ABE for Circuits with Constant-Size Secret Keys and Adaptive Security

Hanjun Li, Huijia Lin, and Ji Luo

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

Abstract. An important theme in the research on attribute-based encryption (ABE) is minimizing the sizes of secret keys and ciphertexts. In this work, we present two new ABE schemes with *constant-size* secret keys, i.e., the key size is independent of the sizes of policies or attributes and dependent only on the security parameter λ .

- We construct the first key-policy ABE scheme for circuits with constant-size secret keys, $|\mathbf{sk}_f| = \operatorname{poly}(\lambda)$, which concretely consist of only three group elements. The previous state-of-the-art scheme by [Boneh et al., Eurocrypt '14] has key size polynomial in the maximum depth d of the policy circuits, $|\mathbf{sk}_f| = \operatorname{poly}(d, \lambda)$. Our new scheme removes this dependency of key size on d while keeping the ciphertext size the same, which grows linearly in the attribute length and polynomially in the maximal depth, $|\mathbf{ct}_{\mathbf{x}}| = |\mathbf{x}| \operatorname{poly}(d, \lambda)$.
- We present the first ciphertext-policy ABE scheme for Boolean formulae that simultaneously has constant-size keys and succinct ciphertexts of size independent of the policy formulae, namely, $|\mathbf{sk}_f| = \text{poly}(\lambda)$ and $|\mathbf{ct}_{\mathbf{x}}| = \text{poly}(|\mathbf{x}|, \lambda)$. Concretely, each secret key consists of only two group elements. Previous ciphertext-policy ABE schemes either have succinct ciphertexts but non-constant-size keys [Agrawal–Yamada, Eurocrypt '20, Agrawal–Wichs–Yamada, TCC '20], or constant-size keys but large ciphertexts that grow with the policy size as well as the attribute length. Our second construction is the first ABE scheme achieving *double succinctness*, where both keys and ciphertexts are smaller than the corresponding attributes and policies tied to them.

Our constructions feature new ways of combining lattices with pairing groups for building ABE and are proven selectively secure based on LWE and in the generic (pairing) group model. We further show that when replacing the LWE assumption with its adaptive variant introduced in [Quach–Wee–Wichs FOCS '18], the constructions become adaptively secure.

1 Introduction

Attribute-based encryption (ABE) [24,37] is a novel generalization of public-key encryption for enforcing fine-grained access control. In this work, we focus on

[©] IACR 2022. This article is the final version submitted by the authors to the IACR and to Springer-Verlag on 20 September 2022.

improving the efficiency of ABE schemes, especially on minimizing the sizes of secret keys while keeping ciphertexts small. In key-policy (KP) ABE, a secret key sk_f is tied to a policy f and a ciphertext $\mathsf{ct}_{\mathbf{x}}$ encrypting a message μ is tied to an attribute \mathbf{x} , so that a secret key is only "authorized" to decrypt a ciphertext if the associated attribute \mathbf{x} satisfies the policy f. At first glance, since a secret key specifies the associated policy f, it appears that the size of the secret key would have to depend at least *linearly* on the (description) size of f. Similarly, a ciphertext would have to grow *linearly* with the length of the associated attribute \mathbf{x} . Secret keys and ciphertexts with linear dependency of their sizes on the policies and attributes they are tied to are said to be *compact*, and most ABE schemes are indeed compact.

However, upon closer examination, as ABE does not guarantee privacy of the policies nor the attributes, it is possible to give a description of the policy fin the clear in the secret key, and the *non-trivial* part of the secret key may be smaller than the policy. In this case, the right measure of efficiency should be the size of the non-trivial part (i.e., the overhead), which we now view as *the* secret key. We can now aim for secret keys of size smaller than that of the policy — i.e., $|\mathbf{sk}_f| = \mathbf{o}(|f|)$ — referred to as *succinct* keys, or even keys of size independent of that of the policy — i.e., $|\mathbf{sk}_f| = O(1)$ — referred to as *constantsize* keys.¹ Similarly, *succinct* ciphertexts have size smaller than the length of the attributes, $|\mathbf{ct}_{\mathbf{x}}| = \mathbf{o}(|\mathbf{x}|)$, and *constant-size* ciphertexts satisfy $|\mathbf{ct}_{\mathbf{x}}| = O(1)$. We further examine the efficiency of ciphertext-policy (CP) ABE [10], which enables instead the ciphertexts \mathbf{ct}_f to specify the policies, so that only secret keys $\mathbf{sk}_{\mathbf{x}}$ with attributes satisfying the policies can decrypt them. Naturally, succinct keys and ciphertexts have size $|\mathbf{sk}_{\mathbf{x}}| = \mathbf{o}(|\mathbf{x}|)$ and $|\mathbf{ct}_f| = \mathbf{o}(|f|)$, and constant size means the same as in KP-ABE.

How close can we get to the *ideal efficiency* of having both constant-size keys and ciphertexts? Despite tremendous effort, the state-of-the-art is still far from the ideal. Current ABE schemes with either succinct keys or succinct ciphertexts can be broadly classified as follows (see Figs. 1 and 2):

- The work of [11] built KP-ABE based on LWE for polynomial-size circuits with succinct keys $|\mathsf{sk}_f| = \operatorname{poly}(d)$ and ciphertexts of size $|\mathsf{ct}_{\mathbf{x}}| = |\mathbf{x}| \operatorname{poly}(d)$, where d is the depth of the circuit.
- Several works [5–7,32,38,42,43] constructed KP-ABE and CP-ABE for lowdepth computations with *either* constant-size secret keys *or* constant-size ciphertexts from pairing, i.e., *either* $|\mathbf{sk}| = O(1)$ or $|\mathbf{ct}| = O(1)$, at the cost of the other component being much larger, of size $\Omega(|f| \cdot |\mathbf{x}|)$.
- The recent works of [3,4] constructed CP-ABE for Boolean formulae with succinct ciphertexts $|\mathsf{ct}_f| = \Theta(|\mathbf{x}|)$ and compact keys $|\mathsf{sk}_{\mathbf{x}}| = \Theta(|\mathbf{x}|)$. These schemes are based on LWE and strong assumptions on pairing groups either the generic (pairing) group model [4] or knowledge assumptions [3].

In this work, we set out to improve the state-of-the-art towards the direction of ideal efficiency. We observe that though there are ABE schemes for low-depth

¹ We always ignore polynomial factors in the security parameter.



Fig. 1. Efficiency comparison for KP-ABE schemes. The pink region highlights succinctness for $|ct_x|$ and $|sk_f|$. This work and $[BGG^+14]$ are KP-ABE schemes for circuits, while the rest of the schemes are for low-depth computation.

Fig. 2. Efficiency comparison for CP-ABE schemes. The pink region highlights succinctness for $|ct_f|$ and $|sk_x|$. All the included schemes are CP-ABE for low-depth computation.

computations with constant-size keys, we do not have such ABE for general circuits. We ask:

Can we construct ABE for circuits with constant-size keys?

Furthermore, all of the above schemes either have succinct keys or succinct ciphertexts, but never both at the same time. If we were to eventually achieve ideal efficiency, we would have to first overcome the intermediate barrier of simultaneously having succinct keys and ciphertexts — we refer to this as *double succinctness*. We thus ask:

> Can we construct ABE for expressive policies with both succinct keys and succinct ciphertexts?

We note that the above questions are unanswered even when assuming the strong primitive of indistinguishability obfuscation (iO). Several works [17,18,27] constructed ABE for circuits (or even functional encryption for circuits) using indistinguishability obfuscation or related primitives. However, they all have large secret keys of size poly(|f|). The only work that manages to obtain ABE for RAM with constant-size keys [20] rely on a strong primitive called extractable witness encryption, which however lacks provably secure instantiation.

Our Results. We address both questions. For the former, we construct the first KP-ABE scheme for circuits with constant-size keys while keeping the ciphertext size the same as in [11]. Concretely, each secret key consists of only 3 group

elements. For the latter, we present the first CP-ABE scheme for Boolean formulae achieving double succinctness — it has constant-size keys and succinct ciphertexts. Concretely, each secret key consists of only 2 group elements. Both constructions rely on LWE and the generic (pairing) group model, similar to [4].

Theorem (KP-ABE). Assuming LWE, in the generic (pairing) group model, there is a KP-ABE for circuits (Construction 2) that achieves selective security and has key size $|\mathbf{sk}_C| = \text{poly}(\lambda)$ (concretely, containing 3 group elements) and ciphertext size $|\mathbf{ct}_{\mathbf{x}}| = |\mathbf{x}| \text{poly}(\lambda, d)$, where d is the maximum depth of the policy circuits.

Theorem (CP-ABE). Assuming LWE, in the generic (pairing) group model, there is a CP-ABE for Boolean formulae [31] that achieves very selective security, and has constant-size keys $|\mathbf{sk}_{\mathbf{x}}| = \operatorname{poly}(\lambda)$ (concretely, containing 2 group elements) and ciphertexts of size $|\mathbf{ct}_f| = |\mathbf{x}|^2 \operatorname{poly}(\lambda)$ independent of the formula size |f|.

Additional Contribution — Adaptive Security. The standard security property of ABE is collusion resistance, which stipulates that no information of the message μ encrypted in a ciphertext should be revealed even when multiple secret keys are issued, as long as none of the keys alone is authorized to decrypt the ciphertext. Adaptive security requires collusion resistance to hold even when attributes and policies tied to the challenge ciphertext and the secret keys are chosen adaptively by the adversary. The weaker selective security restricts the adversary to commit to the attribute (in KP-ABE) or the policy (in CP-ABE) associated with the challenge ciphertext before seeing any parameters of the system, and very selective security further requires all attributes and policies in both the challenge ciphertext and the secret keys to be chosen statically.

Adaptive security guards against more powerful adversaries than selective security. It is known that the latter can be generically lifted to the former via complexity leveraging, at the cost of subexponential hardness assumptions. However, complexity leveraging is undesirable not only because it requires subexponential hardness, but also because it requires scaling the security parameter to be polynomial in the length of the information to be guessed, $\lambda = \text{poly}(|x|)$ in KP-ABE or $\lambda = \text{poly}(|f|)$ in CP-ABE. As a result, complexity leveraging is not a viable solution when aiming for constant-size keys, as key size $\text{poly}(\lambda)$ would already depends on |x| or |f|.

Instead, we show that in our constructions of KP- and CP-ABE, if assume adaptive LWE instead of plain LWE, then they achieve *adaptive security* and our reduction only incurs a polynomial amount of security loss. The adaptive LWE assumption [36] postulates that LWE samples of the form $\{\mathbf{s}^{\mathsf{T}}(\mathbf{A}_i - \mathbf{x}[i]\mathbf{G}) + \mathbf{e}_i^{\mathsf{T}}\}_i$ are pseudorandom, even if the adversary adaptively chooses \mathbf{x} depending on the random matrices $\{\mathbf{A}_i\}_i$.

Theorem (adaptive security). Assuming the polynomial hardness of adaptive LWE (instead of LWE), in the generic (pairing) group model, the KP-ABE scheme (Construction 2) and the CP-ABE scheme [31] are adaptively secure.

In the literature, the ABE schemes for circuits based on lattices [11,21] achieve only selective security (without complexity leveraging). Adapting it to have adaptive security has remained a technical barrier, except for very limited classes of policies such as 3-CNF [39]. Alternatively, there are schemes based on indistinguishability obfuscation or functional encryption for all circuits that are adaptively secure [27,40], but requiring stronger assumptions. Our technique can be viewed as making the lattice-based schemes adaptively secure when combined with pairing. Note that this is not trivial, for instance, the recent CP-ABE schemes in [3,4] that combine [11] with pairing groups inherit the selective security of the former (even if assuming adaptive LWE).

Organization. In Sect. 2, we present an overview of our techniques. In Sect. 3, we introduce preliminaries. In Sect. 4, we define nearly linear secret sharing schemes and non-annihilability, and construct such a scheme for bounded-depth circuits based on the adaptive LWE assumption. In Sect. 5, we present our adaptively secure KP-ABE for bounded-depth circuits. Due to space constraints, we refer the readers to the full version [31] for our doubly succinct CP-ABE as well as the secret sharing scheme and our new analysis of IPFE of [1] for that.

2 Technical Overview

High-Level Ideas. Let's focus on our KP-ABE scheme for circuits first. The first known construction of KP-ABE for circuits from LWE [22] has keys of size $|f| \operatorname{poly}(d, \lambda)$. The scheme of [11] reduces the key size to $\operatorname{poly}(d, \lambda)$. Both schemes achieve only selective security because they rely on the *lattice trapdoor simulation techniques*. Consider the BGG⁺ scheme. Its ciphertext encodes the attributes **x** and message μ as follows.

BGG:
$$\mathbf{s}^{\mathsf{T}}\mathbf{A} + \mathbf{e}^{\mathsf{T}}, \quad \mathbf{s}^{\mathsf{T}}\left(\overbrace{(\mathbf{A}_{1}||\cdots||\mathbf{A}_{\ell})}^{\mathsf{T}} - \mathbf{x} \otimes \mathbf{G}\right) + (\mathbf{e}')^{\mathsf{T}}, \quad \mathbf{s}^{\mathsf{T}}\mathbf{v} + e'' + \mu \lfloor q/2 \rfloor.$$

One can homomorphically evaluate any circuit f on the attribute encoding to obtain $\mathbf{s}^{\mathsf{T}}(\mathbf{B}_f - f(x)\mathbf{G}) + \mathbf{e}_f^{\mathsf{T}}$. To decrypt, the secret key sk_f simply is a *short* vector $\mathbf{r}_{\mathbf{A},f}$ satisfying $(\mathbf{A}||\mathbf{B}_f)\mathbf{r}_{\mathbf{A},f} = \mathbf{v}$, which can be sampled using a trapdoor $\mathbf{T}_{\mathbf{A}}$ for \mathbf{A} . This approach however has two drawbacks:

- <u>Difficulty towards Constant-Size Keys.</u> The short vector $\mathbf{r}_{\mathbf{A},f}$ contained in the secret key sk_f has size $\operatorname{poly}(d, \lambda)$. This is because it has dimension $m = n \log q$ for $\log q = \operatorname{poly}(d, \lambda)$ and entries of magnitude exponential in d.
- Difficulty towards Adaptive Security. The security proof relies on the ability to simulate trapdoors for these matrices $\mathbf{A} || \mathbf{B}_f$ corresponding to secret keys that are *unauthorized* to decrypt the challenge ciphertext with attribute \mathbf{x}^* , that is $f(\mathbf{x}^*) = 1$. However, to do so, current technique plants \mathbf{x}^* in the public matrices \mathbf{A}_i 's (contained in mpk), leading to selective security. Note that even with the stronger adaptive LWE assumption, it is unclear how to simulate these trapdoors in another way.

Towards constant-size keys and adaptive security, our construction circumvents the use of lattice trapdoors all together. At a high level, we turn attention to a much weaker lattice primitive called attribute-based laconic function evaluation (AB-LFE) [36], and lifts it to a KP-ABE scheme for circuits using pairing. AB-LFE is an interactive protocol where a receiver sends a digest of a function, which is exactly the matrix \mathbf{B}_f in BGG. The sender then encodes the attribute \mathbf{x} and message μ as follows.

AB-LFE:
$$\mathbf{s}^{\mathsf{T}}((\mathbf{A}_1 || \cdots || \mathbf{A}_\ell) - \mathbf{x} \otimes \mathbf{G}) + (\mathbf{e}')^{\mathsf{T}}, \quad \mathbf{s}^{\mathsf{T}} \mathbf{B}_f \mathbf{r} + e'' + \mu \lfloor q/2 \rfloor.$$

where $\mathbf{r} = \mathbf{G}^{-1}(\mathbf{a})$ is the bit decomposition of a random vector \mathbf{a} . Security guarantees that the encoding reveals only the output $f(\mathbf{x})$. At a first glance, the LFE encoding appears the same as BGG, but the novelty is in details. Since the LFE encoding depends on \mathbf{B}_f (and hence f), it can be generated without using lattice trapdoors — the short vector \mathbf{r} sampled first, and $\mathbf{B}_f \mathbf{r}$ computed next. When f(x) = 1, the hiding of μ follows directly from the pseudorandomness of LWE samples $\mathbf{s}^{\mathsf{T}}((\mathbf{A}_1 || \cdots || \mathbf{A}_\ell) - \mathbf{x} \otimes \mathbf{G})) + \mathbf{e}'$ and $\mathbf{s}^{\mathsf{T}} \mathbf{a} + e'''$. When \mathbf{x} is adaptively chosen, security follows naturally from adaptive LWE.

However, AB-LFE is able to avoid lattice trapdoor only because it is significantly weaker than ABE, or even 1-key ABE: 1) its message encoding depends on \mathbf{B}_f (unknown at ABE encryption time), and 2) it is only secure for a single function. Our next challenge is lifting AB-LFE back to full ABE, for which we use pairing.

More specifically, we first modify the AB-LFE scheme of [36] to obtain a *nearly linear secret sharing scheme* for circuits. It contains two parts.

Our LSS encoding:
$$L_{\mathbf{x}} = \mathbf{s}^{\mathsf{T}} ((\mathbf{A}_1 || \cdots || \mathbf{A}_\ell) - \mathbf{x} \otimes \mathbf{G}) + (\mathbf{e}')^{\mathsf{T}} \mod q,$$

 $L_f = \mathbf{s}^{\mathsf{T}} \mathsf{Round}(\mathbf{B}_f \mathbf{r}) + e'' + \mu |p/2| \mod p.$

Note that we round $\mathbf{B}_f \mathbf{r}$ from modulus q of poly(d) length to p of poly(λ) length so that the component L_f in the secret sharing that depends on f and μ has *constant* size, which is the key towards constant-size ABE keys. To solve the problem that L_f requires knowledge of \mathbf{B}_f unknown at encryption time, we use a pairing-based inner-product functional encryption (IPFE) to compute L_f in the exponent, by viewing it as as inner product $L_f = \langle \mathbf{s}^{\mathsf{T}} || \mu \lfloor p/2 \rceil$, Round($\mathbf{B}_f \mathbf{r}$)||1 \rangle , where the two vectors are known respectively at ABE encryption and key generation time. To overcome that AB-LFE only guarantees security for a single L_f . We follow the idea of [3,4] to compute $\delta_f \cdot L_f$ in the exponent instead, where δ_f is an independent and random scalar chosen at key generation time. In GGM, the presence of δ_f prevents adversaries from meaningfully "combining" information from multiple L_f for different f.

Comparison with [3,4]. Our way of combining lattice-based LSS with pairingbased IPFE differs from that of [3,4], in order to address unique technical difficulties. To start with, they use an LSS scheme based on the BGG ABE and inherits the selective security. Second, our KP-ABE scheme reveals part of the secret h $L_{\mathbf{x}}$ in the clear (in ciphertext), and only compute L_f in the exponent, whereas [3, 4] computes the entire LSS in the exponent. This is because the decryptor needs to perform the non-linear rounding operation on the result of homomorphic evaluation on $L_{\mathbf{x}}$, in order to obtain Round($\mathbf{s}^{\mathsf{T}}(\mathbf{B}_f - f(x)\mathbf{G})\mathbf{r} + \mathbf{e}_f^{\mathsf{T}}$) for decryption. Keeping $L_{\mathbf{x}}$ in the clear allows rounding, but renders security harder to prove.

Furthermore, the security proof of AB-LFE relies on noise flooding — their technique can only show that $L_f + \tilde{e}$ is secure for a super-polynomially large \tilde{e} . But noise flooding is incompatible with computing L_f in the exponent, since we must keep noises polynomially small in order for decryption to be efficient (which performs discrete logarithm). Without noise flooding, we cannot prove that unauthorized shares are pseudorandom as in [36]. Nevertheless, we show that unauthorized shares are "entropic", captured by a new notion called *nonannihilability*, and that the "entropic" L_f computed in the exponent still hides the message μ . The proof of non-annihilability combines techniques from AB-LFE and leakage simulation [16,25]. The work of [3,4] does not encounter issues with super-polynomial noises.

We add a note on our doubly succinct CP-ABE for Boolean formulae. It is closer to the CP-ABE scheme of [3,4]. However, to obtain constant-size keys, we rely on an IPFE scheme with strong (selective) simulation security — it enables simultaneously simulating a polynomial number k of ciphertexts, by programming k inner products for every secret key, while keeping the secret key constant-size (independent of k). Such strong simulation is impossible in the standard model following an incompressibility argument. We show that this is possible in GGM, in particular, the IPFE scheme of [1] satisfies it. IPFE with such strong simulation may find other applications.

Next, we explain our ideas in more details.

Combining LSS with IPFE. An IPFE scheme enables generating keys isk (\mathbf{v}_j) and ciphertexts ict (\mathbf{u}_i) associated with vectors $\mathbf{v}_j, \mathbf{u}_i \in \mathbb{Z}_p^N$ such that decryption reveals only their inner products $\langle \mathbf{u}_i, \mathbf{v}_j \rangle$ and hides all other information about \mathbf{u}_i encrypted in the ciphertexts (whereas \mathbf{v}_j associated with the keys are public). It can be based on a variety of assumptions such as MDDH, LWE, or DCR [1,2].

A nearly linear secret sharing scheme enables generating shares $L_f, L_0, \{L_b^i\}$ associated with a policy f and some secret μ , such that for any input $\mathbf{x} \in \{0, 1\}^{\ell}$, its corresponding subset of shares $L^{\mathbf{x}} = (L_0, \{L_i^{\mathbf{x}[i]}\})$, together with L_f can be used to approximately reconstruct the secret μ if and only if $f(\mathbf{x}) = 0$:

$$(L_f, L_0, \{L_i^b\}_{i \in [\ell], b \in \{0,1\}}) \leftarrow \mathsf{Share}(f, \mu; \mathbf{r})$$

$$f(x) = 0 \implies \mu \approx \mathsf{Recon}(f, \mathbf{x}, L_f, L^{\mathbf{x}}).$$

Near linearity means that Recon is linear in the shares L_f, L^x and that its output is close to the secret μ .

How can we combine these two primitives to construct a KP-ABE? We require $L_0, \{L_i^b\}$ to be independent of f and μ , and L_f to be linear in μ and the randomness **r** of Share. The first requirement allows us to simply put $L^{\mathbf{x}}$ in the ciphertext.

The second requirement allows us to encode μ , **r** into ict's and the coefficients (of L_f as a function of μ , **r**) into isk's, so that their inner product is exactly L_f . For convenience, we write $[\![x]\!]_i$ for g_i^x and use additive notation for the groups. The idea is as follows:

$$\begin{array}{l} \mathsf{kp.sk}_{f} : \llbracket \delta \rrbracket_{2}, & \mathsf{isk}(\mathsf{coefficients of } \delta L_{f}) \\ \mathsf{kp.ct}_{\mathbf{x}} : & L^{\mathbf{x}}, & \mathsf{ict}(\mu, \mathbf{r}) \end{array} \right\} \quad \llbracket \delta L_{f} \rrbracket_{\mathrm{T}} \text{ and } L^{\mathbf{x}}.$$
(1)

If $f(\mathbf{x}) = 0$, the linear reconstruction can be carried out in the exponents to approximately obtain $[\![\delta\mu]\!]_{\mathrm{T}}$. Decryption enumerates all possible errors to recover μ exactly. We stress again that different from [3,4], we keep $L_{\mathbf{x}}$ in the clear (in the ciphertext), instead of computing the entire secret sharing $L_{\mathbf{x}}, L_f$ in the exponent, which is important for achieving constant-size keys, but makes proving security more difficult.

We construct a secret sharing scheme that features L_f of *constant size*, which translates to KP-ABE with constant-size secret keys.

Combining secret sharing and IPFE to construct CP-ABE is similar. We can encode $L_0, \{L_i^b\}$ in ict's, and a "selection" vector according to \mathbf{x} in isk's, so that their inner products are exactly $L^{\mathbf{x}}$:

$$\begin{array}{l} \mathsf{cp.sk}_{\mathbf{x}} : \ \llbracket \delta \rrbracket_2, & \mathsf{isk}(\delta \cdot \mathsf{selection vector for } \mathbf{x}) \\ \mathsf{cp.ct}_f : \ \llbracket L_f \rrbracket_1, & \mathsf{ict}(L_0, \{L_i^b\}). \end{array} \right\} \quad \llbracket \delta L_f \rrbracket_{\mathrm{T}} \text{ and } \llbracket \delta L^{\mathbf{x}} \rrbracket_{\mathrm{T}}.$$
(2)

We use an IPFE scheme with secret keys of constant size, independent of the vector dimension or the number of ciphertexts, and a secret sharing scheme whose $L_f, L^{\mathbf{x}}$ grows only with the input length $|\mathbf{x}|$. This translates to CP-ABE with double succinctness.

Lattice-Based Nearly Linear Secret Sharing. The BGG⁺ ABE scheme introduces an important homomorphic evaluation procedure: Given public matrices $\mathbf{B} = (\mathbf{A}_1 || \cdots || \mathbf{A}_{|\mathbf{x}|})$, and the following encoding of an input \mathbf{x} , one can homomorphically evaluate any circuit f on the encodings to obtain an encoding of the output.

$$\mathbf{c}^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}}(\mathbf{B} - (1, \mathbf{x}) \otimes \mathbf{G}) + \mathbf{e}_{2}^{\mathsf{T}},$$

EvalCX($\mathbf{c}_{2}, f, \mathbf{x}$) = $\mathbf{c}_{f}^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}}(\mathbf{B}_{f} - f(\mathbf{x})\mathbf{G}) + \mathbf{e}_{f}^{\mathsf{T}},$ where EvalC(\mathbf{B}, f) = \mathbf{B}_{f} . (3)

As discussed before, the BGG⁺ ABE scheme uses lattice trapdoor simulation technique, which we try to avoid in order to get constant-size key and adaptive security.

We hence turn to using the weaker primitive of AB-LFE scheme introduced by [36]. It is a two-party protocol between a sender and a receiver who share the LWE public matrix **B** as the common reference string. The receiver first computes a digest $\mathbf{B}_f = \mathsf{EvalC}(\mathbf{B}, f)$ for a function f and sends it to the sender. Upon receiving the digest, the sender masks a message μ by an LWE sample $c_0 = \mathbf{s}^{\mathsf{T}} \mathbf{v}_f + e + \mu \lfloor q/2 \rfloor$, where $\mathbf{r} = \mathbf{G}^{-1}(\mathbf{a})$ and $\mathbf{v}_f = \mathbf{B}_f \mathbf{r}$ are analogues of $\mathbf{r}_{\mathbf{A},f}$ and \mathbf{v} in BGG⁺. It also encodes an attribute \mathbf{x} into LWE samples \mathbf{c}_1 as described below.

AB-LFE.crs : **B**
AB-LFE.digest :
$$\mathbf{B}_f = \mathsf{EvalC}(\mathbf{B}, f)$$

AB-LFE.digest : $\mathbf{B}_f = \mathsf{EvalC}(\mathbf{B}, f)$
AB-LFE.digest : $\mathbf{B}_f = \mathsf{EvalC}(\mathbf{B}, f)$

To decrypt, first run $\text{EvalCX}(\mathbf{c}_1, f, \mathbf{x})$ to obtain $\mathbf{c}_f = \mathbf{s}^{\mathsf{T}}(\mathbf{B}_f - f(\mathbf{x})\mathbf{G}) + \mathbf{e}_f^{\mathsf{T}}$. If $f(\mathbf{x}) = 0$, the decryptor can compute $c_0 - \mathbf{c}_f^{\mathsf{T}}\mathbf{r} = \mu \lfloor q/2 \rfloor + (e - \mathbf{e}_f^{\mathsf{T}}\mathbf{r})$ and round it to recover μ .

Observe that the above scheme can be viewed as a nearly linear secret sharing scheme, where the shares chosen by \mathbf{x} are exactly $\mathbf{L}^{\mathbf{x}} = \mathbf{c}_1$ and the shares dependent on f and μ is $L_f = c_0$. At the moment, the bit-length of L_f is $\Theta(\log q)$. Since the noise growth during the homomorphic evaluation is exponential to the depth of the computation, q is a poly (d, λ) -bit modulus in order to accommodate for the noise growth. We next turn to reducing the size of L_f to a constant independent of d.

Rounding to Make L_f **Constant-Size.** Since the encrypted message is only a single bit, we can afford to lose a lot of precision in the above decryption process. In particular, the scheme is still correct if we round down the digest \mathbf{B}_f to a much smaller, depth-independent, modulus $p \ll q$, and change c_0 to use the rounded digest (while keeping $\mathbf{c}_1^{\mathsf{T}}$ unchanged):

$$c'_0 = \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p + \mu \lfloor p/2 \rceil + e \quad \text{over } \mathbb{Z}_p.$$

During decryption, one now computes, over \mathbb{Z}_p ,

$$c_{0}^{\prime} - \lfloor \mathbf{c}_{f}^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rfloor_{p} = c_{0}^{\prime} - \lfloor \mathbf{s}^{\mathsf{T}} \mathbf{B}_{f} \mathbf{G}^{-1}(\mathbf{a}) + f(\mathbf{x}) \mathbf{s}^{\mathsf{T}} \mathbf{a} + \mathbf{e}_{f}^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rceil_{p}$$
$$= c_{0}^{\prime} - \left(\mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}_{f} \mathbf{G}^{-1}(\mathbf{a}) \rceil_{p} + f(\mathbf{x}) \lfloor \mathbf{s}^{\mathsf{T}} \mathbf{a} \rceil_{p} + \underbrace{\lfloor \mathbf{e}_{f}^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rceil_{p}}_{e_{f}^{\prime}} + e_{s} \right)$$
$$= \mu \lfloor p/2 \rceil - f(\mathbf{x}) \lfloor \mathbf{s}^{\mathsf{T}} \mathbf{a} \rceil_{p} + (e - e_{f}^{\prime} - e_{s}), \qquad (4)$$

where the rounding error e_s is of magnitude $|e_s| = \Theta(||\mathbf{s}||_1)$. As long as the error terms are much smaller than p/2, when $f(\mathbf{x}) = 0$, one can still recover μ . We can now recast the above rounded AB-LFE scheme into a secret sharing scheme with L_f of bit-length $\Theta(\log p)$, independent of depth d. (Only the larger modulus q will, and thus \mathbf{e}_f itself can, grow with d.)

$$\begin{aligned} & \mathsf{SS.pp}: \quad \mathbf{a}, \mathbf{B} & \mod q; \\ & L_f: \quad c_0' = \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p + \mu \lfloor p/2 \rceil + e & \mod p \ll q; \\ & \mathbf{L}^{\mathbf{x}}: \quad \mathbf{c}_1^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}} \left(\mathbf{B} - (1, \mathbf{x}) \otimes \mathbf{G} \right) + \mathbf{e}_1^{\mathsf{T}} & \mod q. \end{aligned}$$

As shown in Eq. (1), to obtain KP-ABE, we will use a pairing-based IPFE to compute L_f in the exponent. Specifically, the IPFE secret key isk encodes $(\lfloor \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p, \lfloor p/2 \rceil)$, and the IPFE ciphertext ict encodes $(\mathbf{s}^{\intercal}, \mu)$. Together, they decrypt to exactly $\llbracket L_f \rrbracket_T$. Since both vectors live in \mathbb{Z}_p , the KP-ABE key, consisting of only isk, is of size independent of d. Our secret sharing scheme is summarized below. It turns out that arguing security is actually tricky and requires additional modification.

Our Secret Sharing Scheme for KP-ABE $\begin{aligned} & \mathsf{Setup}(1^{\lambda}) : \mathsf{pp} = \mathsf{LFE}.\mathsf{pp} = (\mathbf{a}, \mathbf{B}) = (\mathbf{a}, \mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_{\ell}). \\ & \mathsf{ShareX}(\mathsf{pp}) : \mathsf{Compute LWE samples} \\ & \mathbf{L}_0 = \mathbf{s}^{\mathsf{T}} \mathbf{A}_0 + \mathbf{e}_0, \quad \mathbf{L}_i^b = \mathbf{s}^{\mathsf{T}} (\mathbf{A}_i - b\mathbf{G}) + \mathbf{e}_i^{\mathsf{T}}. \\ & \mathsf{Output} (\mathbf{L}_0, \{\mathbf{L}_i^b\}, \mathbf{s}). \\ & \mathsf{ShareF}(\mathsf{pp}, f, \mu, \mathbf{s}) : \mathsf{Compute B}_f = \mathsf{EvalC}(\mathbf{B}, f). \\ & \mathsf{Output } L_f = \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p + \mu \lfloor p/2 \rceil. \\ & \forall \mathbf{x} \in \{0, 1\}^{\ell} : \mathbf{L}^{\mathbf{x}} = (\mathbf{L}_0, \mathbf{L}_1^{\mathbf{x}[1]}, \dots, \mathbf{L}_{\ell}^{\mathbf{x}[\ell]}) \\ & \mathsf{Recon}(\mathsf{pp}, f, L_f, \mathbf{x}, \mathbf{L}^{\mathbf{x}}) : \mathsf{If} \ f(\mathbf{x}) = 1, \ \mathsf{output } \bot. \\ & \mathsf{Otherwise, \ compute } \mathbf{c}_f = \mathsf{EvalCX}(\mathbf{L}^{\mathbf{x}}, f, \mathbf{x}), \ \mathsf{and} \\ & \mathsf{recover} \ \mu \ \mathrm{from} \ L_f - \lfloor \mathbf{c}_f^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rceil_p \approx \mu \lfloor p/2 \rceil. \end{aligned}$

Non-Annihilability by Leakage Simulation. However, using AB-LFE creates a further complication, as its security relies on flooding the e'_f, e_s terms (which may contain information of **s** and **x**) with e, in order to prove pseudorandomness of L_f . By Eqs. (3, 4), when $f(\mathbf{x}) = 1$ we have

$$L_f = \lfloor \mathsf{EvalCX}(\mathbf{L}^{\mathbf{x}}, f, \mathbf{x})^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rceil_p - \lfloor \mathbf{s}^{\mathsf{T}} \mathbf{a} + e_a \rceil_p + \mu \lfloor p/2 \rceil + (e - e'_f - e_s).$$
(5)

Observe that in the above, for later convenience, an additional polynomial LWE noise e_a is introduced in the term $\lfloor \mathbf{s}^{\mathsf{T}} \mathbf{a} + e_a \rceil_p$ (which by rounding simply equals to $|\mathbf{s}^{\mathsf{T}} \mathbf{a} \rceil_p$).

At this point, in order to show that L_f is pseudorandom, given that **x** is selected before Setup, one could program the public matrices as $\mathbf{A}_i = \mathbf{A}'_i + x_i \mathbf{G}$ according to **x**, where $\mathbf{B}' = (\mathbf{A}'_0, ..., \mathbf{A}'_\ell)$ are sampled at random. And one would hope to apply LWE to argue that

$$\mathbf{L}_{\mathbf{x}} = \mathbf{s}^{\mathsf{T}}(\mathbf{B} + (1, \mathbf{x}) \otimes \mathbf{G}) + \mathbf{e}_{1}^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}}\mathbf{B}' + \mathbf{e}_{1}^{\mathsf{T}},$$

and $(\mathbf{s}^{\mathsf{T}}\mathbf{a} + e_a)$ are jointly pseudorandom. However, the noise terms e'_f and e_s may leak information about \mathbf{e}_2 and \mathbf{s} .

The solution in [36] is noise flooding. By setting e to be super-polynomially larger than $(e'_f + e_s)$, we have $e - e'_f - e_s \approx_s e$. By LWE, we can now switch $\mathbf{L}^{\mathbf{x}}$ and $(\mathbf{s}^{\mathsf{T}}\mathbf{a} + e_a)$ to random and conclude that L_f is pseudorandom.

However, the unique challenge here is that L_f is going to be computed in the exponent of the pairing group, and decryption only recovers $(\mu \lfloor p/2 \rfloor + e)$ in the exponent. When e is super-polynomial, we can no longer extract μ out of the exponent. Our solution is avoiding flooding altogether and remove the noise e from L_f . As such, we cannot prove pseudorandomness of L_f , but only a weaker security notion that we call non-annihilability (for L_f). This notion captures that L_f is still entropic.

Non-Annihilability. Non-annihilability requires that no adversary, after seeing $\mathbf{L}^{\mathbf{x}}$ (but not L_f) can come up with an affine function γ such that $\gamma(L_f) = 0$. As we will see, this security notion, combined with GGM, suffices for our proof.

Towards proving non-annihilability, we want to show that L_f is highly entropic (even without e). Our idea is to view the noises e'_f, e_s as leakage of the randomness that generates $\mathbf{L}^{\mathbf{x}}$ and $(\mathbf{s}^{\mathsf{T}}\mathbf{a} + e_a)$ as well as the other information, and simulate e'_f, e_s using leakage simulation [16,25]. Crucially, because e'_f, e_s have polynomial range, the simulation can run in polynomial time. More precisely, the leakage simulation lemma of [16] states that for any joint distribution $(X, Z) \sim \mathcal{D}$ (Z viewed as leakage of randomness for generating X), adversary size bound s, and error bound ε , there is a simulator h simulating Z as h(X)such that (X, Z) and (X, h(X)) are (s, ε) -indistinguishable. Furthermore, the running time of h is $O(s\varepsilon^{-2}2^{|Z|})$. Suppose for contradiction that there is an adversary \mathcal{A} of size $s = \operatorname{poly}(\lambda)$ winning the non-annihilability game with probability $2\varepsilon \geq 1/\operatorname{poly}(\lambda)$. Consider the joint distribution \mathcal{D} of running the game with \mathcal{A} , defined in the first line below:

$$\mathcal{D} \to \{ X = (\mathsf{pp}, \mathbf{x}, \mathbf{L}^{\mathbf{x}}, f, \mu, \gamma, \psi = \mathbf{s}^{\mathsf{T}} \mathbf{a} + e_a), \qquad Z = e'_f + e_s \}$$

 $\stackrel{s,\varepsilon}{\approx} \text{Hybrid } 1 \to \{ X = (\mathsf{pp}, \mathbf{x}, \mathbf{L}^{\mathbf{x}}, f, \mu, \gamma, \psi = \mathbf{s}^{\mathsf{T}} \mathbf{a} + e_a), \qquad Z = h(X) \}$

 $\approx \text{Hybrid } 2 \to \{X = (\mathsf{pp}, \mathbf{x}, \mathbf{L}^{\mathbf{x}} \text{ random}, f, \mu, \gamma, \psi \text{ random}), \quad Z = h(X) \}.$

Using (X, Z), one can emulate L_f as (cf. Eq. (5) with e removed and $(e'_f + e_s)$ replaced by Z)

$$L_f = \lfloor \mathsf{EvalCX}(\mathbf{L}^{\mathbf{x}}, f, \mathbf{x})\mathbf{G}^{-1}(\mathbf{a}) \rceil_p - \lfloor \psi \rceil_p + \mu \lfloor p/2 \rceil - Z.$$

Since $Z = e'_f + e_s$, and s, ε^{-1} are all polynomially bounded, we can simulate Z by h(X) in polynomial time (Hybrid 1). Now, we can apply LWE to switch $\mathbf{L}^*, \psi = \mathbf{s}^{\mathsf{T}} \mathbf{a} + e_a$ to random (Hybrid 2). At this point, it seems that L_f is just pseudorandom by the pseudorandomness of ψ . However, there is a subtle issue: Z = h(X) depends on ψ contained in X, and hence $(\lfloor \psi \rceil_p - Z)$ may not be pseudorandom, and neither may be L_f . Despite this dependency, thanks again to $(e'_f + e_s)$, thus h(X), being polynomially bounded, $(-\lfloor \psi \rceil_p + h(X))$ still has almost full entropy (up to a logarithmic loss). Therefore, the probability that L_f is annihilated by an affine function γ chosen by \mathcal{A} before ψ is randomly sampled is negligible. This gives a contradiction and concludes the proof of non-annihilability.

Multi-Key Security of KP-ABE in GGM. Our KP-ABE scheme combines an IPFE scheme with the secret sharing scheme described above. As described before, in our KP-ABE scheme, we only compute the L_f part of secret sharing using IPFE, and leave the $\mathbf{L}^{\mathbf{x}}$ part in the clear so that rounding can be performed. To achieve multi-key security, we further employ the idea from [3,4] to "isolate" each ABE secret key in GGM by multiplying it with a fresh random element δ .

kp.sk :
$$\llbracket \delta \rrbracket_2$$
, isk $(\llbracket \delta(\lfloor \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p, \lfloor p/2 \rceil) \rrbracket_2)$
kp.ct : $\mathbf{L}^{\mathbf{x}}$, ict $(\llbracket (\mathbf{s}, \mu) \rrbracket_1)$ decrypt to $\llbracket \delta L_f \rrbracket_T$.

The decryption algorithm first computes IPFE decryption to recover $[\![\delta L_f]\!]_{\mathrm{T}}$. It then computes (homomorphically in the exponent of g_{T})

$$\llbracket \delta L_f \rrbracket_{\mathrm{T}} - \llbracket \delta \rrbracket_{\mathrm{T}} \lfloor \mathbf{c}_f^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rceil_p = \llbracket \delta \left(\mu \lfloor p/2 \rceil - (e'_f + e_s) \right) \rrbracket_{\mathrm{T}}.$$

Since the noise $(e'_f + e_s)$ has a polynomial range, the decryption algorithm enumerates all its possible values to recover μ .

Multi-key security, at a high level, relies on the fact that in GGM, an adversary can only learn information about $[\![\delta L_f]\!]_T$ by submitting zero-test queries of affine functions. When the adversary attacks multiple keys, it essentially submits zero-test queries over the terms $\{\delta_j L_{f_j}\}$. Let $\gamma(\{\delta_j L_{f_j}\})$ be any zero-test query submitted by \mathcal{A} , we can view it as a degree-1 polynomial over δ_i 's:

$$\gamma(\{\delta_j L_{f_j}\}) = \sum_j \gamma_j(L_{f_j})\delta_j + \gamma_0,$$

where $\gamma_j(L_{f_j})$ is the coefficient of δ_j . Since each δ_j is sampled independently at random, by Schwartz–Zippel, with all but negligible probability, γ evaluates to zero only if all γ_j 's evaluate to zero. In other words, the adversary is effectively constrained to annihilate each L_{f_j} individually. By the non-annihilability for L_f , if γ_j is not the zero function, it evaluates to non-zero with overwhelming probability. Hence the adversary learns no information of each L_{f_j} and the message μ encoded in them.

Our KP-ABE Scheme Setup (1^{λ}) : Output mpk = impk for IPFE, pp for secret sharing and msk = imsk for IPFE. KeyGen(msk, C): Sample $\delta \stackrel{\$}{\leftarrow} \mathbb{Z}_p$ and compute $\mathbf{B}_f = \text{EvalC}(\mathbf{B}, f)$. Output sk = $([\![\delta]\!]_2, \text{isk}([\![\delta(\lfloor \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p, \lfloor p/2 \rceil)]\!]_2))$. Enc(mpk, \mathbf{x}, μ): Compute $(\mathbf{L}_0, \{\mathbf{L}_b^b\}, \mathbf{s}) \stackrel{\$}{\leftarrow} \text{ShareX}(\text{pp})$. Output ct = $(\mathbf{L}^{\mathbf{x}}, \text{ict}([\![\mathbf{s}, \mu]\!]_1))$. Dec(mpk, sk, C, ct, \mathbf{x}): Run IPFE decryption to recover $[\![\delta L_f]\!]_{\mathrm{T}}$. Compute $\mathbf{c}_f = \text{EvalCX}(\mathbf{L}^{\mathbf{x}}, f, \mathbf{x})$ and find μ from $[\![\delta L_f]\!]_{\mathrm{T}} - [\![\delta]\!]_{\mathrm{T}} \lfloor \mathbf{c}_f^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) \rceil_p = [\![\delta(\mu \lfloor p/2] - e_s - e'_f)]\!]_{\mathrm{T}}$. **Summary of Our KP-ABE.** Combining the above secret sharing scheme with an IPFE scheme, we obtain a KP-ABE scheme for bounded-depth circuits as summarized above.

We note that our KP-ABE scheme achieves the same asymptotic ciphertext compactness as the BGG⁺ scheme. Let d be an upper bound on the depth of the policy f, then $|\mathbf{ct}| = \text{poly}(\lambda, d)|\mathbf{x}|$. The secret keys of our scheme contains only O(1) group elements, in fact only *three* using the IPFE scheme of [2] in a group of order p. We set $\log p = \text{poly}(\lambda)$ and hence obtain constant size keys.

Security Sketch for KP-ABE. Finally, for completeness, we add a security sketch that puts the previous ideas together. We emphasize that that we only use GGM in the last argument, when we need to isolate the share L_{f_i} for each f.

The selective security game of ABE (summarized in H_0 below) at a high level is as follows: The adversary \mathcal{A} first decides a challenge attribute \mathbf{x}^* before receiving a master public key mpk and a ciphertext ct^* from the challenger \mathcal{C} . It is then allowed to repeatedly query secret keys sk_j for functions f_j . The adversary wins if every queried function f_j satisfies $f_j(\mathbf{x}^*) \neq 0$, and if it guesses the encrypted bit μ correctly.

$$\begin{array}{ccc} \mathsf{H}_{0} & \mathsf{H}_{1} \\ \mathcal{C} & & \mathsf{H}_{1} \\ & & \mathsf{mpk}, \mathsf{ct}^{*} = \mathbf{L}^{\mathbf{x}^{*}}, \mathsf{ict}(\llbracket \mathbf{s}, \mu \rrbracket_{1}) \\ & & \mathsf{f}_{j} \\ & & \mathsf{isk}_{j} = (\llbracket \delta_{j} \rrbracket_{2}, \\ & & \mathsf{isk}(\llbracket \delta_{j}(\lfloor \mathbf{B}_{f} \mathbf{G}^{-1}(\mathbf{a}) \rceil_{p}, \lfloor p/2 \rceil) \rrbracket_{2})) \end{array} \right) \\ \end{array} \right\} \xrightarrow{\operatorname{IPFE}} \begin{array}{c} \mathsf{H}_{1} \\ & & \mathsf{mpk}, \mathsf{ct}^{*} = \mathbf{L}^{\mathbf{x}^{*}}, \mathsf{ict}(\bot) \\ & & \mathsf{f}_{j} \\ & & \mathsf{isk}_{j} = (\llbracket \delta_{j} \rrbracket_{2}, \\ & & \mathsf{isk}(\llbracket \delta_{j} L_{f_{j}} \rrbracket_{2})) \end{array} \right) \\ \end{array}$$
repeat
$$\begin{array}{c} \mathsf{repeat} \\ & & \mathsf{isk}(\llbracket \delta_{j} L_{f_{j}} \rrbracket_{2})) \\ & & \mathsf{isk}(\llbracket \delta_{j} L_{f_{j}} \rrbracket_{2})) \end{array} \right) \\ \end{array}$$
repeat
$$\begin{array}{c} \mathsf{repeat} \\ & & \mathsf{isk}(\llbracket \delta_{j} L_{f_{j}} \rrbracket_{2}) \end{array} \right) \\ \end{array}$$

Note that we can generate the IPFE ciphertext $ict(\llbracket \mathbf{s}, \mu \rrbracket_1)$ before any IPFE secret keys $isk(\llbracket \delta(\lfloor \mathbf{B}_{f_j} \mathbf{G}^{-1}(\mathbf{a}) \rceil_p, \lfloor p/2 \rceil) \rrbracket_2)$. Relying on the selective simulation security of IPFE, we can (as summarized in \mathbf{H}_1 above) replace $ict(\llbracket \mathbf{s}, \mu \rrbracket_1)$ with a simulated ciphertext $ict(\bot)$, and each $isk(\llbracket \delta_j(\lfloor \mathbf{B}_{f_j} \mathbf{G}^{-1}(\mathbf{a}) \rceil_p, \lfloor p/2 \rceil) \rrbracket_2)$ with a simulated secret key $isk(\llbracket \delta_j L_{f_i} \rrbracket_2)$ using their inner products.

In GGM, we can now argue that \mathcal{A} only learns information about μ through zero-test queries over $\{\delta_j L_{f_j}\}$. As argued before, by the non-annihilability of L_f , the adversary learns no information of μ .

Building Doubly Succinct CP-ABE. To build a CP-ABE scheme we need a different secret sharing construction, because the previous rounding solution does not work anymore. As described in Eq. (2), in the CP case, we use IPFE to compute $\mathbf{L}^{\mathbf{x}}$ in the exponent, hence cannot perform rounding on it. Without rounding, the \mathbf{e}_f term, as a result of EvalCX, in Eq. (4) becomes super-polynomial. This again makes the ABE decryption inefficient.

Fortunately, for Boolean formulae, the work of [23] develops specialized homomorphic evaluation procedures EvalF, EvalFX that ensure the evaluation noise \mathbf{e}_f has a polynomial range. Therefore, our secret sharing scheme for CP removes the rounding and replaces EvalC, EvalCX by EvalF, EvalFX. We summarize our modified secret sharing scheme below (Setup, ShareX are kept the same).

Modified Secret Sharing Scheme for CP-ABE ShareF'(pp, f, μ, \mathbf{s}) : Compute $\mathbf{B}_f = \text{EvalF}(\mathbf{B}, f)$. Output $L_f = \mathbf{s}^{\mathsf{T}} \mathbf{B}_f \mathbf{G}^{-1}(\mathbf{a}) + \mu \lfloor p/2 \rfloor + e$. Recon'(pp, $f, L_f, \mathbf{x}, \mathbf{L}^{\mathbf{x}}$) : If $f(\mathbf{x}) = 1$, output \perp . Otherwise, compute $\mathbf{c}_f = \text{EvalFX}(\mathbf{L}^{\mathbf{x}}, f, \mathbf{x})$, and find μ from $L_f - \mathbf{c}_f^{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) = \mu \lfloor p/2 \rfloor + (e - e'_f)$.

As noted before, in our CP-ABE scheme we use IPFE to compute $\mathbf{L}^{\mathbf{x}}$. To achieve double succinctness, we carefully implement a pair of functions Sel, Encode using an IPFE with constant-size isk's, such that $Sel([[\mathbf{x}]]_2)$ and $Encode([[\mathbf{L}_0, {\mathbf{L}_i^b}]]_1)$ decrypts exactly to $[[\mathbf{L}^{\mathbf{x}}]]_{\mathrm{T}}$. We obtain a CP-ABE scheme for Boolean formulae as summarized below.

 $\label{eq:output} \begin{array}{l} \textbf{Our CP-ABE Scheme} \\ & \mathsf{Setup}(1^{\lambda}): \mathsf{Output } \mathsf{mpk} = \mathsf{impk} \text{ for IPFE and pp for secret sharing,} \\ & \mathsf{and } \mathsf{msk} = \mathsf{imsk} \text{ for IPFE.} \\ & \mathsf{KeyGen}(\mathsf{msk}, \mathbf{x}): \mathsf{Sample} \ \delta \overset{\$}{=} \mathbb{Z}_p. \\ & \mathsf{Output } \mathsf{sk} = (\llbracket \delta \rrbracket_2, \mathsf{Sel}(\llbracket \delta \mathbf{x} \rrbracket_2)). \\ & \mathsf{Enc}(\mathsf{mpk}, f, \mu): \mathsf{Compute} \ (\mathbf{L}_0, \{\mathbf{L}_i^b\}, \mathbf{s}) \overset{\$}{=} \mathsf{ShareX}(\mathsf{pp}) \\ & \mathsf{and} \ L_f \overset{\$}{=} \mathsf{ShareF}'(\mathsf{pp}, f, \mu, \mathbf{s}). \\ & \mathsf{Output } \mathsf{ct} = (\llbracket L_f \rrbracket_1, \mathsf{Encode}(\llbracket \mathbf{L}_0, \{\mathbf{L}_i^b\} \rrbracket_1)). \\ & \mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}, \mathbf{x}, \mathsf{ct}, f): \mathsf{Run } \mathsf{IPFE} \ decryption \ \mathsf{to } \mathsf{recover} \ \llbracket \delta \mathbf{L}^{\mathbf{x}} \rrbracket_{\mathsf{T}}. \\ & \mathsf{Compute} \ \llbracket \delta \mathbf{c}_f \rrbracket_{\mathsf{T}} = \mathsf{EvalFX}(\llbracket \delta \mathbf{L}^{\mathbf{x}} \rrbracket_{\mathsf{T}}, f, \mathbf{x}), \\ & \mathsf{and} \ \mathrm{find} \ \mu \ \mathrm{from} \\ & \llbracket L_f \rrbracket_1 \llbracket \delta \rrbracket_2 - \llbracket \delta \mathbf{c}_f^{\mathsf{T}} \rrbracket_{\mathsf{T}} \mathbf{G}^{-1}(\mathbf{a}) = \llbracket \delta(\mu \lfloor p/2 \rceil + (e - e_f')) \rrbracket_{\mathsf{T}}. \end{array}$

We now describe the Sel, Encode functions. Let $\ell = |\mathbf{x}|$ denote the length of **x**. The Sel algorithm first computes the "selection vector" for **x** as

$$\mathbf{v} = (1, 1 - \mathbf{x}[1], \mathbf{x}[1], \dots, 1 - \mathbf{x}[i], \mathbf{x}[i], \dots),$$

and then computes an IPFE secret key $isk([v]_2)$. The Encode algorithm places input shares in the matrix

$$\begin{pmatrix} \mathbf{L}_0 & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_1^0 & \mathbf{L}_1^1 & \cdots & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{L}_\ell^0 & \mathbf{L}_l^1 & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{L}_\ell^0 & \mathbf{L}_\ell^1 \end{pmatrix}$$

and computes one IPFE ciphertext for each row \mathbf{u}_l of the matrix. Our CP-ABE has both succinct keys and ciphertexts: $|\mathbf{sk}| = O(1)$ and $|\mathbf{ct}| = \text{poly}(\lambda)|\mathbf{x}|^2$.

Simulation Security for IPFE in GGM. Similar to the security proof for KP-ABE, our security proof for CP-ABE requires selective simulation security of IPFE.

 $\begin{array}{ccc} \mathsf{Sel}: & \mathsf{isk}(\llbracket \mathbf{v} \rrbracket_2) \\ \mathsf{Encode}: \forall l & \mathsf{ict}(\llbracket \mathbf{u}_l \rrbracket_1) \end{array} \begin{vmatrix} c & \widetilde{\mathsf{isk}}(\llbracket \mathbf{L}^{\mathbf{x}} \rrbracket_2) \\ \approx & \widetilde{\mathsf{ict}}(\bot) \end{vmatrix}$

Note that above we need to simulate *multiple* IPFE ciphertexts and program all their decryption outcome $\mathbf{L}^{\mathbf{x}}$ in each secret key. This is possible using existing IPFE schemes [2,32], but at the cost of having the secret key size proportional to the number $k = |\mathbf{L}^{\mathbf{x}}|$ of ciphertexts to be simulated. However, we aim for constant-size secret keys (independent of k). Unfortunately, in the standard model, it is impossible to achieve simulation security for k ciphertexts if the secret key is shorter than k bits by an incompressibility argument [12]. We show that simulation security for unbounded polynomially many ciphertexts can nevertheless be achieved with constant-size secret keys in the GGM. In particular, the IPFE scheme of [1], whose secret key contains a single group element, satisfies it. Roughly speaking, in the GGM, an adversary only learns information about values in the exponent through zero-test queries over the pairings of keys and ciphertexts, which the simulator can answer by translating them into zero-test queries over the inner products. As a side note, we can in fact prove *adaptive* simulation security for the [1] IPFE scheme, though our ABE scheme only relies on *selective* simulation security.

Achieving Adaptive Security. Examining the security sketch for KP-ABE, we observe that in our construction, the $ict([[s, \mu]]_1)$ component of ciphertext ct^* doesn't depend on the challenge attribute \mathbf{x}^* . This means that even in the adaptive KP-ABE game, where \mathbf{x}^* is decided after some key queries, the $ict([[s, \mu]]_1)$ component of ct^* can be fixed at the beginning of the game, before any key queries. Therefore, we can still rely on selective simulation security of IPFE for the first proof step.

However, when we next need to invoke non-annihilability for L_f , we run into a problem: the security for L_f only holds when \mathbf{x}^* is chosen before the LWE public matrix **B** is revealed in the public parameter **pp** of the secret sharing. To achieve adaptive security, what we need is adaptive non-annihilability property, which allows \mathbf{x}^* to be chosen adaptively dependent on **pp**. We show that this is implied by the *adaptive* LWE assumption formulated in [36].

In summary, we obtain adaptively secure KP-ABE for circuits and CP-ABE for Boolean formulae both with constant-size keys from GGM and Adaptive LWE.

3 Preliminaries

Let λ be the security parameter, which runs through \mathbb{N} . Except in the definitions, we suppress λ for brevity. We write [a..b] for the set $\{a, a + 1, \ldots, b\}$ and [n] for [1..n]. Vectors and matrices are written in boldface, and are always indexed using $[\cdot]$, i.e., $\mathbf{A}[i, j]$ is the (i, j)-entry of \mathbf{A} . The infinity norm of a vector and its induced operator norm of a matrix are denoted by $\|\cdot\|_{\infty}$. We will use the following lemma for various proofs:

Lemma 1 (Schwartz–Zippel). Let $P(\mathbf{z})$ be a non-zero polynomial with Z indeterminates of degree at most d over \mathbb{Z}_p , then $\Pr[\mathbf{z} \stackrel{*}{\leftarrow} \mathbb{Z}_p^Z : P(\mathbf{z}) = 0] \leq d/p$.

3.1 Attribute-Based Encryption

Definition 1 (ABE [24]). Let $\mathcal{P} = \{\mathcal{P}_{\lambda}\}_{\lambda \in \mathbb{N}}$ be a sequence of predicate families with $\mathcal{P}_{\lambda} = \{P : X_P \times Y_P \to \{0,1\}\}$. An attribute-based encryption scheme for \mathcal{P} consists of 4 efficient algorithms:

- Setup $(1^{\lambda}, P)$ takes as input the security parameter 1^{λ} and a predicate $P \in \mathcal{P}_{\lambda}$, and outputs a pair of master public/secret keys (mpk, msk).
- KeyGen(msk, y) takes as input the master secret key msk and some $y \in Y_P$, and outputs a secret key sk.
- $\mathsf{Enc}(\mathsf{mpk}, x, \mu)$ takes as input the master public key mpk , some $x \in X_P$, and a message $\mu \in \{0, 1\}$, and it outputs a ciphertext ct.
- Dec(mpk, sk, y, ct, x) takes as input the master public key mpk, a secret key sk, its associated y, a ciphertext ct, and its associated x, and is supposed to recover the message if P(x, y) = 1.

The scheme is required to be correct, i.e., for all $\lambda \in \mathbb{N}$, $P \in \mathcal{P}_{\lambda}$, $x \in X_P$, $y \in Y_P$, $\mu \in \{0,1\}$ such that P(x,y) = 1, it holds that

$$\Pr \begin{bmatrix} (\mathsf{mpk},\mathsf{msk}) \stackrel{\hspace{0.1em} {\scriptstyle \$}}{\scriptstyle \hspace{0.1em}} \mathsf{Setup}(1^{\lambda},P) \\ \\ \mathsf{sk} \stackrel{\hspace{0.1em} {\scriptstyle \$}}{\scriptstyle \hspace{0.1em}} \mathsf{KeyGen}(\mathsf{msk},y) : \mathsf{Dec}(\mathsf{mpk},\mathsf{sk},y,\mathsf{ct},x) = \mu \\ \\ \\ \mathsf{ct} \stackrel{\hspace{0.1em} {\scriptstyle \$}}{\scriptstyle \hspace{0.1em}} \mathsf{Enc}(\mathsf{mpk},x,\mu) \end{bmatrix} = 1.$$

In KP-ABE, each $y \in Y_P$ describes a function from X_P to $\{0, 1\}$, each $x \in X_P$ is an input (bit-string) to the functions, and P(x, y) evaluates y on x. When we want to emphasize x (resp. y) is a bit-string (resp. circuit), we write \mathbf{x} (resp. C) instead.

Security. Due to space constraints, we refer the readers to [4,29,31] for the definitions of *adaptive*, *selective*, and *very selective* security for ABE.

Computation Model. We will consider KP-ABE for bounded-depth circuits for any polynomial bound. Our CP-ABE for NC^1 can be found in the full version [31].

Definition 2 (KP-ABE for circuits). *A* KP-ABE for (bounded-depth) circuits is *ABE for* $\mathcal{P}^{\mathsf{Ckt}}$.²

$$\begin{aligned} X_{\lambda,\ell,d}^{\mathsf{Ckt}} &= \{0,1\}^{\ell}, \quad Y_{\lambda,\ell,d}^{\mathsf{Ckt}} = \big\{ \text{Boolean circuit } C : \{0,1\}^{\ell} \to \{0,1\} \text{ of depth } d \big\}, \\ P_{\lambda,\ell,d}^{\mathsf{Ckt}}(\mathbf{x},C) &= \neg C(\mathbf{x}), \quad \mathcal{P}_{\lambda}^{\mathsf{Ckt}} = \big\{ P_{\lambda,\ell,d}^{\mathsf{Ckt}} \big| \ell, d \in \mathbb{N}, d \leq D_{\lambda} \big\}, \quad \mathcal{P}^{\mathsf{Ckt}} = \big\{ \mathcal{P}_{\lambda}^{\mathsf{Ckt}} \big\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

Here, D_{λ} is a super-polynomial function (specified by the constructions). As an input to Setup, the predicate $P_{\lambda,\ell,d}^{\mathsf{Ckt}}$ is represented by $(1^{\ell}, 1^{d})$.

Note that since Setup takes the unary representation of ℓ , d, which will be polynomial in λ , as input, they are bounded by *that* polynomial once the system is set up. However, d can be up to D_{λ} , which is super-polynomial in λ , so one can set up the system for any polynomial depth, i.e., our KP-ABE for circuits supports bounded-depth circuits for arbitrary polynomial depth bound.

Compactness and Succinctness. Since KeyGen, Enc run in polynomial time, the lengths of key and ciphertext could grow polynomially in |y|, |x|, respectively. Moreover, the input length is an argument passed into Setup, so both keys and ciphertexts could have polynomial size dependency on it. We are interested in ABE schemes with short keys and ciphertexts:

Definition 3 (ABE efficiency). A KP-ABE for circuits (of depth at most d) has

- succinct keys if $|\mathbf{sk}| = \text{poly}(\lambda, d)$ is independent of $|C|, |\mathbf{x}|$;
- compact ciphertexts if $|\mathbf{ct}| = |\mathbf{x}| \operatorname{poly}(\lambda, d)$ is independent of |C|.

We remark that an ideally succinct component should be of length $poly(\lambda)$. Nevertheless, our version defined above is still meaningful as the circuit size can be much larger than its depth.

3.2 Lattice Tools

Homomorphic Evaluation. We use the following abstraction of homomorphic evaluation for ABE over lattices, developed in a series of works [11,19,23] with the syntax in [13,14]. The actual algorithm we use is a slightly changed version of that for ABE for circuits in [11]. In our version, instead of using **G** as the gadget matrix, we consider **QG** for any invertible **Q**. Note that $\mathbf{G}^{-1}(\mathbf{Q}^{-1} \times \cdot)$ is a right inverse of **QG** with binary output. We replace any invocation of $\mathbf{G}^{-1}(\cdot)$ in the original algorithms by $\mathbf{G}^{-1}(\mathbf{Q}^{-1} \times \cdot)$ to obtain the following:

² When working with lattices, it is more convenient to indicate authorization of decryption by zero, thus the negation of $C(\mathbf{x})$.

Lemma 2 (homomorphic evaluation for circuits, adapted from [11]). EvalC and EvalCX are two efficient deterministic algorithms. Let n, ℓ, q be positive integers, $m = n \lceil \log_2 q \rceil$, **G** the gadget matrix, **B** a matrix over \mathbb{Z}_q of shape $n \times (\ell + 1)m$, **Q** an invertible matrix over \mathbb{Z}_q of shape $n \times n$, **x** an ℓ -bit string (row vector), and C a circuit of depth d with input length ℓ . The algorithms work as follows:

- EvalC(B, Q, C) outputs $\mathbf{H}_C \in \mathbb{Z}^{(\ell+1)m \times m}$;
- EvalCX(**B**, **Q**, *C*, **x**) outputs $\widehat{\mathbf{H}}_{C,\mathbf{x}} \in \mathbb{Z}^{(\ell+1)m \times m}$.

The outputs satisfy

 $\|\mathbf{H}_{C}^{\mathsf{T}}\|_{\infty}, \|\widehat{\mathbf{H}}_{C,\mathbf{x}}^{\mathsf{T}}\|_{\infty} \leq (m+1)^{d}, \quad (\mathbf{B} - (1,\mathbf{x}) \otimes \mathbf{QG})\widehat{\mathbf{H}}_{C,\mathbf{x}} = \mathbf{BH}_{C} - C(\mathbf{x})\mathbf{QG}.$

Gadget Matrix [33]. Let n, q be positive integers and $m = n \lceil \log_2 q \rceil$. The gadget matrix is $\mathbf{G} = \mathbf{g}^{\mathsf{T}} \otimes \mathbf{I}_n$, where $\mathbf{g}^{\mathsf{T}} = (2^0, 2^1, \dots, 2^{\lceil \log_2 q \rceil - 1})$. There exists an efficiently computable function $\mathbf{G}^{-1} : \mathbb{Z}_q^n \to \{0, 1\}^m$ such that $\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{u}) = \mathbf{u}$ for all $\mathbf{u} \in \mathbb{Z}_q^n$.

Assumption. We rely on the following assumption, a small-secret version of adaptive learning with errors (LWE), which itself is a natural variant of LWE first proposed in [36]:

Definition 4 (small-secret adaptive LWE). We suppress the security parameter λ and all the parameters are dependent on λ . Let n be the dimension, q the modulus, χ the error distribution, $m = n \lceil \log_2 q \rceil$, and **G** the gadget matrix. The smallsecret adaptive LWE assumption sALWE_{n,q,χ} states that $\mathsf{Exp}_{sALWE}^0 \approx \mathsf{Exp}_{sALWE}^1$, where $\mathsf{Exp}_{sALWE}^b(1^n, q, \chi)$ with adversary \mathcal{A} proceeds as follows:

- Setup. The challenger launches \mathcal{A} and receives $(1^{\ell}, 1^{m'})$ from it. The challenger samples $\mathbf{A} \stackrel{*}{\leftarrow} \mathbb{Z}_q^{n \times m'}, \mathbf{B} \stackrel{*}{\leftarrow} \mathbb{Z}_q^{n \times (\ell+1)m}$, and a uniformly random invertible $\mathbf{Q} \in \mathbb{Z}_q^{n \times n}$. It sends $\mathbf{A}, \mathbf{B}, \mathbf{Q}$ to \mathcal{A} .
- Challenge. \mathcal{A} submits $\mathbf{x} \in \{0, 1\}^{\ell}$. Depending on b,

$$\begin{split} \text{if } b &= 0: \qquad \boxed{\mathbf{s} \stackrel{\$}{\leftarrow} \chi^n}, \ \mathbf{e} \stackrel{\$}{\leftarrow} \chi^{m'}, \ \mathbf{f} \stackrel{\$}{\leftarrow} \chi^{(\ell+1)m}, \\ \mathbf{c}^{\mathsf{T}} &= \mathbf{s}^{\mathsf{T}} \mathbf{A} + \mathbf{e}^{\mathsf{T}}, \qquad \mathbf{d}^{\mathsf{T}} = \mathbf{s}^{\mathsf{T}} \left(\mathbf{B} - (1, \mathbf{x}) \otimes \mathbf{Q} \mathbf{G} \right) + \mathbf{f}^{\mathsf{T}}; \\ \text{if } b &= 1: \qquad \mathbf{c}^{\mathsf{T}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{m'}, \qquad \mathbf{d}^{\mathsf{T}} \stackrel{\$}{\leftarrow} \mathbb{Z}_q^{(\ell+1)m}. \end{split}$$

The challenger sends \mathbf{c}, \mathbf{d} to \mathcal{A} .

- Guess. \mathcal{A} outputs a bit, which is the outcome of the experiment.

Lemma 3 (small-secret adaptive LWE). The small-secret adaptive LWE assumption holds if the adaptive LWE assumption [36] holds for the same parameters.

The proof of Lemma 3 can be found in the full version [31].

Parameter Settings. We rely on the hardness of small-secret LWE with subexponential modulus-to-noise ratio. For some $0 < \delta < \frac{1}{2}$, the small-secret LWE assumption is assumed to be hard when the dimension is $n = \text{poly}(\lambda)$, the *prime* modulus is $q = O(2^{n^{\delta}})$, and the error distribution χ is the discrete Gaussian over \mathbb{Z} of width \overline{B}/λ truncated within $[-\overline{B}..\overline{B}]$ for $\overline{B} = \text{poly}(\lambda)$.³ Hereafter we default to these parameters.

3.3 Pairing Groups and Generic Asymmetric Pairing Group Model

We construct our ABE using pairing groups and prove its security in the generic pairing group model.

Pairing Groups. Throughout the paper, we use a sequence of pairing groups

$$\mathcal{G} = \{ (p_{\lambda}, G_{\lambda,1}, G_{\lambda,2}, G_{\lambda,T}, g_{\lambda,1}, g_{\lambda,2}, g_{\lambda,T}, e_{\lambda}) \}_{\lambda \in \mathbb{N}},$$

where $G_{\lambda,1}$ (resp. $G_{\lambda,2}, G_{\lambda,T}$) is a cyclic group generated by $g_{\lambda,1}$ (resp. $g_{\lambda,2}, g_{\lambda,T}$) of prime order $p_{\lambda} = 2^{\lambda^{\Theta(1)}}$ and $e_{\lambda} : G_{\lambda,1} \times G_{\lambda,2} \to G_{\lambda,T}$ is the pairing operation, satisfying $e_{\lambda}(g_{\lambda,1}^{a}, g_{\lambda,2}^{b}) = g_{\lambda,T}^{ab}$ for all integers a, b. We require the group operations as well as the pairing operation to be efficiently computable.

For a fixed security parameter λ , we denote $g_{\lambda,i}^x$ by $[\![x]\!]_i$ for $i \in \{1, 2, T\}$. The notation extends to matrices, $[\![A]\!]_i = g_{\lambda,i}^A$, where exponentiation is done component-wise. With these notations, the group operations are written additively and the pairing operation multiplicatively. For example, $[\![A]\!]_1 - \mathbf{B}[\![C]\!]_1 \mathbf{D} =$ $[\![A - \mathbf{B}\mathbf{C}\mathbf{D}]\!]_1$ and $[\![\mathbf{X}]\!]_2[\![\mathbf{Y}]\!]_1 = [\![\mathbf{X}\mathbf{Y}]\!]_T$.

Generic Asymmetric Pairing Group. The security of our ABE scheme holds in the generic asymmetric pairing group model (GGM), where the pairing groups can only be accessed via (non-unique) handles representing group elements and oracles for operating the handles. Due to space constraints, we refer the readers to the full version [31] for the formal definition of the version we use in this work.

3.4 Inner-Product Functional Encryption

Inner-product functional encryption schemes enable generating keys and ciphertexts tied to vectors. Decryption reveals the inner product and nothing more about the plaintext vector. In this work, we consider IPFE schemes based on pairing, where keys and ciphertexts are encoded in the two source groups and decryption recovers inner products encoded in the target group.

Definition 5 (group-based IPFE). Let \mathcal{G} be a sequence of pairing groups of order $\{p_{\lambda}\}_{\lambda \in \mathbb{N}}$. An inner-product functional encryption (IPFE) scheme based on \mathcal{G} consists of 4 efficient algorithms:

³ This truncation only introduces an exponentially small statistical error.

- Setup(1^λ, 1^N) takes as input the security parameter 1^λ and the vector dimension 1^N. It outputs a pair of master public/secret keys (impk, imsk).
- KeyGen(imsk, $[\![v]\!]_2$) takes as input the master secret key and a vector (encoded in G_2), and outputs a secret key isk.
- $\mathsf{Enc}(\mathsf{impk}, [\![\mathbf{u}]\!]_1)$ takes as input the master public key and a vector (encoded in G_1), and outputs a ciphertext ict.
- Dec(isk, [[v]]₂, ict) takes a secret key, the vector in the secret key, and a ciphertext as input, and is supposed to compute the inner product encoded in G_T.

The scheme is required to be correct, meaning that for all $\lambda, N \in \mathbb{N}, \mathbf{u}, \mathbf{v} \in \mathbb{Z}_{p_{\lambda}}^{N}$,

$$\Pr \begin{bmatrix} (\mathsf{impk},\mathsf{imsk}) \stackrel{\hspace{0.1em} \leftarrow \hspace{0.1em}}{} \mathsf{Setup}(1^{\lambda},1^{N}) \\ \mathsf{isk} \stackrel{\hspace{0.1em} \leftarrow \hspace{0.1em}}{} \mathsf{KeyGen}(\mathsf{imsk},[\![\mathbf{v}]\!]_{2}) : \mathsf{Dec}(\mathsf{isk},[\![\mathbf{v}]\!]_{2},\mathsf{ict}) = [\![\mathbf{u}^{\mathsf{T}}\mathbf{v}]\!]_{\mathsf{T}} \\ \mathsf{ict} \stackrel{\hspace{0.1em} \leftarrow \hspace{0.1em}}{} \mathsf{Enc}(\mathsf{impk},[\![\mathbf{u}]\!]_{1}) \end{bmatrix} = 1.$$

Definition 6 (key-succinct IPFE). An IPFE scheme (Definition 5) is (key-)succinct if the length of isk is a fixed polynomial in λ , independent of N.

Security. Our basic security notion is selective simulation:

Definition 7 (selective simulation [32,41]). A simulator for an IPFE scheme (Definition 5) consists of 3 efficient algorithms:

- Setup $(1^{\lambda}, 1^{N})$ takes the same input as Setup, and outputs simulated keys (impk, imsk).
- KeyGen(imsk, $[v]_2, [z_i]_2$) takes as input the simulated master secret key, a vector encoded in G_2 , and an inner product encoded in G_2 . It outputs a simulated key isk.
- Enc(imsk) takes as input the simulated master secret key. It outputs a simulated ciphertext ict.

The IPFE scheme is selectively simulation-secure if there exists a simulator such that $\mathsf{Exp}_{real} \approx \mathsf{Exp}_{sim}$, where $\mathsf{Exp}_{real}(1^{\lambda})$ or $\mathsf{Exp}_{sim}(1^{\lambda})$ with \mathcal{A} proceeds as follows:

- **Challenge.** The challenger launches $\mathcal{A}(1^{\lambda})$ and receives from it the vector dimension 1^N and the challenge vector $\mathbf{u} \in \mathbb{Z}_p^N$.
- **Setup.** The challenger runs

$$\begin{split} & \text{in } \mathsf{Exp}_{\mathrm{real}}: \quad (\mathsf{impk},\mathsf{imsk}) \xleftarrow{\hspace{0.1cm}} \mathsf{Setup}(1^{\lambda},1^{N}), \quad \mathsf{ict} \xleftarrow{\hspace{0.1cm}} \mathsf{Enc}(\mathsf{impk},\llbracket \mathbf{u} \rrbracket_{1}); \\ & \text{in } \mathsf{Exp}_{\mathrm{sim}}: \quad (\mathsf{impk},\widetilde{\mathsf{imsk}}) \xleftarrow{\hspace{0.1cm}} \widetilde{\mathsf{Setup}}(1^{\lambda},1^{N}), \quad \mathsf{ict} \xleftarrow{\hspace{0.1cm}} \widetilde{\mathsf{Enc}}(\widetilde{\mathsf{imsk}}); \end{split}$$

and sends impk, ict to \mathcal{A} .

- Query. The following is repeated for arbitrarily many rounds determined by \mathcal{A} : In each round, \mathcal{A} submits a vector $[\![\mathbf{v}_j]\!]_2$ encoded in G_2 . Upon receiving the query, the challenger runs

$$\begin{array}{ll} & \text{in } \mathsf{Exp}_{\text{real}} \text{:} & \text{isk}_{j} \stackrel{\hspace{0.1cm} \leftarrow \hspace{0.1cm}}{} \mathsf{KeyGen}(\mathsf{imsk}, \llbracket \mathbf{v}_{j} \rrbracket_{2}); \\ & \text{in } \mathsf{Exp}_{\text{sim}} \text{:} & \text{isk}_{j} \stackrel{\hspace{0.1cm} \leftarrow \hspace{0.1cm}}{} \widetilde{\mathsf{KeyGen}}(\widetilde{\mathsf{imsk}}, \llbracket \mathbf{v}_{j} \rrbracket_{2}, \mathbf{u}^{\mathsf{T}} \llbracket \mathbf{v}_{j} \rrbracket_{2}); \end{array}$$

and sends isk_j to \mathcal{A} .

- **Guess.** \mathcal{A} outputs a bit b, which is the output of the experiment.

Lemma 4 ([2,41]). Assuming the MDDH assumption (true in GGM), there exists a succinct selectively simulation-secure IPFE scheme. Its components have sizes

$ impk = k(k+1+N) G_1 ,$	$ imsk = (k+1)N\log_2 p,$	
$ isk = (k+1) G_2 ,$	$ ict = (k+1+N) G_1 $,

where p is the modulus, k is the MDDH parameter (can be 1 in GGM), N is the dimension, and $|G_i|$ is the bit-length of an element in G_i .

4 Computational Secret Sharing with Adaptive Security

Secret sharing schemes have been used extensively to construct ABE schemes. The seminal work of [24] and a long line of follow-up works ([5,28–30,35] to name a few) used *linear* secret sharing schemes to construct ABE schemes in pairing groups. Policies with polynomial-sized shares and information-theoretic security are in NC [8,9,15,26,34].

The works of [3,4] introduced the notion of *nearly linear* secret sharing with *computational* security. The relaxations enabled greater expressiveness and better efficiency. Assuming LWE, such a scheme exists for all polynomial-sized circuits [3,4,11,23] and the shares are *succinct*, i.e., they only grow with the circuit depth, but not the circuit size. However, the scheme is only *selectively* secure. Furthermore, due to technical reasons, when combined with pairing to obtain ABE, it only applies to Boolean formulae (equivalent to 5-PBP).

This work follows the blueprint of [3,4] for the notions of secret sharing schemes, but departs from them in three important aspects. First, we consider a different security notion, *adaptive non-annihilability*, which is incomparable⁴ to selective pseudorandomness considered in [3,4] and enables us to prove adaptive security of ABE. Second, we further relax the linearity requirement so that it could apply to KP-ABE for polynomial-sized circuits. Third, we refine the syntax to separate encodings of input and function.

Definition 8 (secret sharing). Let $\mathcal{F} = \{\mathcal{F}_{\lambda,\ell,\text{param}}\}_{\lambda,\ell\in\mathbb{N},\text{param}}$ be an ensemble of Boolean function families such that for all $\lambda, \ell\in\mathbb{N}$ and param, every $f\in\mathcal{F}_{\lambda,\ell,\text{param}}$ is a function mapping $\{0,1\}^{\ell}$ to $\{0,1\}$. A secret sharing scheme for \mathcal{F} consists of 4 efficient algorithms:

- Setup(1^λ, 1^ℓ, param) takes the security parameter 1^λ, the input length 1^ℓ, and additional parameters param as input. It outputs some public parameter pp.
- ShareX(pp) takes the public parameter pp as input. It outputs $1 + 2\ell$ shares, $L_0, \{L_i^b\}_{i \in [\ell]}^{b \in \{0,1\}}$, and some shared randomness r. For $x \in \{0,1\}^{\ell}$, we denote by $L^{\mathbf{x}}$ the set of shares $L_0, \{L_i^{\mathbf{x}[i]}\}_{i \in [\ell]}$.

⁴ It is stronger in that it is adaptive, but weaker in that the shares are not necessarily pseudorandom.

- ShareF(pp, f, μ, r) takes the public parameter pp, a function $f \in \mathcal{F}_{\lambda,\ell,\text{param}}$, a secret $\mu \in \{0, 1\}$, and the shared randomness r (output by ShareX) as input. It outputs a share L_f .
- Recon(pp, $f, \mathbf{x}, L_f, L^{\mathbf{x}}$) takes the public parameter pp, the Boolean function $f \in \mathcal{F}_{\lambda,\ell,\text{param}}$, the input $\mathbf{x} \in \{0,1\}^{\ell}$ to f, and the shares $L_f, L^{\mathbf{x}}$ as input. It is supposed to recover the secret μ if $f(\mathbf{x}) = 0.5^{5}$

The scheme is required to be correct, i.e., for all $\lambda, \ell \in \mathbb{N}$, param, $\mathbf{x} \in \{0, 1\}^{\ell}$, $f \in \mathcal{F}_{\lambda,\ell,\text{param}}, \mu \in \{0, 1\}$ such that $f(\mathbf{x}) = 0$, it holds that

$$\Pr \begin{bmatrix} \mathsf{pp} \stackrel{\hspace{0.1em} {\scriptstyle \bullet \hspace{0.1em}}}{\scriptstyle \mathsf{Setup}(1^{\lambda}, 1^{\ell}, \mathsf{param}) \\ (L_0, \{L_i^b\}_{i \in [\ell]}^{b \in \{0,1\}}, r) \stackrel{\hspace{0.1em} {\scriptstyle \bullet \hspace{0.1em}}}{\scriptstyle \bullet \hspace{0.1em}} \mathsf{Share} \mathsf{X}(\mathsf{pp}) & : \mathsf{Recon}(\mathsf{pp}, f, \mathbf{x}, L_f, L^{\mathbf{x}}) = \mu \\ L_f \stackrel{\hspace{0.1em} {\scriptstyle \bullet \hspace{0.1em}}}{\scriptstyle \bullet \hspace{0.1em}} \mathsf{Share} \mathsf{F}(\mathsf{pp}, f, \mu, r) \end{bmatrix} = 1.$$

Definition 9 (succinct shares). A secret sharing scheme is succinct if the size of each share output by ShareX, ShareF is a fixed polynomial in λ , independent of the length of \mathbf{x} or the description size of f, i.e., $|L_f|, |L_0|, |L_i^b|$ are all $poly(\lambda, |param|)$, where $i \in [\ell], b \in \{0, 1\}$.⁶

While correctness (Definition 8) and succinctness (Definition 9) are defined similarly to that of [3], our linearity and security notions are different.

4.1 Secret Sharing for Bounded-Depth Circuits from Adaptive LWE

In our KP-ABE construction, we need a secret sharing scheme with two linearity properties. The first is a relaxation of the nearly linear reconstruction requirement in [3]. requirement on reconstruction. Our relaxed version (Definition 10) only stipulates it to be linear in \mathbf{L}_{f} (and possibly non-linear in $L^{\mathbf{x}}$).

Definition 10 (weakly nearly linear reconstruction). A secret sharing scheme (Definition 8) is weakly nearly linear if it satisfies the following requirements:

- Let $\{p_{\lambda}\}_{\lambda \in \mathbb{N}}$ be a sequence of prime numbers. $\mathbf{L}_{f} = L_{f}$ is a vector over $\mathbb{Z}_{p_{\lambda}}$.
- There is an efficient coefficient-finding algorithm FindCoef(pp, $f, \mathbf{x}, L^{\mathbf{x}}$), taking as input the public parameter pp, a Boolean function $f \in \mathcal{F}_{\lambda,\ell,\text{param}}$, an input $\mathbf{x} \in \{0,1\}^{\ell}$ to f, and the shares $L^{\mathbf{x}}$. It outputs an affine function γ and a noise bound 1^{B} . For all $\lambda, \ell \in \mathbb{N}$, param, $\mathbf{x} \in \{0,1\}^{\ell}, f \in \mathcal{F}_{\lambda,\ell,\text{param}}, \mu \in \{0,1\}$ such that $f(\mathbf{x}) = 0$, it holds that

$$\Pr \begin{bmatrix} \mathsf{pp} \overset{\$}{\leftarrow} \mathsf{Setup}(1^{\lambda}, 1^{\ell}, \mathsf{param}) & \\ (L_0, \{L_i^b\}_{i \in [\ell]}^{b \in \{0,1\}}, r) \overset{\$}{\leftarrow} \mathsf{ShareX}(\mathsf{pp}) & 4B + 1 < p_{\lambda} \text{ and} \\ & : \exists e \in [-B..B] \text{s.t.} \\ \mathbf{L}_f \leftarrow \mathsf{ShareF}(\mathsf{pp}, f, \mu, r) & \gamma(\mathbf{L}_f) = \mu \lfloor p/2 \rceil + e \end{bmatrix} = 1$$

$$(\gamma, 1^B) \overset{\$}{\leftarrow} \mathsf{FindCoef}(\mathsf{pp}, f, \mathbf{x}, L^{\mathbf{x}})$$

⁵ We use $f(\mathbf{x}) = 0$ to express authorization.

⁶ There are $2|\mathbf{x}| + 2$ shares, so the total share size is linear in the length of \mathbf{x} .

The second is an additional linearity requirement on ShareF.

Definition 11 (linear function sharing). Let $\{p_{\lambda}\}_{\lambda \in \mathbb{N}}$ be a sequence of primes. A secret sharing scheme (Definition 8) has linear function sharing if $\mathbf{r} = r$ is a vector over $\mathbb{Z}_{p_{\lambda}}$ and ShareF(pp, f, μ, \mathbf{r}) is deterministic and linear in (μ, \mathbf{r}) .

A (weakly) nearly linear scheme is by definition correct. Given FindCoef, we let Recon call FindCoef to obtain γ , B and output the unique $\mu \in \{0, 1\}$ satisfying $\gamma(\mathbf{L}_f) - \mu \lfloor p/2 \rfloor \in [-B..B]$. The constructed Recon is efficient and correct. Since Recon is implied by FindCoef, we will only specify FindCoef and omit Recon when constructing (weakly) nearly linear secret sharing schemes.

Security. We consider a new security notion called *non-annihilability*. Unlike [3], which fixes the choice of policy f before Setup is run, we allow the adversary to adaptively choose f after seeing the public parameters pp and the input shares $L^{\mathbf{x}}$. Another difference is that instead of requiring all shares $(L_f, L^{\mathbf{x}})$ to look random, we only require that efficient adversaries cannot find a non-trivial affine function (potentially dependent on $L^{\mathbf{x}}$) that evaluates to zero on \mathbf{L}_f . This notion suffices for the security proofs of our KP-ABE scheme.

Definition 12 (non-annihilability for \mathbf{L}_f). Let $\{p_{\lambda}\}_{\lambda \in \mathbb{N}}$ be a sequence of prime numbers. A secret sharing scheme (Definition 8) is adaptively non-annihilable for \mathbf{L}_f if the output \mathbf{L}_f of ShareF is a vector over $\mathbb{Z}_{p_{\lambda}}$ and all efficient adversary wins Exp_{ANN-f} with negligible probability, where in $\mathsf{Exp}_{ANN-f}^{\mathcal{A}}(1^{\lambda})$, the adversary \mathcal{A} interacts with the challenger as follows:

- Setup. The challenger launches $\mathcal{A}(1^{\lambda})$ and receives from it the input length 1^{ℓ} and the additional parameter param. The challenger sets up the system by running pp $\stackrel{\text{\tiny{\$}}}{\stackrel{\text{\tiny{\$}}}$ Setup $(1^{\lambda}, 1^{\ell}, \text{param})$, and sends pp to \mathcal{A} .
- Share. \mathcal{A} first submits an input $\mathbf{x} \in \{0, 1\}^{\ell}$. Upon receiving it, the challenger creates the input shares by running $(L_0, \{L_i^b\}_{i \in [\ell]}^{b \in \{0,1\}}, r) \stackrel{\text{s}}{\leftarrow} \mathsf{ShareX}(\mathsf{pp})$ and sends $L^{\mathbf{x}}$ to \mathcal{A} .
- Challenge. \mathcal{A} outputs a Boolean function $f \in \mathcal{F}_{\lambda,\ell}$, a message bit $\mu \in \{0,1\}$, and an affine function γ . Upon receiving them, the challenger runs $\mathbf{L}_f \stackrel{s}{\leftarrow} \mathsf{ShareF}(\mathsf{pp}, f, \mu, r)$ and determines the outcome of the experiment. \mathcal{A} wins if i) $f(\mathbf{x}) = 1$; ii) γ is not the zero function; and iii) $\gamma(\mathbf{L}_f) = 0$. Otherwise, \mathcal{A} loses.

Furthermore, a secret sharing scheme is selectively non-annihilable if it satisfies the above conditions, with the change that the adversary must choose the input \mathbf{x} before receiving pp.

We now construct a succinct secret sharing scheme, satisfying the above linearity and adaptive annihilability for bounded-depth circuits from small-secret adaptive LWE. Our construction is based on the attribute-based laconic function evaluation scheme [11,36].

Construction 1 (secret sharing for circuits). Let *n* be the LWE dimension, $p = 2^{\omega(\log \lambda)}$ a fixed prime modulus, (the LWE parameters will be chosen during Setup). We construct a weakly nearly linear and succinct secret sharing scheme, with linear function sharing, for the family of bounded-depth circuits (see Definition 2):

 $\mathsf{Ckt}_{\lambda,\ell,d} = \{ \text{Boolean circuit } C : \{0,1\}^{\ell} \to \{0,1\} \text{ of depth at most } d \},\$

where $d \leq \frac{p^{\delta/4} - \log_2 p}{(1+\delta^{-1})\Theta(1)}$. Let (EvalC, EvalCX) be the algorithms in Lemma 2.

- $\mathsf{Setup}(1^{\lambda}, 1^{\ell}, 1^{d})$ takes the input length ℓ in unary as input. It sets

$$n = \overline{B} = \left((d+1)(\delta^{-1}+1) + \log_2 p + \mathcal{O}(d) \right)^{2/\delta}, \quad q = 2^{n^{\delta}}, \quad m = n \lceil \log_2 q \rceil,$$

and picks χ to be \overline{B} -bounded. It next samples and sets

$$\mathbf{a} \stackrel{\hspace{0.1em}\mathsf{\scriptscriptstyle\$}}{\leftarrow} \mathbb{Z}_q^n, \quad \mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_\ell \stackrel{\hspace{0.1em}\mathsf{\scriptscriptstyle\$}}{\leftarrow} \mathbb{Z}_q^{n \times m}, \quad \mathbf{B} = (\mathbf{A}_0, \mathbf{A}_1, \dots, \mathbf{A}_\ell).$$

It finally samples a random invertible matrix $\mathbf{Q} \in \mathbb{Z}_q^{n \times n}$, and outputs $pp = (n, q, m, \overline{B}, \chi, \mathbf{a}, \mathbf{B}, \mathbf{Q})$.

Note: Recall that δ is a constant depending on the underlying adaptive LWE assumption. The choice of n, \overline{B}, q are subject to the requirement of the underlying adaptive LWE assumption as well as correctness and efficiency of the scheme. They satisfy $q/\overline{B} \geq (m+1)^{d+1}$ and $4((n+1)\overline{B}+3)+1 < p$.

- $\mathsf{ShareX}(pp)$ takes the public parameter pp as input. It samples and sets

$$\mathbf{s} \stackrel{\mathbf{s}}{\leftarrow} \chi^{n}, \quad \mathbf{e}_{0}, \mathbf{e}_{1}, \dots, \mathbf{e}_{\ell} \stackrel{\mathbf{s}}{\leftarrow} \chi^{m},$$
$$\mathbf{L}_{0} = \mathbf{s}^{\mathsf{T}}(\mathbf{A}_{0} - \mathbf{Q}\mathbf{G}) + \mathbf{e}_{0}^{\mathsf{T}}, \quad \{\mathbf{L}_{i}^{b} = \mathbf{s}^{\mathsf{T}}(\mathbf{A}_{i} - b\mathbf{Q}\mathbf{G}) + \mathbf{e}_{i}^{\mathsf{T}}\}_{i \in [\ell]}^{b \in \{0,1\}},$$

and outputs $(\mathbf{L}_0, {\{\mathbf{L}_i^b\}}_{i \in [\ell]}^{b \in \{0,1\}}, \mathbf{s}).$

- ShareF(pp, f, μ, \mathbf{s}) takes as input the public parameter pp, some $f \in \mathsf{Ckt}_{\lambda,\ell,d}$, a secret bit $\mu \in \{0, 1\}$, and the shared randomness \mathbf{s} . It runs $\mathbf{H}_f \leftarrow \mathsf{EvalC}(\mathbf{B}, \mathbf{Q}, f)$, and sets and outputs

$$\mathbf{L}_f = \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B} \mathbf{H}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p + \mu \lfloor p/2 \rceil, \text{ where } \lfloor x \rceil_p = \lfloor px/q \rceil.$$

Note: The scheme indeed has linear function sharing (Definition 11) because ShareF is a deterministic linear function over μ , **s** with coefficients $\lfloor p/2 \rfloor$, $\lfloor \mathbf{BH}_{f}\mathbf{G}^{-1}(\mathbf{a}) \rfloor_{p}$. The scheme is also succinct as \mathbf{L}_{f} contains 1 element in \mathbb{Z}_{p} , and each share output by ShareX contains m elements in \mathbb{Z}_{q} . Note that m is a fixed polynomial in λ , d and is independent of the description size of f and the input length ℓ .

- FindCoef(pp, $f, \mathbf{x}, \mathbf{L}^{\mathbf{x}}$) takes as input the public parameter pp, some $\mathbf{x} \in \{0, 1\}^{\ell}$, some $f \in \mathsf{Ckt}_{\lambda,\ell,d}$, and the shares $\mathbf{L}^{\mathbf{x}}$. If $f(\mathbf{x}) = 1$, it outputs \perp and terminates. Otherwise, it runs $\widehat{\mathbf{H}}_{f,\mathbf{x}} \leftarrow \mathsf{EvalCX}(\mathbf{B}, \mathbf{Q}, f, \mathbf{x})$, and defines

$$\gamma(\mathbf{L}_f) = \mathbf{L}_f - [\mathbf{L}^{\mathbf{x}} \widehat{\mathbf{H}}_{f,\mathbf{x}} \mathbf{G}^{-1}(\mathbf{a})]_p, \quad B = (n+1)\overline{B} + 3,$$

The algorithm outputs $(\gamma, 1^B)$.

Note: The procedure is indeed efficient since n, \overline{B} are polynomials in λ, d . We show that FindCoef is correct, i.e., if $f(\mathbf{x}) = 0$, then $4B + 1 \leq p$ and $\gamma(\mathbf{L}_f) = \mu \lfloor p/2 \rfloor + e$ for some $e \in [-B, B]$. First, by the choice of n, \overline{B} ,

$$4B + 1 = 4((n+1)\overline{B} + 3) + 1 \le p.$$

Next, by construction we have

~

$$\gamma(\mathbf{L}_{f}) = \mathbf{L}_{f} - [\mathbf{L}^{\mathbf{x}} \widehat{\mathbf{H}}_{f,\mathbf{x}} \mathbf{G}^{-1}(\mathbf{a})]_{p}$$

$$= \underbrace{\mathbf{s}^{\mathsf{T}} [\mathbf{B}\mathbf{H}_{f} \mathbf{G}^{-1}(\mathbf{a})]_{p} + \mu[p/2]}_{\mathbf{L}_{f}}$$

$$- [\underbrace{(\mathbf{s}^{\mathsf{T}}(\mathbf{B} - (1, \mathbf{x}) \otimes \mathbf{Q}\mathbf{G}) + (\mathbf{e}_{0}^{\mathsf{T}}, \mathbf{e}_{1}^{\mathsf{T}}, \dots, \mathbf{e}_{\ell}^{\mathsf{T}})}_{\mathbf{L}^{\mathbf{x}}} \widehat{\mathbf{H}}_{f,\mathbf{x}} \mathbf{G}^{-1}(\mathbf{a})]_{p}$$

(Lemma 2) = $\mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B} \mathbf{H}_f \mathbf{G}^{-1}(\mathbf{a}) \rceil_p + \mu \lfloor p/2 \rceil$

$$- \lfloor \mathbf{s}^{\mathsf{T}} (\mathbf{B}\mathbf{H}_{f} - \underbrace{f(\mathbf{x})}_{=0} \mathbf{Q}\mathbf{G})\mathbf{G}^{-1}(\mathbf{a}) + \underbrace{(\mathbf{e}_{0}^{\mathsf{T}}, \mathbf{e}_{1}^{\mathsf{T}}, \dots, \mathbf{e}_{\ell}^{\mathsf{T}})\widehat{\mathbf{H}}_{f,\mathbf{x}}\mathbf{G}^{-1}(\mathbf{a})}_{=e_{f}} \rceil_{p}$$
$$= \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) \rceil_{p} + \mu \lfloor p/2 \rceil - \lfloor \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) + e_{f} \rceil_{p}$$

Since $\mathbf{G}^{-1}(\mathbf{a}) \in \{0,1\}^m$, by the definition of EvalCX (Lemma 2), we have

$$|e_f| \le m \cdot \|\widehat{\mathbf{H}}_{f,\mathbf{x}}^{^{\mathsf{T}}}\|_{\infty} \cdot \|(\mathbf{e}_0^{^{\mathsf{T}}}, \mathbf{e}_1^{^{\mathsf{T}}}, \dots, \mathbf{e}_{\ell}^{^{\mathsf{T}}})^{^{\mathsf{T}}}\|_{\infty} \le (m+1)^{(d+1)}\overline{B}$$

Note that we can break a rounded sum into a sum of individually rounded terms, at the expense of some rounding errors:

$$[\mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) + e_{f}] = [\mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a})]_{p} + [e_{f}]_{p} + \varepsilon, \text{ where } |\varepsilon| \leq 3, \\ [\mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a})]_{p} = \mathbf{s}^{\mathsf{T}}[\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a})]_{p} + e_{s}, \text{ where } |e_{s}| \leq n \cdot ||\mathbf{s}||_{\infty} \leq n\overline{B}.$$

Finally, we have

$$\gamma(\mathbf{L}_{f}) = \mu \lfloor p/2 \rceil + \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) \rceil_{p} - \lfloor \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) + e_{f} \rceil_{p}$$
$$= \mu \lfloor p/2 \rceil + \mathbf{s}^{\mathsf{T}} \lfloor \mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) \rceil_{p} - \lfloor \mathbf{s}^{\mathsf{T}}\mathbf{B}\mathbf{H}_{f}\mathbf{G}^{-1}(\mathbf{a}) \rceil_{p} - \lfloor e_{f} \rceil_{p} - \varepsilon$$
$$= \mu \lfloor p/2 \rceil \underbrace{-e_{s} - \lfloor e_{f} \rceil_{p} - \varepsilon}_{-e_{s}}.$$

By the definition of e_f, e_s, ε , and the setting of q, we have

$$|e| \le |e_s| + |\lfloor e_f \rceil_p| + |\varepsilon| \le \left\lceil \frac{(m+1)^{(d+1)}}{q/p} \overline{B} \right\rceil + n\overline{B} + 3 \le B.$$

Efficiency. In the above construction, the public parameters **pp** mainly consists of three matrices $\mathbf{a} \in \mathbb{Z}_q^n$, $\mathbf{B} \in \mathbb{Z}_q^{n \times (m\ell)}$, $\mathbf{Q} \in \mathbb{Z}_q^{n \times n}$, where $n = \text{poly}(\lambda, d)$, $q = 2^{n^{\delta}}$, and $m = n \lceil \log q \rceil = \text{poly}(\lambda, d)$. Therefore, the bit length of **pp** is $|\mathbf{pp}| = \text{poly}(\lambda, d)$.

 ℓ . The shares \mathbf{L}_0 and $\{\mathbf{L}_i^b\}$ are $2\ell + 1$ vectors in \mathbb{Z}_q^m . Therefore $|\mathbf{L}_0| = |\mathbf{L}_i^b| = \text{poly}(\lambda, d)$. Finally, \mathbf{L}_f is a single element in \mathbb{Z}_p , where $p = 2^{\omega(\log \lambda)}$. Therefore, $|\mathbf{L}_f| = \text{poly}(\lambda)$.

We next state non-annihilability security for \mathbf{L}_f of the scheme. The proof can be found in the full version [31].

Proposition 5. Assuming the small-secret adaptive LWE assumption, Construction 1 is non-annihilable for \mathbf{L}_f .

5 KP-ABE for Bounded-Depth Circuits

In this section, we combine a succinct and weakly nearly linear secret sharing scheme that has linear function sharing, with a succinct and selectively simulationsecure IPFE scheme to obtain a compact and adaptively secure KP-ABE scheme.

Construction 2 (KP-ABE). All variables x_{λ} are indexed by λ . For simplicity of notations, we suppress λ in subscripts. Our construction uses the following two ingredients:

- A group based IPFE scheme (IPFE.Setup, IPFE.KeyGen, IPFE.Enc, IPFE.Dec) with modulus p given by Lemma 4.
- A secret sharing scheme (SS.Setup, SS.ShareX, SS.ShareF, SS.FindCoef) for bounded-depth circuits as in Construction 1. Recall that the scheme has three properties. First, the shares are succinct: \mathbf{L}_0 and \mathbf{L}_i^b are vectors in \mathbb{Z}_q of length $m = \text{poly}(\lambda, d)$, and \mathbf{L}_C is a single element in \mathbb{Z}_p . Second, the scheme has weakly nearly linear reconstruction: the algorithm SS.FindCoef outputs an affine function γ over \mathbf{L}_C that approximately evaluates to $\mu \lfloor p/2 \rfloor$. Third, the scheme has linear function sharing: SS.ShareF_{SS.pp,C}(\cdot, \cdot) is a deterministic linear function over \mathbb{Z}_p .

Our KP-ABE for circuits (see Definition 2) works as follows:

- $\mathsf{Setup}(1^{\lambda}, P)$ takes as input the security parameter λ in unary, and a predicate $P \in \mathsf{Ckt}$. Let ℓ, d be the attribute length and depth for P. The algorithm runs and sets

$$\begin{split} & \mathsf{SS.pp} \xleftarrow{\hspace{0.15cm}} \mathsf{SS.Setup}(1^{\lambda}, 1^{\ell}, 1^{d}), \\ & (\mathsf{impk}, \mathsf{imsk}) \xleftarrow{\hspace{0.15cm}} \mathsf{IPFE.Setup}(1^{\lambda}, 1^{N}) \text{ for dimension } N = n+1, \\ & \mathsf{mpk} = (\mathsf{SS.pp}, \mathsf{impk}), \quad \mathsf{msk} = \mathsf{imsk}. \end{split}$$

It outputs $\mathsf{mpk}, \mathsf{msk}.$

- KeyGen(msk, C) takes as input the master secret key msk and a policy $C \in \mathsf{Ckt}_{\ell,d}$. Since the secret sharing scheme has linear function sharing (Definition 11), the SS.ShareF_{SS.pp,C}(\cdot, \cdot) function is a deterministic linear function with coefficients $\mathbf{c} = (c_{\mu}, \mathbf{c_r})$. The KeyGen algorithm samples $\delta \stackrel{*}{\leftarrow} \mathbb{Z}_p \setminus \{0\}$, runs

isk \leftarrow IPFE.KeyGen(imsk, $[\![\delta \mathbf{c}]\!]_2$),

and outputs $\mathsf{sk} = (\llbracket \delta \rrbracket_2, \mathsf{isk})$ as the secret key for C.

- $\mathsf{Enc}(\mathsf{mpk}, \mathbf{x}, \mu)$ takes as input the master public key mpk , an attribute $\mathbf{x} \in \{0, 1\}^{\ell}$, and a message $\mu \in \{0, 1\}$. The algorithm runs

$$(\mathbf{L}_0, \{\mathbf{L}_i^b\}_{i \in [\ell]}^{b \in \{0,1\}}, \mathbf{r}) \stackrel{\hspace{0.1em} \ast}{\leftarrow} \mathsf{SS.ShareX}(\mathsf{SS.pp}), \quad \mathsf{ict} \stackrel{\hspace{0.1em} \ast}{\leftarrow} \mathsf{IPFEEnc}(\mathsf{impk}, \llbracket(\mu, \mathbf{r})\rrbracket_1),$$

and outputs $ct = (\mathbf{L}^{\mathbf{x}}, ict)$.

- Dec(mpk, sk, C, ct, x) takes as input the master public key mpk, a secret key sk, its associated policy C, a ciphertext ct, and its associated attribute x. If $P(\mathbf{x}, C) = 0$, the algorithm outputs \perp and terminates. Otherwise, it parses sk = ($[\![\delta]\!]_2$, isk), and computes the coefficients $\mathbf{c} = (c_{\mu}, \mathbf{c_r})$ for ShareF_{SS.pp,C}(\cdot, \cdot) as in KeyGen. The algorithm next parses ct into $\mathbf{L}^{\mathbf{x}}$, ict, and runs

$$\Lambda_C \stackrel{\text{\tiny{\$}}}{\leftarrow} \mathsf{IPFE}.\mathsf{Dec}(\mathsf{isk}, \llbracket \delta \rrbracket_2 \mathbf{c}, \mathsf{ict}), \quad (\gamma, 1^B) \stackrel{\text{\tiny{\$}}}{\leftarrow} \mathsf{SS}.\mathsf{FindCoef}(\mathsf{SS}.\mathsf{pp}, C, \mathbf{x}, \mathbf{L}^{\mathbf{x}}).$$

The algorithm applies the affine function γ homomorphically in the exponent of $G_{\rm T}$ to compute $\gamma(\Lambda_C)$. It then finds and outputs the unique $\mu' \in \{0, 1\}$ (as the decrypted message) such that $\gamma(\Lambda_C) = \llbracket \mu' \lfloor p/2 \rceil + e \rrbracket_1 \llbracket \delta \rrbracket_2$, for some $e \in [-B..B]$, by enumerating over all possible e.

Note: We show that the scheme is correct. By the correctness of IPFE and by linear function sharing of the secret sharing scheme, we have

$$\Lambda_C = \llbracket \delta(c_{\mu} \cdot \mu + \mathbf{c_r} \cdot \mathbf{r}) \rrbracket_{\mathrm{T}} = \llbracket \delta \mathsf{SS}.\mathsf{ShareF}_{\mathsf{SS}.\mathsf{pp},C}(\mu, \mathbf{r}) \rrbracket_{\mathrm{T}} = \llbracket \delta \mathbf{L}_C \rrbracket_{\mathrm{T}}.$$

Therefore, $\gamma(\Lambda_C) = [\![\delta\gamma(\mathbf{L}_C)]\!]_{\mathrm{T}} = [\![\gamma(\mathbf{L}_C)]\!]_1[\![\delta]\!]_2$. By the correctness of the weakly nearly linear secret sharing scheme, the decryption algorithm outputs the correct bit $\mu' = \mu$.

Efficiency. By Lemma 4, for MDDH dimension $k = \text{poly}(\lambda)$ and input vector length N = n + 1, the IPFE components have bit lengths $|\mathsf{impk}|, |\mathsf{imsk}|, |\mathsf{ict}| = \text{poly}(\lambda, d), |\mathsf{isk}| = \text{poly}(\lambda)$. Also recall that the secret sharing components have bit lengths $|\mathsf{SS.pp}| = \text{poly}(\lambda, d) \cdot \ell, |\mathbf{L}_0| = |\mathbf{L}_i^b| = \text{poly}(\lambda, d), |\mathbf{L}_C| = \text{poly}(\lambda)$. In the above construction,

- the master public key consists of SS.pp and impk, hence has bit length $|\mathsf{mpk}| = |\mathsf{SS.pp}| + |\mathsf{impk}| = \mathrm{poly}(\lambda, d) \cdot \ell.$
- The master secret key consists of imsk, hence has bit length $|\mathsf{msk}| = |\mathsf{imsk}| = \mathrm{poly}(\lambda, d).$
- A secret key consists of a single isk, and $[\![\delta]\!]_2$ in G_2 , hence has bit length $|\mathsf{sk}| = |\mathsf{isk}| + |G_2| = \operatorname{poly}(\lambda)$.
- A ciphertext consists of a single ict, and $\ell + 1$ shares, hence has bit length $|\mathsf{ct}| = |\mathsf{ict}| + (\ell + 1)|\mathbf{L}_0| = \operatorname{poly}(\lambda, d) \cdot \ell.$

We now state adaptive IND-CPA security of the scheme. The proof can be found in the full version [31].

Proposition 6. Suppose in Construction 2, the IPFE scheme is selectively simulation-secure, and the secret sharing scheme is non-annihilable for \mathbf{L}_f . Then the constructed KP-ABE scheme is adaptively IND-CPA in GGM.

Acknowledgement. The authors were supported by NSF grants CNS-1528178, CNS-1929901, CNS-1936825 (CAREER), CNS-2026774, a Hellman Fellowship, a JP Morgan AI Research Award, the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-15-C-0236, and a subcontract No. 2017-002 through Galois. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government. The authors thank the anonymous reviewers for their valuable comments.

References

- Abdalla, M., Bourse, F., De Caro, A., Pointcheval, D.: Simple functional encryption schemes for inner products. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 733– 751. Springer, Heidelberg (Mar / Apr 2015).
- Agrawal, S., Libert, B., Stehlé, D.: Fully secure functional encryption for inner products, from standard assumptions. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part III. LNCS, vol. 9816, pp. 333–362. Springer, Heidelberg (Aug 2016).
- Agrawal, S., Wichs, D., Yamada, S.: Optimal broadcast encryption from LWE and pairings in the standard model. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part I. LNCS, vol. 12550, pp. 149–178. Springer, Heidelberg (Nov 2020).
- Agrawal, S., Yamada, S.: Optimal broadcast encryption from pairings and LWE. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part I. LNCS, vol. 12105, pp. 13–43. Springer, Heidelberg (May 2020).
- Attrapadung, N.: Dual system encryption framework in prime-order groups via computational pair encodings. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part II. LNCS, vol. 10032, pp. 591–623. Springer, Heidelberg (Dec 2016).
- Attrapadung, N., Libert, B., de Panafieu, E.: Expressive key-policy attribute-based encryption with constant-size ciphertexts. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 90–108. Springer, Heidelberg (Mar 2011).
- Attrapadung, N., Tomida, J.: Unbounded dynamic predicate compositions in ABE from standard assumptions. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part III. LNCS, vol. 12493, pp. 405–436. Springer, Heidelberg (Dec 2020).
- 8. Beimel, A.: Secure Schemes for Secret Sharing and Key Distribution. Ph.D. thesis, Technion–Israel Institute of Technology (1996)
- 9. Berkowitz, S.J.: On computing the determinant in small parallel time using a small number of processors. Information Processing Letters **18**(3), 147–150 (1984).
- Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: 2007 IEEE Symposium on Security and Privacy. pp. 321–334. IEEE Computer Society Press (May 2007).
- Boneh, D., Gentry, C., Gorbunov, S., Halevi, S., Nikolaenko, V., Segev, G., Vaikuntanathan, V., Vinayagamurthy, D.: Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 533–556. Springer, Heidelberg (May 2014).
- Boneh, D., Sahai, A., Waters, B.: Functional encryption: Definitions and challenges. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 253–273. Springer, Heidelberg (Mar 2011).

- Brakerski, Z., Tsabary, R., Vaikuntanathan, V., Wee, H.: Private constrained PRFs (and more) from LWE. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 264–302. Springer, Heidelberg (Nov 2017).
- Brakerski, Z., Vaikuntanathan, V.: Constrained key-homomorphic PRFs from standard lattice assumptions - or: How to secretly embed a circuit in your PRF. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 1–30. Springer, Heidelberg (Mar 2015).
- Buntrock, G., Damm, C., Hertrampf, U., Meinel, C.: Structure and importance of logspace-MOD class. Mathematical Systems Theory 25(3), 223–237 (1992).
- Chen, Y.H., Chung, K.M., Liao, J.J.: On the complexity of simulating auxiliary input. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 371–390. Springer, Heidelberg (Apr / May 2018).
- Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS. pp. 40–49. IEEE Computer Society Press (Oct 2013).
- Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 467– 476. ACM Press (Jun 2013).
- Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 75–92. Springer, Heidelberg (Aug 2013).
- Goldwasser, S., Kalai, Y.T., Popa, R.A., Vaikuntanathan, V., Zeldovich, N.: How to run turing machines on encrypted data. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 536–553. Springer, Heidelberg (Aug 2013).
- Gorbunov, S., Vaikuntanathan, V., Wee, H.: Functional encryption with bounded collusions via multi-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 162–179. Springer, Heidelberg (Aug 2012).
- Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute-based encryption for circuits. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 545–554. ACM Press (Jun 2013).
- Gorbunov, S., Vinayagamurthy, D.: Riding on asymmetry: Efficient ABE for branching programs. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 550–574. Springer, Heidelberg (Nov / Dec 2015).
- 24. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for finegrained access control of encrypted data. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006. pp. 89–98. ACM Press (Oct / Nov 2006). Available as Cryptology ePrint Archive Report 2006/309
- Jetchev, D., Pietrzak, K.: How to fake auxiliary input. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 566–590. Springer, Heidelberg (Feb 2014).
- Karchmer, M., Wigderson, A.: On span programs. In: Proceedings of Structures in Complexity Theory. pp. 102–111 (1993)
- Kitagawa, F., Nishimaki, R., Tanaka, K., Yamakawa, T.: Adaptively secure and succinct functional encryption: Improving security and efficiency, simultaneously. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 521–551. Springer, Heidelberg (Aug 2019).
- Kowalczyk, L., Wee, H.: Compact adaptively secure ABE for NC¹ from k-Lin. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 3–33. Springer, Heidelberg (May 2019).

- Lewko, A.B., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 62–91. Springer, Heidelberg (May / Jun 2010).
- Lewko, A.B., Waters, B.: New proof methods for attribute-based encryption: Achieving full security through selective techniques. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 180–198. Springer, Heidelberg (Aug 2012).
- Li, H., Lin, H., Luo, J.: ABE for circuits with constant-size secret keys and adaptive security. Cryptology ePrint Archive, Report 2022/659 (2022), https://eprint.iacr. org/2022/659
- Lin, H., Luo, J.: Succinct and adaptively secure ABE for ABP from k-lin. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part III. LNCS, vol. 12493, pp. 437–466. Springer, Heidelberg (Dec 2020).
- Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer, Heidelberg (Apr 2012).
- Mulmuley, K.: A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. Combinatorica 7(1), 101–104 (1987).
- Okamoto, T., Takashima, K.: Fully secure functional encryption with general relations from the decisional linear assumption. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 191–208. Springer, Heidelberg (Aug 2010).
- Quach, W., Wee, H., Wichs, D.: Laconic function evaluation and applications. In: Thorup, M. (ed.) 59th FOCS. pp. 859–870. IEEE Computer Society Press (Oct 2018).
- Sahai, A., Waters, B.R.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (May 2005).
- Takashima, K.: Expressive attribute-based encryption with constant-size ciphertexts from the decisional linear assumption. In: Abdalla, M., Prisco, R.D. (eds.) SCN 14. LNCS, vol. 8642, pp. 298–317. Springer, Heidelberg (Sep 2014).
- Tsabary, R.: Fully secure attribute-based encryption for t-CNF from LWE. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 62–85. Springer, Heidelberg (Aug 2019).
- 40. Waters, B.: A punctured programming approach to adaptively secure functional encryption. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 678–697. Springer, Heidelberg (Aug 2015).
- Wee, H.: Attribute-hiding predicate encryption in bilinear groups, revisited. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 206–233. Springer, Heidelberg (Nov 2017).
- Yamada, S., Attrapadung, N., Hanaoka, G., Kunihiro, N.: A framework and compact constructions for non-monotonic attribute-based encryption. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 275–292. Springer, Heidelberg (Mar 2014).
- Zhang, K., Gong, J., Tang, S., Chen, J., Li, X., Qian, H., Cao, Z.: Practical and efficient attribute-based encryption with constant-size ciphertexts in outsourced verifiable computation. In: Chen, X., Wang, X., Huang, X. (eds.) ASIACCS 16. pp. 269–279. ACM Press (May / Jun 2016)