_c\ &\mathbf{A}, \mathbf{u}^\top, \{\mathbf{r}_i^\top\}_{i\in[Q]}, \{\boxed{\mathbf{c}'_{1,i}}, \qquad\qquad \boxed{c'_{0,i}}\}_{i\in[Q]} \end{aligned}$$
Formally, we account for $\mathbf{u}^\top$ by taking tensor LWE with parameter $\ell + 1$ and padding $\mathbf{x}_1, \ldots, \mathbf{x}_\ell$ with a 0.

6 Discussion on Evasive LWE

Recall our informal statement of evasive LWE: for every efficient samplable distributions over $(\mathbf{A}', \mathbf{P}, \mathsf{aux})$,

if $(\mathbf{A}', \mathbf{B}, \mathbf{P}, \boxed{\mathbf{sA} + \mathbf{e}'}, \boxed{\mathbf{sB} + \mathbf{e}}, \mathbf{sP} + \mathbf{e}'', \mathsf{aux}) \approx_c (\mathbf{A}', \mathbf{B}, \mathbf{P}, \boxed{\mathbf{c}'}, \boxed{\mathbf{c}}, \boxed{\mathbf{c}''}, \mathsf{aux})$,

then $(\mathbf{A}', \mathbf{B}, \boxed{\mathbf{sA} + \mathbf{e}'}, \boxed{\mathbf{sB} + \mathbf{e}}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{A}', \mathbf{B}, \boxed{\mathbf{c}'}, \boxed{\mathbf{c}}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux})$

Examples. We begin with three quick examples:

- if $\mathbf{P}$ is drawn from the uniform distribution over $\mathbb{Z}_q^{n\times t}$, then evasive LWE holds unconditionally, since $\mathbf{B}^{-1}(\mathbf{P})$ is distributed according to a random Gaussian.
- if $\mathbf{P} = \mathbf{0}$, then both the pre and post conditions are false, so evasive LWE is vacuously true.
- if $\mathbf{P} = [\mathbf{U} \mid \mathbf{U}]$ where $\mathbf{U} \leftarrow \mathbb{Z}_q^{n\times t/2}$, then the pre-condition is false, and evasive LWE does not provide any security guarantees. In fact, we know that the post-condition is false if $t \gg n \log q$, since $\mathbf{B}^{-1}(\mathbf{U} \mid \mathbf{U})$ would then leak a basis for $\mathbf{B}$.

As an additional example, suppose $\mathbf{P}$ is a uniformly random block-diagonal matrix, that is, $\mathbf{P} = \begin{pmatrix} \mathbf{U}_0 & \\ & \mathbf{U}_1 \end{pmatrix}$, where $\mathbf{U}_0, \mathbf{U}_1 \leftarrow \mathbb{Z}_q^{n/2\times t/2}$. It is easy to see that the pre-condition holds via LWE, and in this case, we can also show that the post-condition holds assuming LWE. Concretely, let $\mathbf{B}_0, \mathbf{B}_1$ denote the top and bottom halves of the matrix $\mathbf{B}$. Then, $\mathbf{B}^{-1}(\mathbf{P}) \approx_s (\mathbf{B}_1^{-1}(\mathbf{0}), \mathbf{B}_0^{-1}(\mathbf{0}))$ via [24], and the post-condition boils down to showing that $(\mathbf{B}, \mathbf{sB} + \mathbf{e}')$ is pseudorandom given trapdoors for $\mathbf{B}_0, \mathbf{B}_1$. As shown in [21, Theorem 5.3], this follows from LWE, where in the reduction, we sample $\mathbf{B}_0 = [\mathbf{A}_0 \mid \mathbf{A}_0\mathbf{R} + \mathbf{G}], \mathbf{B}_1 = [\mathbf{A}_1 \mid \mathbf{A}_1\mathbf{R} - \mathbf{G}]$, where $\mathbf{R}$ is low-norm.

Algorithmic attacks. The known algorithmic attacks on the post-condition essentially fall into one of two categories:

- Attacks on LWE ignoring $\mathbf{B}^{-1}(\mathbf{P})$: this is ruled out via the pre-condition;
- Attacks computing $\mathbf{c}^* = (\mathbf{sB} + \mathbf{e}') \cdot \mathbf{B}^{-1}(\mathbf{P}) \approx \mathbf{sP}$: suppose given aux, an attacker can find a low-norm $\mathbf{z}$ such that $\mathbf{P} \cdot \mathbf{z}^\top = \mathbf{0}$; we can then use $\mathbf{z}$ to distinguish $\mathbf{sP} + \mathbf{e}''$ from $\mathbf{c}''$, thereby violating the pre-condition. Zeroizing attacks on multi-linear map and obfuscation candidates fall into this category. The attacks on naive approaches to LWE-based ABE via secret-sharing in [2, Section 6] also falls into this category.

Acknowledgments. I would like to thank Yilei Chen and Vinod Vaikuntanathan for numerous illuminating discussions about LWE and zeroizing attacks, and Ivy Woo for helpful feedback on the write-up. Special thanks to Pepita Coffee (@pepitacoffeeco) as well as Shakespeare and Company Café (@shakespeareandcocafeparis), where most of the work was done.

References

1. S. Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 191–225. Springer, Heidelberg, May 2019.
2. S. Agrawal, X. Boyen, V. Vaikuntanathan, P. Voulgaris, and H. Wee. Functional encryption for threshold functions (or fuzzy ibe) from lattices. In M. Fischlin, J. Buchmann, and M. Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 280–297. Springer, Heidelberg, May 2012.
3. S. Agrawal, B. Libert, and D. Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 333–362. Springer, Heidelberg, Aug. 2016.
4. S. Agrawal and A. Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 110–140. Springer, Heidelberg, May 2020.
5. 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.
6. 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.
7. 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.
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. B. Barak, Z. Brakerski, I. Komargodski, and P. K. Kothari. Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation). In J. B. Nielsen and V. Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 649–679. Springer, Heidelberg, Apr. / May 2018.
10. J. Bartusek, J. Guan, F. Ma, and M. Zhandry. Return of GGH15: Provable security against zeroizing attacks. In A. Beimel and S. Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 544–574. Springer, Heidelberg, Nov. 2018.
11. J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy*, pages 321–334. IEEE Computer Society Press, May 2007.
12. W. Beullens and H. Wee. Obfuscating simple functionalities from knowledge assumptions. In D. Lin and K. Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 254–283. Springer, Heidelberg, Apr. 2019.
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. D. Boneh and B. Waters. A fully collusion resistant broadcast, trace, and revoke system. In A. Juels, R. N. Wright, and S. De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 211–220. ACM Press, Oct. / Nov. 2006.
17. D. Boneh, B. Waters, and M. Zhandry. Low overhead broadcast encryption from multilinear maps. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 206–223. Springer, Heidelberg, Aug. 2014.
18. D. Boneh and M. Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Heidelberg, Aug. 2014.
19. Z. Brakerski, N. Döttling, S. Garg, and G. Malavolta. Candidate iO from homomorphic encryption schemes. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 79–109. Springer, Heidelberg, May 2020.
20. Z. Brakerski, N. Döttling, S. Garg, and G. Malavolta. Factoring and pairings are not necessary for io: Circular-secure lwe suffices. Cryptology ePrint Archive, Report 2020/1024, 2020.
21. Z. Brakerski and V. Vaikuntanathan. Lattice-inspired broadcast encryption and succinct ciphertext-policy ABE. In *ITCS*, pages 28:1–28:20, 2022.
22. R. Canetti and Y. Chen. Constraint-hiding constrained PRFs for NC^1 from LWE. In J.-S. Coron and J. B. Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 446–476. Springer, Heidelberg, Apr. / May 2017.
23. J. Chen, R. Gay, and H. Wee. Improved dual system ABE in prime-order groups via predicate encodings. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 595–624. Springer, Heidelberg, Apr. 2015.
24. Y. Chen, V. Vaikuntanathan, and H. Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 577–607. Springer, Heidelberg, Aug. 2018.

25. J. H. Cheon, K. Han, C. Lee, H. Ryu, and D. Stehlé. Cryptanalysis of the multilinear map over the integers. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12. Springer, Heidelberg, Apr. 2015.
26. J.-S. Coron, M. S. Lee, T. Lepoint, and M. Tibouchi. Cryptanalysis of GGH15 multilinear maps. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 607–628. Springer, Heidelberg, Aug. 2016.
27. P. Datta, I. Komargodski, and B. Waters. Decentralized multi-authority ABE for DNFs from LWE. In A. Canteaut and F.-X. Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 177–209. Springer, Heidelberg, Oct. 2021.
28. A. Fiat and M. Naor. Broadcast encryption. In D. R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 480–491. Springer, Heidelberg, Aug. 1994.
29. S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17. Springer, Heidelberg, May 2013.
30. S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, Oct. 2013.
31. S. Garg, E. Miles, P. Mukherjee, A. Sahai, A. Srinivasan, and M. Zhandry. Secure obfuscation in a weak multilinear map model. In M. Hirt and A. D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 241–268. Springer, Heidelberg, Oct. / Nov. 2016.
32. R. Gay and R. Pass. Indistinguishability obfuscation from circular security. In S. Khuller and V. V. Williams, editors, *53rd ACM STOC*, pages 736–749. ACM Press, June 2021.
33. C. Gentry, S. Gorbunov, and S. Halevi. Graph-induced multilinear maps from lattices. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Heidelberg, Mar. 2015.
34. C. Gentry and B. Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In A. Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 171–188. Springer, Heidelberg, Apr. 2009.
35. 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.
36. S. Halevi, T. Halevi, V. Shoup, and N. Stephens-Davidowitz. Implementing BP-obfuscation using graph-induced encoding. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *ACM CCS 2017*, pages 783–798. ACM Press, Oct. / Nov. 2017.
37. S. B. Hopkins, A. Jain, and H. Lin. Counterexamples to new circular security assumptions underlying iO. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 673–700, Virtual Event, Aug. 2021. Springer, Heidelberg.
38. A. Jain, H. Lin, and A. Sahai. Indistinguishability obfuscation from well-founded assumptions. Cryptology ePrint Archive, Report 2020/1003, 2020.
39. S. Ling, D. H. Phan, D. Stehlé, and R. Steinfeld. Hardness of k-LWE and applications in traitor tracing. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Heidelberg, Aug. 2014.
40. A. Lombardi and V. Vaikuntanathan. Limits on the locality of pseudorandom generators and applications to indistinguishability obfuscation. In Y. Kalai and L. Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 119–137. Springer, Heidelberg, Nov. 2017.
41. E. Miles, A. Sahai, and M. Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 629–658. Springer, Heidelberg, Aug. 2016.
42. 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.
43. 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.
44. V. Vaikuntanathan, H. Wee, and D. Wichs. Witness encryption and null-iO from evasive LWE, 2022.
45. H. Wee. Broadcast encryption with size $N^{1/3}$ and more from k-lin. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 155–178, Virtual Event, Aug. 2021. Springer, Heidelberg.
46. H. Wee and D. Wichs. Candidate obfuscation via oblivious LWE sampling. In A. Canteaut and F.-X. Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 127–156. Springer, Heidelberg, Oct. 2021.
47. H. Wee and D. J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In *EUROCRYPT*, 2023.