FABEO: Fast Attribute-Based Encryption with Optimal Security

Doreen Riepel
Ruhr-Universität Bochum
Bochum, Germany
doreen.riepel@rub.de

Hoeteck Wee
NTT Research
Sunnyvale, CA, USA
hoeteck.wee@ntt-research.com

ABSTRACT

Attribute-based encryption (ABE) enables fine-grained access control on encrypted data and has a large number of practical applications. This paper presents FABEO: faster pairing-based ciphertext-policy and key-policy ABE schemes that support expressive policies and put no restriction on policy type or attributes, and the first to achieve optimal, adaptive security with multiple challenge ciphertexts. We implement our schemes and demonstrate that they perform better than the state-of-the-art (Bethencourt et al. S&P 2007, Agrawal et al., CCS 2017 and Ambrona et al., CCS 2017) on all parameters of practical interest.

CCS CONCEPTS

• **Security and privacy** → *Cryptography*.

KEYWORDS

attribute-based encryption; generic group model; tightness

ACM Reference Format:
Doreen Riepel and Hoeteck Wee. 2022. FABEO: Fast Attribute-Based Encryption with Optimal Security. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22), November 7–11, 2022, Los Angeles, CA, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3548606.3560699

1 INTRODUCTION

Attribute-based encryption (ABE) [30, 45] extends classical public-key encryption to support fine-grained access control on encrypted data. ABE has applications in a variety of settings including electronic medical records [5], messaging systems [41], online social networks [9] and information-centric networking [35]. Companies like Cloudflare already use ABE to distribute private key storage across data centers [48].

ABE comes in two variants: ciphertext-policy (CP-ABE) and key-policy (KP-ABE), depending on whether access policies are attached to ciphertexts or to keys [12, 30]. In CP-ABE, keys are associated with sets of attributes, and a key is able to recover the message hidden in a ciphertext if and only if the set of attributes satisfy the access policy attached to the ciphertext. For instance, a policy P could say '(Zipcode:90210 OR City:BeverlyHills) AND (AgeGroup:18-25)' and an individual A could have a key for Zipcode:90210, AgeGroup:Over65, in which case A would not be able to decrypt any message encrypted under P. A KP-ABE is the dual of CP-ABE with ciphertexts attached to attribute sets and keys associated with access policies.

There is by now a vast body of research on ABE realizing a broad spectrum of trade-offs between efficiency, expressiveness, security and hardness assumptions. The state of the art for practical ABE schemes are encapsulated by the following pairing-based schemes: (i) BSW CP-ABE scheme (Bethencourt, Sahai and Waters [12]), (ii) FAME CP-ABE and KP-ABE schemes (Agrawal and Chase [2]), and (iii) ABGW CP-ABE and KP-ABE schemes (Ambrona, Barthe, Gay and Wee [6]). These schemes simultaneously achieve the following properties that are highly desirable in practice:

(1) support expressive policies described by boolean formula and monotone span programs (MSP);
(2) put no restriction on size of policies or attribute sets;
(3) allow any arbitrary string such as street addresses to be used as an attribute;
(4) achieve the strong and natural notion of adaptive security, the defacto standard for ABE.

However, these schemes achieve incomparable efficiency guarantees, and deciding which one to deploy requires making complex performance trade-offs that depend on the policies that arise in the specific context.

1.1 Our Contributions

We present FABEO, new pairing-based KP-ABE and CP-ABE schemes achieving properties (1) – (4), with improved efficiency and quantitatively stronger security guarantees. FABEO uses asymmetric (Type-III) prime-order bilinear groups ($\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$) which support efficient hashing to $\mathbb{G}_1$ [24, 46]. Ciphertexts and secret keys in FABEO comprise mostly of elements in the smaller and faster group $\mathbb{G}_1$, plus 1 or 2 elements in $\mathbb{G}_2$. Computation for key generation, encryption and decryption are mostly carried out in $\mathbb{G}_1$, with 2 to 3 pairings for decryption. We prove *optimal* security bounds for FABEO against adversaries that get an arbitrary number of ciphertexts and keys: in particular, when instantiated over the popular BLS12-381 curve, FABEO achieves close to 128-bit security.

FABEO subsumes BSW, FAME and ABGW on all parameters of practical interest. We improve upon the ciphertext and key sizes of all three schemes, as well as the running times. In particular, our ciphertexts are 66% smaller; encryption is (at least) 33% faster; and decryption uses fewer pairings. FABEO also supports multi-use of attributes like in BSW and ABGW (without an a-prior bound during set-up), with a small additive overhead in the multi-use parameter. See Table 1 for a property-wise comparison of our schemes against BSW, FAME, ABGW and other prominent schemes in the literature, as well as Tables 2 and 3 for a theoretical analysis and comparison for efficiency.

CCS '22, November 7–11, 2022, Los Angeles, CA, USA

ACM ISBN 978-1-4503-9450-5/22/11.
https://doi.org/10.1145/3548606.3560699

Scheme	*Unrestricted policies*	*Arbitrary attributes*	*Fast decryption*	*Attribute multi-use*	*Security bounds*
CP-ABE					
BSW [12, §4.2] [2, §D]	✓	✓	×	✓	t^3/p
Waters [51, §3] [2, §E]	✓	×	×	✓	–
ABGW [6, §5.3]	✓	✓	×	✓	t^4/p
FAME [2, §3]	✓	✓	✓	×	t^4/p
Ours FABEO [Fig 1]	✓	✓	✓	✓	t^2/p
KP-ABE					
GPSW [30, §A.1] [2, §F]	✓	×	×	✓	–
ABGW [6, §5.3]	✓	✓	×	✓	t^4/p
FAME [2, §B]	✓	✓	✓	×	t^4/p
Ours FABEO [Fig 1]	✓	✓	✓	✓	t^2/p

Table 1: A property-wise comparison of the various ABE schemes we consider. The BSW, Waters and GPSW schemes were specified using symmetric pairings in the original works; throughout, we refer to the asymmetric variants from [2]. The last column shows the security bounds on the adversary's advantage in the multi-ciphertext setting for the adaptively secure schemes, with dashes indicating selectively secure schemes.

FABEO achieves properties (2) and (3) by hashing attributes to $\mathbb{G}_1$; smaller ciphertext/key sizes and fast decryption via randomness reuse (in CP-ABE ciphertexts and KP-ABE keys); and adaptive security without efficiency penalties by considering "generic" adversaries, a widely accepted model that captures all known attacks. While each of these techniques is already present in BSW, FAME, ABGW and prior works, FABEO is the first to combine them in a single design, along with a novel analysis establishing optimal security.

Optimal security. We prove security of our schemes in the generic bilinear group model (GGM) [14, 43, 47] (as with BSW and ABGW), where we model the underlying hash function as a random oracle [11] (as with BSW and FAME). We show that any generic, adaptive adversary running in time t and sees at most t ciphertexts and keys breaks our schemes with probability at most $O(t^2/p)$, where p is the order of the underlying group. This bound is optimal, since an adversary can break discrete log with the same probability. Prior ABE schemes, including BSW, FAME, and ABGW, achieve a bound of $O(t^3/p)$ or worse, since the security proofs only consider a single challenge ciphertext, and a hybrid argument is needed to achieve multi-ciphertext security.

Proof framework and application. In both our CP-ABE and KP-ABE schemes, the ciphertext ct for x and secret key sk for y are of the form:

$$\mathsf{ct} = \left(g_1^{c_x^1(\mathbf{s},\mathbf{b})}, g_2^{c_x^2(\mathbf{s})}, e(g_1,g_2)^{\alpha s_1}\cdot M\right), \mathsf{sk} = \left(g_1^{k_y^1(\alpha,\mathbf{b},\mathbf{r})}, g_2^{k_y^2(\mathbf{r})}\right).$$

Here, $\mathbf{s} = (s_1, \ldots)$ and $\mathbf{r}$ are fresh randomness; $g_1^{\mathbf{b}}$ contains the hash of every attribute in the universe[1], and $c_x^1, c_x^2, k_y^1, k_y^2$ are simple functions of degree 1 or 2. Roughly speaking, we show that for schemes of this form[2], security for a single ciphertext-key pair implies optimal, adaptive security against generic adversaries with an arbitrary number of ciphertexts and keys. Our modular proof framework extends and generalizes an analogous statement shown in ABGW in several ways: (i) we allow $c_x^2(\mathbf{s})$ to have arbitrary length instead of length 1, as is necessary to capture our CP-ABE scheme and the one below, (ii) we consider security with multiple ciphertexts, and (iii) we achieve optimal security.

Next, we describe an additional application of our proof framework that pertains to property (1). A limitation of boolean formula and monotone span programs is they do not capture computation over data of arbitrary, unbounded size, which arise settings such as genome sequencing, processing network and event logs, tax returns and virus scanners; such computation are better captured by regular languages, or deterministic finite automata (DFA). As a secondary contribution, we prove that Waters' KP-ABE scheme for DFA [52] achieves optimal, adaptive security.[3] In this scheme, $c_x^2(\mathbf{s})$ has arbitrary length that grows with x. Compared to prior adaptively secure KP-ABE for DFA [3, 7, 8, 29, 42], we obtain (at least) a 50% improvement in ciphertext and key sizes as well as running times.

Implementation and evaluation. We implement FABEO in the Charm framework [4]. Our experiments validate our theoretical analysis in Table 2 showing that FABEO improves on the performance of BSW, FAME and ABGW, for all of key generation, encryption and decryption. FABEO compares favorably even against the Waters CP-ABE [51] and GPSW KP-ABE [30], even though these schemes do not achieve property (3). See Figure 3 in Section 7 for the performance of the algorithms of each scheme under various test cases. Our code is available on GitHub [44] (we plan to also release our code to open source).

All computations are performed on an ordinary laptop and we achieve practical results, even for large attribute sets and policies. Specifically for our CP-ABE with the MNT224 curve, set-up takes less than 0.02s, and it takes around 0.09s to generate a key for 100 attributes, and 0.18s to encrypt data under a policy that requires all 100 attributes. Decryption then takes only 0.02s. As a comparison,

[1] Ignoring for now the fact that $\mathbf{b}$ has exponential length.

[2] The ABGW CP-ABE and KP-ABE schemes we compare with are not of this form since the k_y^1 computes a rational function.

[3] Waters only proved weaker, selective security for his scheme. More precisely, we consider a variant of Waters' scheme with smaller keys from [28].

the ABGW CP-ABE scheme takes 0.63s to generate a key for the same number of attributes, 0.33s to encrypt and 0.48s to decrypt. In FAME, decryption takes 0.03s, and key generation and encryption are slower than ABGW.

Summary of Contributions. To summarize, our contributions are as follows:

- We present new KP-ABE and CP-ABE schemes for MSP with improved efficiency guarantees and the first to achieve optimal, adaptive security with multiple challenge ciphertexts.
- We provide a more general and modular framework for proving optimal ABE security in the GGM.
- We implement our KP-ABE and CP-ABE schemes for MSP and evaluate their performance for various parameters.
- We present and implement a new KP-ABE for DFA with optimal, adaptive security in the GGM.

1.2 Discussion and Related Work

We discuss additional context and related works.

Choosing curve parameters. When choosing curve parameters for a pairing-based scheme, practitioners often base the decisions on the hardness of the discrete log problem, and ignore the security bounds provided in security proof for the scheme. This is in part due to the limited number of pairing-friendly curves that are available in practice [46], and the possibly prohibitive performance penalty from using a curve with larger bit security. In particular, there is an implicit expectation that a scheme instantiated over a curve with 128-bit security should also achieve close to 128-bit security. Our work takes a step towards rigorously justifying this expectation in the context of pairing-based ABE.

Achieving adaptive security. There are two main approaches for realizing adaptive security for ABE schemes in the literature: (1) prove security against generic adversaries as was done in BSW, ABGW and this work, and (2) adopt the dual system encryption framework [7, 40, 50, 53] as used in FAME, which allows us to base security on SXDH and DLIN (and in some settings, with the additional use of q-type assumptions). While the latter yields theoretically stronger results, it incurs a huge penalty in efficiency: for security from k-LIN ($k = 1$ corresponds to SXDH and $k = 2$ to DLIN), it requires (at least) a factor $k + 1$ blow-up in ciphertext and key sizes as well as running times for encryption, key generation and decryption [3, 8, 16]. Moreover, the schemes have a more complex structure, and the security proofs are also substantially more complex. Another drawback is that the proofs typically require a hybrid argument over the keys and the ciphertexts, so we cannot hope for a security bound better than $O(t^4/p)$.

GGM security. We argue that GGM security is sufficient for most practical applications. The reasoning is two-fold: First, our understanding of pairing curves has advanced substantially over the past two decades, with increasing adoption (e.g. Cloudflare and ZCash) as well as on-going standardization [46]. The known attacks fall broadly into two categories: (1) attacks on discrete log, most notably the exTNFS in [36], rendering the curves unsuitable for any applications, (2) attacks that are captured by the GGM [19]. In short, there is in practice no discernible distinction between the standard assumptions like SXDH and GGM security. Second, it is much easier to break a real-world system via side channel attacks or poor security practices (e.g. phishing attacks or weak passwords) than to come up with an attack outside of the GGM. Indeed, a large number of recent works also use the GGM to analyze practical cryptosystems, e.g. [10, 31, 32].

Optimal and tight security. This work falls under a broader cryptographic research agenda of achieving optimal security and tight security reductions, for instance, recent works on symmetric-key encryption [33], signature schemes [22] and TLS 1.3 [23, 26]. In the context of ABE, optimal security was only previously known in very limited settings, namely identity-based encryption and its hierachical variant [13, 15, 17, 18, 25, 27, 34, 38]; we clarify that these works focus on the more challenging goal of basing security on static assumptions such as DLIN. In particular, these are the only settings where we know how to carry out a dual system encryption proof with security bound better than $O(t^4/p)$.

Benchmarking ABE schemes. Two very recent works [20, 21] looked into benchmarking pairing-based ABE schemes, focusing on low-level optimizations (whereas our work focuses on high-level design as well as new security guarantees and proof techniques): the first for CP-ABE covering BSW, Waters, and FAME but not ABGW, and the second for broadcast encryption. Both works highlight the complexity of effective benchmarking due to incomparable trade-offs between efficiency, expressiveness, security and hardness assumptions, which we alluded to at the beginning of the paper. Our results, together with those in ABGW and the preceding discussion, support the thesis that one should consider GGM-based schemes for benchmarking, since we can achieve the strongest notion of adaptive security without efficiency penalties.

1.3 Technical Overview

Let $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ be an (asymmetric) bilinear group of prime order p, along with a pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ and generators g_1, g_2 for $\mathbb{G}_1, \mathbb{G}_2$ respectively. In general, the bit sizes of group elements in $\mathbb{G}_2$ are 2-3 times that of $\mathbb{G}_1$ and group operations in $\mathbb{G}_2$ take (at least) twice as much time. In addition, we can securely hash into $\mathbb{G}_1$ at the cost of roughly one exponentiation in $\mathbb{G}_1$.

High-level design. We begin with a high-level overview of our KP-ABE scheme described in Figure 1. An MSP is given by a matrix $\mathbf{M}$ and a function π that maps each row of $\mathbf{M}$ to an attribute (for this overview, assume π is injective, i.e., no attribute multi-use). Following [30], we design the ciphertexts and secret keys so that for each row i in $\mathbf{M}$ such that $\pi(i)$ appears in the attribute set, decryption will compute

$$e(g_1, g_2)^{s_1 \alpha_i} \tag{1}$$

where α_i is a share of the master secret key α and $s_1 \leftarrow \mathbb{Z}_p$ is the encryption randomness. The values in (1) can then be combined to recover the blinding factor $e(g_1, g_2)^{s_1 \alpha}$.

To realize the above invariant, we have

$$\left(g_2^{s_1}, \mathsf{H}(\pi(i))^{s_1}\right) \in \mathsf{ct}, \quad \left(g_2^{r}, g_1^{\alpha_i}\mathsf{H}(\pi(i))^{r}\right) \in \mathsf{sk}$$

so that we can compute (1) using $e(g_1^{\alpha_i}\mathsf{H}(\pi(i))^r, g_2^{s_1})/e(\mathsf{H}(\pi(i))^{s_1}, g_2^r)$. In addition, we use the same r across all the rows in $\mathbf{M}$, to keep the

Setup(1^λ). Run GroupGen(1^λ) to obtain $\mathcal{G} := (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$. Pick $\alpha \xleftarrow{\$} \mathbb{Z}_p$ and a hash function $\mathsf{H} : [|\mathcal{U}|+1] \to \mathbb{G}_1$. Compute the master public key as

$$\mathsf{mpk} := (\mathcal{G}, \mathsf{H}, e(g_1, g_2)^\alpha)$$

Let msk $:= \alpha$ be the master secret key.

KeyGen(msk, $\mathcal{S} \subseteq \mathcal{U}$). Pick $r \xleftarrow{\$} \mathbb{Z}_p$. Compute

$$\mathsf{sk}_1 := g_1^\alpha \cdot \mathsf{H}(|\mathcal{U}|+1)^r \qquad \mathsf{sk}_{2,u} := \mathsf{H}(u)^r \qquad \mathsf{sk}_3 := g_2^r$$

for each $u \in \mathcal{S}$. Output $\mathsf{sk} := (\mathsf{sk}_1, \{\mathsf{sk}_{2,u}\}_{u \in \mathcal{S}}, \mathsf{sk}_3)$.

Enc(mpk, $(\mathbf{M}, \pi)$). Pick $s_1 \xleftarrow{\$} \mathbb{Z}_p$, $\mathbf{v} \xleftarrow{\$} \mathbb{Z}_p^{n_2-1}$, $\mathbf{s}' \xleftarrow{\$} \mathbb{Z}_p^\tau$. Compute

$$\mathsf{ct}_1 := g_2^{s_1} \qquad \mathsf{ct}_{2,j} := g_2^{\mathbf{s}'[j]}$$

for $j \in [\tau]$, as well as

$$\mathsf{ct}_{3,i} := \mathsf{H}(|\mathcal{U}|+1)^{\mathbf{M}_i (s_1 \| \mathbf{v})^\top} \cdot \mathsf{H}(\pi(i))^{\mathbf{s}'[\rho(i)]}$$

for each row $i \in [n_1]$. Output $\mathsf{ct} := (\mathsf{ct}_1, (\mathsf{ct}_{2,j})_{j \in [\tau]}, (\mathsf{ct}_{3,i})_{i \in [n_1]})$ and $d := e(g_1, g_2)^{\alpha s_1}$.

Dec(mpk, $(\mathbf{M}, \pi)$, $\mathcal{S}$, ct, sk). If $\mathcal{S}$ satisfies $(\mathbf{M}, \pi)$, there exist constants $\{\gamma_i\}_{i \in I}$ s.t. $\sum_{i \in I} \gamma_i \mathbf{M}_i = (1, 0, \ldots, 0)$. Reconstruct d by computing

$$e(\mathsf{sk}_1, \mathsf{ct}_1) \cdot \frac{\prod_{j \in [\tau]} e(\prod_{i \in I, \rho(i)=j} (\mathsf{sk}_{2,\pi(i)})^{\gamma_i}, \mathsf{ct}_{2,j})}{e(\prod_{i \in I} (\mathsf{ct}_{3,i})^{\gamma_i}, \mathsf{sk}_3)}$$

and output the result.

Setup(1^λ). Run GroupGen(1^λ) to obtain $\mathcal{G} := (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$. Pick $\alpha \xleftarrow{\$} \mathbb{Z}_p$ and a hash function $\mathsf{H} : \mathcal{U} \to \mathbb{G}_1$. Compute the master public key as

$$\mathsf{mpk} := (\mathcal{G}, \mathsf{H}, e(g_1, g_2)^\alpha)$$

Let msk $:= \alpha$ be the master secret key.

KeyGen(msk, $(\mathbf{M}, \pi)$). Pick $\mathbf{r}' \xleftarrow{\$} \mathbb{Z}_p^\tau$, $\mathbf{v} \xleftarrow{\$} \mathbb{Z}_p^{n_2-1}$. Compute

$$\mathsf{sk}_{1,j} := g_2^{\mathbf{r}'[j]}$$

for $j \in [\tau]$, as well as

$$\mathsf{sk}_{2,i} := g_1^{\mathbf{M}_i (\alpha \| \mathbf{v})^\top} \cdot \mathsf{H}(\pi(i))^{\mathbf{r}'[\rho(i)]}$$

for each row $i \in [n_1]$. Output $\mathsf{sk} := ((\mathsf{sk}_{1,j})_{j \in [\tau]}, (\mathsf{sk}_{2,i})_{i \in [n_1]})$.

Enc(mpk, $\mathcal{S} \subseteq \mathcal{U}$). Pick $s \xleftarrow{\$} \mathbb{Z}_p$. For each $u \in \mathcal{S}$ compute

$$\mathsf{ct}_{1,u} := \mathsf{H}(u)^s \qquad \mathsf{ct}_2 := g_2^s$$

Output $\mathsf{ct} := ((\mathsf{ct}_{1,u})_{u \in \mathcal{S}}, \mathsf{ct}_2)$ and $d := e(g_1, g_2)^{\alpha s}$.

Dec(mpk, $\mathcal{S}$, $(\mathbf{M}, \pi)$, ct, sk). If $\mathcal{S}$ satisfies $(\mathbf{M}, \pi)$, there exist constants $\{\gamma_i\}_{i \in I}$ s.t. $\sum_{i \in I} \gamma_i \mathbf{M}_i = (1, 0, \ldots, 0)$. Reconstruct d by computing

$$\frac{e(\prod_{i \in I} (\mathsf{sk}_{2,i})^{\gamma_i}, \mathsf{ct}_2)}{\prod_{j \in [\tau]} e(\prod_{i \in I, \rho(i)=j} (\mathsf{ct}_{1,\pi(i)})^{\gamma_i}, \mathsf{sk}_{1,j})}$$

and output the result.

Figure 1: Our CP-ABE (left) and KP-ABE (right) scheme for monotone span programs ($\mathbf{M} \in \mathbb{Z}_p^{n_1 \times n_2}, \pi : [n_1] \to \mathcal{U}$). We define $\rho(i) := |\{z \mid \pi(z) = \pi(i), z \leq i\}|$ and $\tau = \max_{i \in [n_1]} \rho(i)$ corresponding to maximum number of times an attribute is used in M.

key size small. This way, we can also carry out decryption using two pairings[4]. and with most of the computation in the faster group $\mathbb{G}_1$. In contrast,

- BSW uses a different r_i for each share, namely $(g_2^{r_i}, g_1^{\alpha_i} \mathsf{H}(\pi(i))^{r_i}) \in \mathsf{sk}$ (here, we are describing the KP-ABE analogue of the BSW CP-ABE). This incurs a factor 2 blow-up in key size, and decryption requires computing a pairing for each row of $\mathbf{M}$.
- ABGW uses
$$\left(g_1^{s-s_i}, g_1^{s_i(b_1+\pi(i)b_2)}\right) \in \mathsf{ct}, \quad \left(g_2^{\frac{\alpha_i}{b_1+\pi(i)b_2}}, g_2^{\alpha_i}\right) \in \mathsf{sk}$$
where $g_1^{b_1}, g_1^{b_2}$ comes from mpk. This incurs (at least) a factor 2 blow-up in ciphertext and key sizes, and an extra exponentation per attribute during encryption. Decryption requires computing a pairing for each attribute.
- FAME replaces $g_2^{s_1}, g_2^r$ with DLIN-tuples in $\mathbb{G}_2^3$ in order to achieve security under the DLIN assumption using the dual system encryption framework as described in Section 1.2. Overall, this incurs a factor 3 blow-up in ciphertext and key sizes, as well as a factor 3-6 blow-up in running time for encryption and key generation.

Our CP-ABE scheme is conceptually the dual of our KP-ABE, though algebraically more intricate and less intuitive (the same holds for BSW, FAME, and ABGW). Briefly, instead of (1), decryption computes $e(g_1, g_2)^{\mu_i b' r}$ where μ_i is a share of the encryption randomness s_1; $g_1^{b'}$ is specified in the public key; and r comes from key generation randomness. These values can then be combined to compute $e(g_1, g_2)^{s_1 b' r}$, which is in turn used to recover the blinding factor $e(g_1, g_2)^{s_1 \alpha}$. Our CP-ABE scheme is the same as the AC17-LU-OK and AC17-LU-CP schemes in the independent work [21], which asserts selective security under q-type assumptions without a formal security proof.

Proof strategy. We provide a unified proof security of our KP-ABE and CP-ABE schemes in the GGM, where we model H as a random oracle. At a high level, we follow the framework in [6]. Both our KP-ABE and CP-ABE schemes have the following structure where the ciphertext is associated with a label x and the secret key with a label y (x is an attribute set for KP-ABE and a policy for CP-ABE, and vice-versa for y)

$$\mathsf{ct}_x = \left(g_1^{c_x^1(\mathbf{s} \otimes \mathbf{b})}, g_2^{c_x^2(\mathbf{s})}, e(g_1, g_2)^{\alpha s_1} \cdot M\right), \mathsf{sk}_y = \left(g_1^{k_y^1(\alpha, \mathbf{r}, \mathbf{b} \otimes \mathbf{r})}, g_2^{k_y^2(\mathbf{r})}\right).$$

where[5]

- $\mathbf{s} = (s_1, \ldots)$ and $\mathbf{r}$ are random vectors over $\mathbb{Z}_p$ corresponding to randomnness for encryption and key generation;
- $g_1^{\mathbf{b}}$ contains the hash of every attribute in the universe, along with $g_1^{b'}$ for our CP-ABE (note that the length of $\mathbf{b}$ is exponential, but the c_x^1, k_y^1 only depend on a polynomial number of entries of $\mathbf{b}$);
- $c_x^1, c_x^2, k_y^1, k_y^2$ are linear functions over $\mathbb{Z}_p$ (therefore c_x^1 and k_y^1 computes degree 2 functions of $\mathbf{s}, \mathbf{b}, \mathbf{r}, \alpha$);
- decryption uses the pairing to compute $e(g_1, g_2)^{c_x^1(\mathbf{s} \otimes \mathbf{b}) \otimes k_y^2(\mathbf{r})}$ and $e(g_1, g_2)^{k_y^1(\alpha, \mathbf{r}, \mathbf{b} \otimes \mathbf{r}) \otimes c_x^2(\mathbf{s})}$, followed by additional linear

[4] by writing $\prod_i (e(g_1, g_2)^{s_1 \alpha_i})^{\gamma_i}$ as $e(\prod_i (g_1^{\alpha_i} \mathsf{H}(\pi(i))^r)^{\gamma_i}, g_2^{s_1}) \cdot e(\prod_i (\mathsf{H}(\pi(i))^{s_1})^{\gamma_i}, g_2^r)^{-1}$

[5] The tensor product $\mathbf{u} \otimes \mathbf{v}$ of two vectors $\mathbf{u} = (u_1, u_2, \ldots)$ and $\mathbf{v} = (v_1, v_2, \ldots)$ is a vector $(u_1 v_1, u_1 v_2, \ldots)$ containing all pairwise products of the entries in u and v.

computation in the exponent to recover the blinding factor $e(g_1, g_2)^{\alpha s_1}$.

We refer to ABE schemes with the above structure as a PES-ABE (PES is short for pair encoding schemes [7]). Towards proving GGM security, we consider notions of symbolic security for PES-ABE, where an adversary sees abstract expressions for group elements in the form of polynomials. The proof of security of our KP-ABE and CP-ABE schemes follows the following modular framework:

Step 1. We show that our KP-ABE and CP-ABE schemes satisfy the syntax of a PES-ABE and (1,1) symbolic security, a relaxation of ABE security where the adversary is selective[6] and only receives a single ciphertext and single secret key.

Step 2. We prove that any PES-ABE satisfying (1,1) symbolic security also satisfies strong symbolic security, where the adversary is still selective but can see the public key as well as an arbitrary number of ciphertexts and secret keys.

Step 3. We prove that any PES-ABE satisfying strong symbolic security is adaptively secure in the GGM with optimal security.

We now describe the key differences between our framework and the one in ABGW:

- The syntax for PES-ABE is different: (i) both ct_x and sk_y contain elements from both $\mathbb{G}_1$ and $\mathbb{G}_2$, and (ii) we generate $g_1^{\mathbf{b}}$ using a random oracle.
- We introduce a strengthening of (1,1) symbolic security where we essentially require that all of $[\alpha c_x^1(\mathbf{s})]_T$ are pseudorandom, and not just $[\alpha s_1]_T$. Our notion is also weaker in that the proof only needs to reason about the terms $e(g_1, g_2)^{c_x^1(\mathbf{s}\otimes\mathbf{b})\otimes k_y^2(\mathbf{r})}$ and $e(g_1, g_2)^{k_y^1(\alpha,\mathbf{r},\mathbf{b}\otimes\mathbf{r})\otimes c_x^2(\mathbf{s})}$.
- The ABGW KP-ABE and CP-ABE schemes for MSP do not satisfy the syntax of a PES-ABE since k_y^2 computes rational functions with linear functions $\mathbf{b}$ in the denominator and therefore proving security of these schemes require directly establishing strong symbolic security;
- The analogue of strong symbolic securty in ABGW in Steps 2 and 3 considers only a single challenge ciphertext.
- We achieve a security bound of $O(t^2/p)$ in Step 3, whereas ABGW achieves $O(t^3/p)$. Our proof crucially relies on the fact that $c_x^1, c_x^2, k_y^1, k_y^2$ compute functions of degree at most 2 in the inputs so that we only need to apply Schwartz-Zippel to constant-degree polynomials. The proof in ABGW applies Schwartz-Zippel to polynomials of degree t in order to "clear the denominators" across t keys.

2 PRELIMINARIES

We will first fix some notation that we will use throughout the paper. For integers m, n where $m < n$, $[m, n]$ denotes the set $m, m+1, ..., n$. For $m = 1$, we simply write $[n]$. For a prime p, let $\mathbb{Z}_p$ denote the set $[0, p-1]$, where addition and multiplication are computed modulo p. For a set $\mathcal{S}$, $s \xleftarrow{\$} \mathcal{S}$ denotes that s is sampled uniformly and independently at random from $\mathcal{S}$. $y \leftarrow \mathcal{A}(x_1, x_2, \ldots)$ denotes that on input $x_1, x_2, \ldots$ the probabilistic algorithm $\mathcal{A}$ returns y. $\mathcal{A}^{\mathcal{O}}$ denotes that algorithm $\mathcal{A}$ has access to oracle $\mathcal{O}$. An adversary is a probabilistic algorithm. A probabilistic algorithm is called *efficient* or PPT if its running time is bounded by some polynomial in the length of its input.

We use lower case bold-face letters for row vectors, where $\|$ denotes concatenation of row vectors. $\mathbf{v}[i]$ denotes the i-th coordinate of the vector $\mathbf{v}$. Given a vector $\mathbf{v}$ of polynomials of length m over $\mathbb{Z}_p$, we write $\mathrm{span}(\mathbf{v})$ to denote $\{\mathbf{v} \cdot \mathbf{e}^\top : \mathbf{e} \in \mathbb{Z}_p^m\}$. Formal variables are marked with a tilde. We write $\tilde{\mathbf{v}} \leftarrow \mathsf{Var}^n$ to pick n formal variables.

2.1 Pairing Groups

Let GroupGen be a PPT algorithm that takes a security parameter 1^λ as input and returns a group description $\mathcal{G} := (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$, where p is a prime of $\Theta(\lambda)$ bits, $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ are cyclic groups of order p, $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is a non-degenerate bilinear map (also called pairing) and g_1 resp. g_2 or generators of $\mathbb{G}_1$ resp. $\mathbb{G}_2$. The generator g_T of $\mathbb{G}_T$ can be computed as $e(g_1, g_2)$. We require that the group operations in $\mathbb{G}_1$, $\mathbb{G}_2$, $\mathbb{G}_T$ and the bilinear map e are computable in deterministic polynomial time in λ. In this work, we only consider asymmetric (or Type-III) pairing groups where there exists no efficiently computable homomorphism between $\mathbb{G}_1$ and $\mathbb{G}_2$. In some cases we will use *implicit representation* of group elements: for a vector $\mathbf{v}$ over $\mathbb{Z}_p$, we define $[\mathbf{v}]_1 := g_s^{\mathbf{v}}$ for $s \in \{1, 2, T\}$, where exponentiation is carried out component-wise.

2.2 Attribute-based Encryption

Throughout the paper, we will use a KEM-style definition of ABE. However note that it is implied by the corresponding definition in the PKE setting.

Syntax. An attribute-based encryption (ABE) scheme for some class P consists of four algorithms:

$\mathsf{Setup}(1^\lambda, \mathsf{P}) \to (\mathsf{mpk}, \mathsf{msk})$. The setup algorithm gets as input the security parameter 1^λ and class description P. It outputs the master public key mpk and the master secret key msk. We assume mpk defines the key space $\mathcal{K}$.

$\mathsf{Enc}(\mathsf{mpk}, x) \to (\mathsf{ct}_x, d)$. The encryption algorithm gets as input mpk and an input x. It outputs a ciphertext ct_x and an encapsulated key $d \in \mathcal{K}$.

$\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, y) \to \mathsf{sk}_y$. The key generation algorithm gets as input mpk, msk and $y \in \mathsf{P}$. It outputs a secret key sk_y.

$\mathsf{Dec}(\mathsf{mpk}, x, y, \mathsf{ct}_x, \mathsf{sk}_y) \to m$. The decryption algorithm gets as input sk_y and ct_x such that $\mathsf{P}(x, y) = 1$ along with mpk. It outputs a key d.

Correctness. For all input x and y with $\mathsf{P}(x) = 1$, we require

$$\Pr\left[\mathsf{Dec}(\mathsf{mpk}, x, y, \mathsf{ct}_x, \mathsf{sk}_y) = d : \begin{array}{l} (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathsf{P}) \\ \mathsf{sk}_y \leftarrow \mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, y) \\ (\mathsf{ct}_x, d) \leftarrow \mathsf{Enc}(\mathsf{mpk}, x) \end{array}\right] = 1.$$

Many-Ciphertext CPA Security. We define security by a game between a challenger and an adversary $\mathcal{A}$. The challenger picks a random challenge bit β and provides the following oracles to $\mathcal{A}$.

- Setup oracle $\mathcal{O}_{\mathsf{mpk}}$: This oracle can only be queried once and it must be the first query. The challenger runs Setup to obtain $(\mathsf{msk}, \mathsf{mpk})$ and outputs mpk to $\mathcal{A}$.

[6] In this overview, we use selective to refer to an adversary that specifies all of its ciphertext and key queries in advance.

- Ciphertext (or challenge) oracle $\mathcal{O}_{\mathsf{ct}}$: On the i-th query, $\mathcal{A}$ provides $x_i \in \mathcal{X}$. The challenger runs $(\mathsf{ct}_i, d_i^{(0)}) \leftarrow \mathsf{Enc}(\mathsf{mpk}, x_i)$, chooses a random key $d_i^{(1)} \overset{\$}{\leftarrow} \mathcal{K}$ and outputs $(\mathsf{ct}_i, d_i^{(\beta)})$.
- Secret key oracle $\mathcal{O}_{\mathsf{sk}}$: On the j-th query, $\mathcal{A}$ provides $y_j \in \mathcal{Y}$. The challenger runs $\mathsf{sk}_j \leftarrow \mathsf{KeyGen}(\mathsf{msk}, y_j)$ and outputs sk_j.

$\mathcal{O}_{\mathsf{ct}}$ and $\mathcal{O}_{\mathsf{sk}}$ can be queried adaptively and an arbitrary polynomial number of times. Finally, $\mathcal{A}$ outputs a bit β'. We say that $\mathcal{A}$ wins the game if $\beta = \beta'$ and $\mathsf{P}(x_i, y_j) = 0$ for all queries x_i and y_j.

Definition 2.1. An ABE scheme is adaptively many-ciphertext secure if for all efficient $\mathcal{A}$,

$$\mathsf{Adv}_{\mathsf{ABE},\mathcal{A}}(\lambda) := \left|\Pr[\beta = \beta'] - \tfrac{1}{2}\right|$$

is negligible in λ.

Boolean formulae and MSP. Boolean formulae are a common way to model access control. A (monotone) boolean formula consists of **and** and **or** gates, where each input is associated with an attribute in the universe of attributes denoted by $\mathcal{U}$. Monotone means that an authorized user who acquires more attributes will not lose any privileges. Let $\mathcal{S} \subseteq \mathcal{U}$ be a set of attributes. We say that $\mathcal{S}$ satisfies a boolean formula if we set all inputs of the formula that map to an attribute in $\mathcal{S}$ to true and the others to false and the formula evaluates to true.

Monotone span programs (MSP) are a more general class of functions and include boolean formulae. We encode an access structure by a policy $(\mathbf{M}, \pi)$, where $\mathbf{M} \in \mathbb{Z}_p^{n_1 \times n_2}$ and $\pi : [n_1] \rightarrow \mathcal{U}$. Note that we can compute $(\mathbf{M}, \pi)$ for any (monotone) boolean formula in polynomial time [39]. Then every row $\mathbf{M}_i$ corresponds to an input to the formula and the number of columns is the same as the number of **and** gates. If the mapping π is not injective, we use the notation $\rho(i) := |\{z \mid \pi(z) = \pi(i), z \leq i\}|$ to denote the $\rho(i)$-th occurrence of attribute $\pi(i)$.

Let $\mathcal{S} \subseteq \mathcal{U}$ be a set of attributes and $I = \{i \mid i \in [n_1], \pi(i) \in \mathcal{S}\}$ be the indices of rows in $\mathbf{M}$ that are associated with $\mathcal{S}$. We say that $(\mathbf{M}, \pi)$ accepts $\mathcal{S}$ if the vector $(1, 0, \ldots, 0)$ lies in the span of rows associated with $\mathcal{S}$. This means, there exist constants $\gamma_i \in \mathbb{Z}_p$ for $i \in I$ such that $\sum_{i \in I} \gamma_i \mathbf{M}_i = (1, 0, \ldots, 0)$. These constants can be computed in time polynomial in the size of $\mathbf{M}$. On the contrary, $(\mathbf{M}, \pi)$ does not accept $\mathcal{S}$ if there exist a vector $\mathbf{w} \in \mathbb{Z}_p^{n_2}$ such that $\mathbf{w}$ is orthogonal to all rows $\mathbf{M}_i$ for $\pi(i) \in \mathcal{S}$, but not to $(1, 0, \ldots, 0)$. That means $\langle \mathbf{w}, \mathbf{M}_i \rangle = 0$. W.l.o.g. we can set $\mathbf{w}[1] = 1$.

Polynomials. Let p be a prime and $n \in \mathbb{N}$. We denote the set of multi-variate polynimals over $\mathbb{Z}_p$ with indeterminates $\tilde{x}_1, \ldots, \tilde{x}_n$ by $\mathbb{Z}_p[\tilde{x}_1, \ldots, \tilde{x}_n]$.

3 PES-ABE

We consider PES-ABE, which is a standard ABE scheme augmented with 3 deterministic algorithms Setup_0, Enc_0, KeyGen_0 used in Setup, Enc, KeyGen, Dec respectively, where:

- $\mathsf{Setup}_0(1^\lambda, \mathcal{X}, \mathcal{Y})$ outputs $n \in \mathbb{N}$,
- $\mathsf{Enc}_0(x)$ outputs linear functions $c^1 : \mathbb{Z}_p^{wn} \rightarrow \mathbb{Z}_p^{w_1}$, $c^2 : \mathbb{Z}_p^{w} \rightarrow \mathbb{Z}_p^{w_2}$,
- $\mathsf{KeyGen}_0(y)$ outputs linear functions $k^1 : \mathbb{Z}_p^{1+m+mn} \rightarrow \mathbb{Z}_p^{m_1}$, $k^2 : \mathbb{Z}_p^{m} \rightarrow \mathbb{Z}_p^{m_2}$,

and

- $\mathsf{Setup}(1^\lambda)$: Run $\mathcal{G} := (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2) \leftarrow \mathsf{GroupGen}(1^\lambda)$, $n \leftarrow \mathsf{Setup}_0$. Pick $\alpha \overset{\$}{\leftarrow} \mathbb{Z}_p$ and a hash function $\mathsf{H} : [n] \rightarrow \mathbb{G}_1$ Output
$$\mathsf{mpk} := (\mathcal{G}, \mathsf{H}, [\alpha]_T),\ \mathsf{msk} := \alpha$$
Using H, we implictly define $\mathbf{b} \in \mathbb{Z}_p^n$ via $[\mathbf{b}[i]]_1 = \mathsf{H}(i)$.
- Enc: Run $(c^1, c^2) \leftarrow \mathsf{Enc}_0(x)$. Pick $\mathbf{s} \leftarrow \mathbb{Z}_p^w$. Compute $[\mathbf{c}^1]_1 := c^1([\mathbf{s} \otimes \mathbf{b}]_1)$, $[\mathbf{c}^2]_2 := c^2([\mathbf{s}]_2)$ where $\mathbf{c}^2[1] = \mathbf{s}[1]$. Output
$$\mathsf{ct} := ([\mathbf{c}^1]_1, [\mathbf{c}^2]_2),\ \mathsf{kem} := [\alpha \mathbf{s}[1]]_T$$
- KeyGen: Run $(k^1, k^2) \leftarrow \mathsf{KeyGen}_0(y)$. Pick $\mathbf{r} \leftarrow \mathbb{Z}_p^m$. Compute $[\mathbf{k}^1]_1 := k^1([\alpha]_1, [\mathbf{r}]_1, [\mathbf{b} \otimes \mathbf{r}]_1)$, $[\mathbf{k}^2]_2 := k^2([\mathbf{r}]_2)$. Output
$$\mathsf{sk} := ([\mathbf{k}^1]_1, [\mathbf{k}^2]_2)$$

Note that Enc and KeyGen compute the linear functions c^1, k^1 “in the exponent” since it only knows $[\mathbf{b}]_1$ and not $\mathbf{b}$. We also require that c^1, k^1 depend only on a polynomial number of entries in $\mathbf{b}$, so that Enc, KeyGen only need to make a polynomial number of calls to H to compute $[c^1(\mathbf{s} \otimes \mathbf{b})]_1$ and $[k^1(\alpha, \mathbf{r}, \mathbf{b} \otimes \mathbf{r})]_1$ respectively. Depending on the application, some of these calls to H can also be pre-computed.

Remark 1 (Decryption). *Note that we can augment PES-ABE with an additional deterministic algorithm* Dec_0 *used in* Dec *where*

- $\mathsf{Dec}_0(x, y)$ *outputs* $\mathbf{e} \in \mathbb{Z}_p^{w_1 m_2}, \mathbf{e}' \in \mathbb{Z}_p^{w_2 m_1}$;
- $\mathsf{Dec}(\mathsf{mpk}, x, y, \mathsf{ct} = ([\mathbf{c}^1]_1, [\mathbf{c}^2]_2), \mathsf{sk} = ([\mathbf{k}^1]_1, [\mathbf{k}^2]_2)$: *Run* $(\mathbf{e}, \mathbf{e}') \leftarrow \mathsf{Dec}_0(x, y)$. *Compute* $[\mathbf{k}^1 \otimes \mathbf{c}^2]_T, [\mathbf{c}^1 \otimes \mathbf{k}^2]_T$ *using e, and output* $[(\mathbf{k}^1 \otimes \mathbf{c}^2) \cdot \mathbf{e}^\top + (\mathbf{c}^1 \otimes \mathbf{k}^2) \cdot \mathbf{e}'^\top]_T$.

It would then follow from ABE correctness that if $\mathsf{P}(x, y) = 1$, $(\mathbf{k}^1 \otimes \mathbf{c}^2) \cdot \mathbf{e}^\top + (\mathbf{c}^1 \otimes \mathbf{k}^2) \cdot \mathbf{e}'^\top = \alpha \mathbf{s}[1]$. *We omit* Dec_0 *in our presentation and instead, specify and analyze* Dec *for correctness directly. This does not affect our security notions and proofs which only refer to* Enc, KeyGen, Enc_0, KeyGen_0.

4 SYMBOLIC SECURITY OF PES-ABE

Following previous work [3, 6], we define symbolic security for PES-ABE, where we replace the inputs $(\alpha, \mathbf{b}, \mathbf{s}, \mathbf{r}) \leftarrow \mathbb{Z}_p \times \mathbb{Z}_p^n \times \mathbb{Z}_p^w \times \mathbb{Z}_p^m$ to the linear functions (c^1, c^2, k^1, k^2) with vectors of formal variables

$$(\tilde{\alpha}, \tilde{\mathbf{b}}, \tilde{\mathbf{s}}, \tilde{\mathbf{r}}) \leftarrow \mathsf{Var} \times \mathsf{Var}^n \times \mathsf{Var}^w \times \mathsf{Var}^m$$

In particular, $c^1(\tilde{\mathbf{s}} \otimes \tilde{\mathbf{b}}), c^2(\tilde{\mathbf{s}}), k^1(\tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}), k^2(\tilde{\mathbf{r}})$ are now (vectors of) polynomials in $\mathbb{Z}_p[\tilde{\alpha}, \tilde{\mathbf{b}}, \tilde{\mathbf{s}}, \tilde{\mathbf{r}}]$.

4.1 Definitions

Fix $x \in \mathcal{X}, y \in \mathcal{Y}$. ABE correctness tells us that if $\mathsf{P}(x, y) = 1$, then

$$\tilde{\alpha}\tilde{\mathbf{s}}[1] \in \mathrm{span}\big(c^1(\tilde{\mathbf{s}} \otimes \tilde{\mathbf{b}}) \otimes k^2(\tilde{\mathbf{r}}) \| k^1(\tilde{\alpha}, \tilde{\mathbf{r}}, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}) \otimes c^2(\tilde{\mathbf{s}})\big)$$

On the other hand, if $\mathsf{P}(x, y) = 0$, it should be the case that

$$\tilde{\alpha}\tilde{\mathbf{s}}[1] \notin \mathrm{span}\big(c^1(\tilde{\mathbf{s}} \otimes \tilde{\mathbf{b}}) \otimes k^2(\tilde{\mathbf{r}}) \| k^1(\tilde{\alpha}, \tilde{\mathbf{r}}, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}) \otimes c^2(\tilde{\mathbf{s}})\big)$$

Our basic formulation of symbolic security stipulates something stronger, where we basically replace $\tilde{\alpha}\tilde{\mathbf{s}}[1]$ with $\tilde{\alpha} \otimes c^2(\tilde{\mathbf{s}})$ and require $c^2(\tilde{\mathbf{s}})[1] = \tilde{\mathbf{s}}[1]$. In the special case where $c^2(\tilde{\mathbf{s}}) = \tilde{\mathbf{s}}[1]$ (as is the case when $w_2 = 1$), these two requirements are equivalent.

Definition 4.1 ((1, 1) Symbolic Security). For all $x \in \mathcal{X}, y \in \mathcal{Y}$ such that $\mathsf{P}(x, y) = 0$: we have

$$\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}^2\big) \cap \mathrm{span}\big(\mathbf{c}^1 \otimes \mathbf{k}^2 \parallel \mathbf{k}^1 \otimes \mathbf{c}^2\big) = \{0\}$$

where

$$\begin{aligned}
&(\tilde{\alpha}, \tilde{\mathbf{b}}) \leftarrow \mathsf{Var} \times \mathsf{Var}^n \\
&(\mathbf{c}^1, \mathbf{c}^2) := (c^1(\tilde{\mathbf{s}} \otimes \tilde{\mathbf{b}}), c^2(\tilde{\mathbf{s}})),\ \tilde{\mathbf{s}} \leftarrow \mathsf{Var}^w, (c^1, c^2) \leftarrow \mathsf{Enc}_0(x) \\
&(\mathbf{k}^1, \mathbf{k}^2) := (k^1(\tilde{\alpha}, \tilde{\mathbf{r}}, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}), k^2(\tilde{\mathbf{r}})),\ \tilde{\mathbf{r}} \leftarrow \mathsf{Var}^m, \\
&\qquad (k^1, k^2) \leftarrow \mathsf{KeyGen}_0(y).
\end{aligned}$$

The symbolic property captured by this definition will be required to prove many-ciphertext CPA security of our ABE scheme. To capture the ABE security experiment more closely, we extend the definition such that it also include many secret keys, many ciphertexts as well as the public key. Also we consider that in the ABE security experiment the adversary may ask for the same x or y multiple times. In Lemma 4.3 below, we show that this stronger symbolic property is actually implied by the weaker one above.

Definition 4.2 (Strong Symbolic Security). For all $Q_{\mathrm{ct}}, Q_{\mathrm{sk}} \in \mathbb{N}$, $X \in \mathcal{X}^{Q_{\mathrm{ct}}}, Y \in \mathcal{Y}^{Q_{\mathrm{sk}}}$ such that $\mathsf{P}(X[i], Y[j]) = 0$ for all $i \in [Q_{\mathrm{ct}}]$, $j \in [Q_{\mathrm{sk}}]$, we have

$$\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}_X^2\big) \cap \mathrm{span}\big(\tilde{\alpha} \parallel (1\|\tilde{\mathbf{b}}\|\mathbf{c}_X^1\|\mathbf{k}_Y^1) \otimes (1\|\mathbf{c}_X^2\|\mathbf{k}_Y^2)\big) = \{0\}$$

where

$$\begin{aligned}
&(\tilde{\alpha}, \tilde{\mathbf{b}}) \leftarrow \mathsf{Var} \times \mathsf{Var}^n \\
&(\mathbf{c}_i^1, \mathbf{c}_i^2) := (c^1(\tilde{\mathbf{s}}_i \otimes \tilde{\mathbf{b}}), c^2(\tilde{\mathbf{s}}_i)), \\
&\qquad \tilde{\mathbf{s}}_i \leftarrow \mathsf{Var}^w, (c_i^1, c_i^2) \leftarrow \mathsf{Enc}_0(X[i]), \forall i \in [Q_{\mathrm{ct}}], \\
&(\mathbf{k}_j^1, \mathbf{k}_j^2) := (k^1(\tilde{\alpha}, \tilde{\mathbf{r}}_j, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}_j), k^2(\tilde{\mathbf{r}}_j)), \\
&\qquad \tilde{\mathbf{r}}_j \leftarrow \mathsf{Var}^m, (k_j^1, k_j^2) \leftarrow \mathsf{KeyGen}_0(Y[j]), \forall j \in [Q_{\mathrm{sk}}], \\
&\mathbf{c}_X^1 := (\mathbf{c}_1^1 \| \cdots \| \mathbf{c}_{Q_{\mathrm{ct}}}^1), \mathbf{c}_X^2 := (\mathbf{c}_1^2 \| \cdots \| \mathbf{c}_{Q_{\mathrm{ct}}}^2) \\
&\mathbf{k}_Y^1 := (\mathbf{k}_1^1 \| \cdots \| \mathbf{k}_{Q_{\mathrm{sk}}}^1), \mathbf{k}_Y^2 := (\mathbf{k}_1^2 \| \cdots \| \mathbf{k}_{Q_{\mathrm{sk}}}^2) .
\end{aligned}$$

4.2 Relations

Now we can establish the desired implication in the following lemma.

Lemma 4.3. *If a PES-ABE scheme satisfies* $(1, 1)$ *symbolic security (Definition 4.1), then it also satisfies strong symbolic security (Definition 4.2).*

The proof follows the high-level strategy laid out in [6, Theorem 4.1] with two main differences: (i) the proof of Claim 2 where we handle $w_2 > 1$ (see also Remark 2) and (ii) Step 2 where we handle many-ciphertext security.

Proof. Fix a PES-ABE satisfying $(1, 1)$ symbolic security as well as $Q_{\mathrm{sk}}, Q_{\mathrm{ct}}, X, Y$ satisfying the conditions in Definition 4.2. We want to show that

$$\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}_X^2\big) \cap \mathrm{span}\big(\tilde{\alpha} \parallel (1\|\tilde{\mathbf{b}}\|\mathbf{c}_X^1\|\mathbf{k}_Y^1) \otimes (1\|\mathbf{c}_X^2\|\mathbf{k}_Y^2)\big) = \{0\} \quad (2)$$

The proof proceeds in three steps.

Step 1. First, we show that for all $i \in [Q_{\mathrm{ct}}]$,

$$\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}_i^2\big) \cap \mathrm{span}\big(\mathbf{c}_i^1 \otimes \mathbf{k}_Y^2 \parallel \mathbf{k}_Y^1 \otimes \mathbf{c}_i^2\big) = \{0\}$$

The proof proceeds by contradiction. Suppose on the contrary that there exist $i^* \in [Q_{\mathrm{ct}}]$, $\mathbf{e}^* \in \mathbb{Z}_p^{w_2}$, $\mathbf{e}_j \in \mathbb{Z}_p^{w_1 m_2}$, $\mathbf{e}'_j \in \mathbb{Z}_p^{m_1 w_2}$ for all $j \in [Q_{\mathrm{sk}}]$ such that $\mathbf{e}^* \neq \mathbf{0}$ and

$$(\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}^{*\top} = \sum_{j \in [Q_{\mathrm{sk}}]} (\mathbf{c}_{i^*}^1 \otimes \mathbf{k}_j^2) \cdot \mathbf{e}_j^\top + (\mathbf{k}_j^1 \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_j'^\top . \quad (3)$$

We claim that $\{\mathbf{e}_j, \mathbf{e}'_j\}_{j \in [Q_{\mathrm{sk}}]}$ then satisfies

- Claim 1: $(\mathbf{c}_{i^*}^1 \otimes \mathbf{k}_j^2) \cdot \mathbf{e}_j^\top + (k_j^1(0, \tilde{\mathbf{r}}_j, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}_j) \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_j'^\top = 0$ for all $j \in [Q_{\mathrm{sk}}]$.
- Claim 2: there exists $j^* \in [Q_{\mathrm{sk}}], \boldsymbol{\mu} \in \mathbb{Z}_p^{w_2}$ such that $\boldsymbol{\mu} \neq \mathbf{0}$ and $(\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2) \cdot \boldsymbol{\mu}^\top = (k_{j^*}^1(\tilde{\alpha}, \mathbf{0}, \mathbf{0}) \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_{j^*}'^\top$.

Combining the two claims with the fact that $\mathbf{k}_{j^*}^1 = k_{j^*}^1(0, \tilde{\mathbf{r}}_{j^*}, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}_{j^*}) + k_{j^*}^1(\tilde{\alpha}, \mathbf{0}, \mathbf{0})$, we have

$$(\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2) \cdot \boldsymbol{\mu}^\top = (\mathbf{c}_{i^*}^1 \otimes \mathbf{k}_{j^*}^2) \cdot \mathbf{e}_{j^*}^\top + (\mathbf{k}_{j^*}^1 \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_{j^*}'^\top .$$

which contradicts $(1, 1)$ symbolic security since $\mathsf{P}(X[i^*], Y[j^*]) = 0$. It remains to establish Claims 1 and 2 to complete the proof:

- Fix $j \in [Q_{\mathrm{sk}}]$. Claim 1 follows from evaluating (3) on $\tilde{\alpha} = 0, \tilde{\mathbf{r}}_{j'} = \mathbf{0}\ \forall j' \in [Q_{\mathrm{sk}}] \setminus \{j\}$.
- Next, evaluating (3) on $\tilde{\mathbf{r}}_j = \mathbf{0}\ \forall j \in [Q_{\mathrm{sk}}]$ yields

$$0 \neq (\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}^{*\top} = \sum_{j \in [Q_{\mathrm{sk}}]} (k_j^1(\tilde{\alpha}, \mathbf{0}, \mathbf{0}) \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_j'^\top .$$

Therefore, there exists $j^* \in Y$ such that $(k_{j^*}^1(\tilde{\alpha}, 0, 0) \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_{j^*}'^\top \neq 0$. Moreover, since the polynomial $k_{j^*}^1(\tilde{\alpha}, \mathbf{0}, \mathbf{0})$ is linear in $\tilde{\alpha}$, there exists $\boldsymbol{\mu} \neq \mathbf{0}$ such that $(k_{j^*}^1(\tilde{\alpha}, 0, 0) \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_{j^*}'^\top = (\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2) \cdot \boldsymbol{\mu}^\top$ and Claim 2 follows.

Step 2. We show that

$$\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}_X^2\big) \cap \mathrm{span}\big(\mathbf{c}_X^1 \otimes \mathbf{k}_Y^2 \parallel \mathbf{k}_Y^1 \otimes \mathbf{c}_X^2\big) = \{0\} .$$

As in the previous step, the proof proceeds by contradiction. Suppose the above statement is false, which means there exist $\{\mathbf{e}_i^* \in \mathbb{Z}_p^{w_2}, \mathbf{e}_i \in \mathbb{Z}_p^{Q_{\mathrm{sk}} \cdot w_1 m_2}, \mathbf{e}'_i \in \mathbb{Z}_p^{Q_{\mathrm{sk}} \cdot m_1 w_2}\}_{i \in [Q_{\mathrm{ct}}]}$ and $i^* \in [Q_{\mathrm{ct}}]$ such that

$$\sum_{i \in [Q_{\mathrm{ct}}]} (\tilde{\alpha} \otimes \mathbf{c}_i^2) \cdot \mathbf{e}_i^{*\top} = \sum_{i \in [Q_{\mathrm{ct}}]} (\mathbf{c}_i^1 \otimes \mathbf{k}_Y^2) \cdot \mathbf{e}_i^\top + (\mathbf{k}_Y^1 \otimes \mathbf{c}_i^2) \cdot \mathbf{e}_i'^\top , \quad (4)$$

and $\mathbf{e}_{i^*}^* \neq \mathbf{0}$. We evaluate (4) on $\tilde{\mathbf{s}}_{i'} = \mathbf{0}\ \forall i' \in [Q_{\mathrm{ct}}] \setminus \{i^*\}$ and get

$$(\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_{i^*}^{*\top} = (\mathbf{c}_{i^*}^1 \otimes \mathbf{k}_Y^2) \cdot \mathbf{e}_{i^*}^\top + (\mathbf{k}_Y^1 \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_{i^*}'^\top .$$

That is, $\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}_{i^*}^2\big) \cap \mathrm{span}\big(\mathbf{c}_{i^*}^1 \otimes \mathbf{k}_Y^2 \| \mathbf{k}_Y^1 \otimes \mathbf{c}_{i^*}^2\big) \neq \{0\}$, which contradicts what we showed in Step 1.

Step 3. We now prove (2), which also proceeds by contradiction. Suppose on the contrary that

$$\mathrm{span}\big(\tilde{\alpha} \otimes \mathbf{c}_X^2\big) \cap \mathrm{span}\big(\tilde{\alpha} \| (1\|\tilde{\mathbf{b}}\|\mathbf{c}_X^1\|\mathbf{k}_Y^1) \otimes (1\|\mathbf{c}_X^2\|\mathbf{k}_Y^2)\big) \neq \{0\} .$$

Then there exist $\mathbf{e}^* \in \mathbb{Z}_p^{Q_{\mathrm{ct}} \cdot w_2}$, $\mathbf{e}_{PK} \in \mathbb{Z}_p^{2+n}$, $\mathbf{e}_X \in \mathbb{Z}_p^{Q_{\mathrm{ct}} \cdot (w_1 + w_2 + w_1 w_2 + n w_2)}$, $\mathbf{e}_Y \in \mathbb{Z}_p^{Q_{\mathrm{sk}} \cdot (m_1 + m_2 + m_1 m_2 + n m_2)}$, $\mathbf{e}_{XY} \in \mathbb{Z}_p^{Q_{\mathrm{ct}} \cdot Q_{\mathrm{sk}} \cdot (w_1 m_2 + m_1 w_2)}$ such that

$$\begin{aligned}(\tilde{\alpha} \otimes \mathbf{c}_X^2) \cdot \mathbf{e}^{*\top} = (1 \| \tilde{\alpha} \| \tilde{\mathbf{b}}) \cdot \mathbf{e}_{PK}^\top &+ (\mathbf{c}_X^1 \| \mathbf{c}_X^2 \| \mathbf{c}_X^1 \otimes \mathbf{c}_X^2 \| \tilde{\mathbf{b}} \otimes \mathbf{c}_X^2) \cdot \mathbf{e}_X^\top \\ &+ (\mathbf{k}_Y^1 \| \mathbf{k}_Y^2 \| \mathbf{k}_Y^1 \otimes \mathbf{k}_Y^2 \| \tilde{\mathbf{b}} \otimes \mathbf{k}_Y^2) \cdot \mathbf{e}_Y^\top \qquad (5) \\ &+ (\mathbf{c}_X^1 \otimes \mathbf{k}_Y^2 \| \mathbf{k}_Y^1 \otimes \mathbf{c}_X^2) \cdot \mathbf{e}_{XY}^\top\end{aligned}$$

and $\mathbf{e}^* \neq \mathbf{0}$. First, we look at the first three terms on the RHS of (5):

- Evaluating (5) on $\tilde{\alpha} = 0$, $\tilde{\mathbf{r}}_Y = \mathbf{0}$, and $\tilde{\mathbf{s}}_X = 0$ yields $(1 \| 0 \| \tilde{\mathbf{b}}) \cdot \mathbf{e}_{PK}^\top = 0$.
- Evaluating (5) on $\tilde{\alpha} = 0$, $\tilde{\mathbf{r}}_Y = \mathbf{0}$ yields $(1 \| 0 \| \tilde{\mathbf{b}}) \cdot \mathbf{e}_{PK}^\top + (\mathbf{c}_X^1 \| \mathbf{c}_X^2 \| \mathbf{c}_X^1 \otimes \mathbf{c}_X^2 \| \tilde{\mathbf{b}} \otimes \mathbf{c}_X^2) \cdot \mathbf{e}_X^\top = 0$.
- Evaluating (5) on $\tilde{\mathbf{s}}_X = 0$ yields $(1 \| \tilde{\alpha} \| \tilde{\mathbf{b}}) \cdot \mathbf{e}_{PK}^\top + (\mathbf{k}_Y^1 \| \mathbf{k}_Y^2 \| \mathbf{k}_Y^1 \otimes \mathbf{k}_Y^2 \| \tilde{\mathbf{b}} \otimes \mathbf{k}_Y^2) \cdot \mathbf{e}_Y^\top = 0$.

Subtracting the first equality from the sum of the second and third implies that the sum of the first three terms on the RHS of (5) is 0. This means

$$(\tilde{\alpha} \otimes \mathbf{c}_X^2) \cdot \mathbf{e}^{*\top} = (\mathbf{c}_X^1 \otimes \mathbf{k}_Y^2 \| \mathbf{k}_Y^1 \otimes \mathbf{c}_X^2) \cdot \mathbf{e}_{XY}^\top$$

which contradicts what we showed in Step 2. □

Remark 2 (handling $w_1 > 1$). *In the proof of the analogue of Claim 2 in [6], they start with*

$$\tilde{\alpha}\mathbf{c}_{i^*}^2[1] = \sum_{j \in [Q_{\mathsf{sk}}]} (k_j^1(\tilde{\alpha}, \mathbf{0}, \mathbf{0}) \otimes \mathbf{c}_{i^*}^2) \cdot \mathbf{e}_j'^\top \ .$$

They show that if $w_2 = 1$ (a requirement mentioned in the proof[7] but not in the theorem statement), then there exists $j^* \in Y$ such $\mu \cdot \mathbf{c}_{i^*}^2[1] = k_{j^*}^1(1, \mathbf{0}, \mathbf{0}) \otimes \mathbf{c}_{i^*}^2 \cdot \mathbf{e}_{j^*}'^\top$ and $\mu \neq 0$. However, if we allow $w_2 > 1$, then this claim does not hold in general. In particular, it could be that for all j, $\mathbf{c}_{i^*}^2[1]$ only appears in a linear combination with other elements of $\mathbf{c}_{i^*}^2$, which then all together sum up to $\mathbf{c}_{i^*}^2[1]$. For this reason, we need to strengthen our definition accordingly.

5 OPTIMAL ABE SECURITY IN THE GGM

We prove symbolic security of PES-ABE implies optimal, adaptive security in the generic group model (GGM). For that, we first recall the generic group model.

5.1 Generic Group Model

In the generic group model, an adversary can perform group operations only via oracle access. We adopt the model by Maurer [43] extended to the pairing group setting, where apart from the group operation, the adversary can also compute the pairing via an oracle. A third party implements the pairing group and maintains a list for $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$. Each list stores group elements of queries by the adversary. Depending on the query, one or multiple entries are appended to the different lists. The adversary can access each entry of the lists by a handle, which is a list index $i \in \mathbb{N}$ and a list identifier $s \in \{1, 2, T\}$. It can also perform equality queries to check if two entries of the same list contain the same group element.

In game G_0 in Figure 2, we model the ABE security game from Section 2.2 in the GGM. That is, the adversary also gets access to oracles $\mathcal{O}_{\mathsf{mpk}}$, $\mathcal{O}_{\mathsf{ct}}$ and $\mathcal{O}_{\mathsf{sk}}$. On each query, the corresponding oracle returns the current length of all modified lists from which the adversary can deduce the corresponding handles since length of ciphertexts and secret keys follow from the definition of the scheme. Furthermore, we model the hash function in our scheme as random oracle, so we additionally provide an oracle H, which also modify the lists. The adversary can then use these indices in further group operation and equality queries as described above.

5.2 Security

The following theorem states that symbolic security implies optimal, adaptive security in GGM.

Theorem 5.1. *Let $\lambda \in \mathbb{N}$ be the security parameter and $\mathcal{A}$ be an adversary that on input $(1^\lambda, p)$ makes $Q_{\mathsf{add}}, Q_{\mathsf{pair}}, Q_{\mathsf{ct}}, Q_{\mathsf{sk}}, Q_{\mathsf{eq}}$ queries to oracles $\mathcal{O}_{\mathsf{add}}, \mathcal{O}_{\mathsf{pair}}, \mathcal{O}_{\mathsf{ct}}, \mathcal{O}_{\mathsf{sk}}, \mathcal{O}_{\mathsf{eq}}$ and Q_{H} queries to the random oracle H. If PES-ABE is (1,1) symbolically secure (Definition 4.1), then it is adaptively Many-CT secure in the GGM. In particular,*

$$\mathsf{Adv}_{\mathsf{ABE},\mathcal{A}}^{\mathsf{GGM}}(\lambda) \leq \frac{3 \cdot (Q_{\mathsf{H}} + (w'+1) \cdot Q_{\mathsf{ct}} + m' \cdot Q_{\mathsf{sk}} + Q_{\mathsf{add}} + Q_{\mathsf{pair}})^2}{p}$$

where $w' := w_1 + w_2$ and $m' := m_1 + m_2$.

We provide the games for the proof in Figure 2. The full proof can be found in the full version of the paper. In fact, it is similar to that in [6, Theorem 3.3]. The latter only considers single-ciphertext ABE security, and achieves an additional loss of Q_{sk}, since they apply Schwartz-Zippel to polynomials of degree Q_{sk} in order to handle rational fractions arising in their schemes.

6 OUR SCHEMES: PUTTING EVERYTHING TOGETHER

We now show that the FABEO CP-ABE and KP-ABE schemes for monotone span programs described in Figure 1 satisfy the PES-ABE framework and $(1,1)$ symbolic security described in Section 2.2. Combined with the statements from Lemma 4.3 and Theorem 5.1, this establishes optimal, adaptive security of our CP-ABE and KP-ABE schemes in GGM (Corollaries 6.1 and 6.2).

6.1 CP-ABE

Our CP-ABE scheme is shown in Figure 1. It builds upon the pair encoding scheme 11 in [7] and that in Appendix B.1 in [3] and extends them by attribute hashing and multi-use of attributes. In particular, we can describe the underlying PES-ABE as follows.

- $\underline{\mathsf{Setup}_0.}$ Output $n := |\mathcal{U}| + 1$.
- $\underline{\mathsf{Enc}_0(\mathbf{M}, \pi).}$ Set $w = n_1 + \tau$, $w_1 = n_1$, $w_2 = \tau + 1$, and output (c^1, c^2) where we parse $\mathbf{s}$ as $(s_1 \| \mathbf{v} \| \mathbf{s}')$ and

$$c^1(\mathbf{s} \otimes \mathbf{b}) := (\mathbf{M}_i(s_1 \| \mathbf{v})^\top \cdot \mathbf{b}[|\mathcal{U}|+1] + \mathbf{s}'[\rho(i)] \cdot \mathbf{b}[\pi(i)])_{i \in [n_1]},$$
$$c^2(s_1) := (s_1 \| \mathbf{s}')$$

- $\underline{\mathsf{KeyGen}_0(\mathcal{S}).}$ Set $m = 1$, $m_1 = |\mathcal{S}| + 1$, $m_2 = 1$, and output (k^1, k^2) where we parse $\mathbf{r}$ as (r) and

$$k^1(\alpha, \mathbf{r}, \mathbf{b} \otimes \mathbf{r}) := (\alpha + r\mathbf{b}[|\mathcal{U}|+1] \| (r\mathbf{b}[u])_{u \in \mathcal{S}}),$$
$$k^2(\mathbf{r}) := (r)$$

Correctness. Let $\mathsf{ct} = (\mathsf{ct}_1, (\mathsf{ct}_{2,j})_{j \in [\tau]}, (\mathsf{ct}_{3,i})_{i \in [n_1]})$ be a ciphertext for $(\mathbf{M}, \pi)$ and $\mathsf{sk} = (\mathsf{sk}_1, (\mathsf{sk}_{2,u})_{u \in \mathcal{S}}, \mathsf{sk}_3)$ be a secret key for $\mathcal{S}$ as defined in Figure 1. Further let $\mathbf{b}[u]$ such that $\mathsf{H}(u) = g_1^{\mathbf{b}[u]}$

[7] On page 662, they wrote "since we assumed $w_1 = 0$". Here, $c^2(\tilde{s})$ corresponds to $\overrightarrow{S} = (S_0, \dots, S_{w_1})$ in [6].

```
Games $\{\mathrm{G}_\mu\}_{\mu\in[0,Q_{\mathrm{eq}}]}$
$i = j := 0,\ \nu := 0,\ X = Y = \mathcal{H} := \emptyset$
for $s \in \{1,2,T\}$: $L_s := \emptyset,\ L^{\sim}_s := \emptyset$
$\beta \xleftarrow{\$} \{0,1\}$
$\beta' \leftarrow \mathcal{A}^{O_{\mathrm{mpk}},O_{\mathrm{add}},O_{\mathrm{pair}},O_{\mathrm{ct}},O_{\mathrm{sk}},O_{\mathrm{eq}},\mathsf{H}}(1^\lambda, p)$
return $[\![\beta = \beta']\!]$ and $[\![\mathsf{P}(X[i], Y[j]) = 0\ \forall i \in [Q_{\mathrm{ct}}],\ j \in [Q_{\mathrm{sk}}]]\!]$

$O_{\mathrm{mpk}}$                                    // first query, only once
$n \leftarrow$ Param, $(\alpha, \mathbf{b}) \xleftarrow{\$} \mathbb{Z}_p \times \mathbb{Z}_p^n$, $(\tilde{\alpha}, \tilde{\mathbf{b}}) \leftarrow \mathsf{Var} \times \mathsf{Var}^n$
$L_1$.append$(1)$, $L_2$.append$(1)$, $L_T$.append$(\alpha)$
$L^{\sim}_1$.append$(1)$, $L^{\sim}_2$.append$(1)$, $L^{\sim}_T$.append$(\tilde{\alpha})$
return $|L_1|, |L_2|, |L_T|$

$O_{\mathrm{add}}(s \in \{1,2,T\}, i', j' \in \mathbb{N})$
$L_s$.append$(L_s[i'] + L_s[j'])$
$L^{\sim}_s$.append$(L^{\sim}_s[i'] + L^{\sim}_s[j'])$
return $|L_s|$

$O_{\mathrm{pair}}(i', j' \in \mathbb{N})$
$L_T$.append$(L_1[i'] \cdot L_2[j'])$
$L^{\sim}_T$.append$(L^{\sim}_1[i'] \cdot L^{\sim}_2[j'])$
return $|L_T|$

$O_{\mathrm{eq}}(s \in \{1,2,T\}, i', j')$
$\nu := \nu + 1$
if $\nu \le \mu$ : return $L^{\sim}_s[i'] = L^{\sim}_s[j']$
return $L_s[i'] = L_s[j']$

$O_{\mathrm{ct}}(x \in \mathcal{X})$
$(c_i^1(\tilde{\mathbf{s}}_i \otimes \tilde{\mathbf{b}}), c_i^2(\tilde{\mathbf{s}}_i)) \leftarrow \mathsf{Enc}_0(x)$
$\mathbf{s}_i \xleftarrow{\$} \mathbb{Z}_p^w, \tilde{\mathbf{s}}_i \leftarrow \mathsf{Var}^w$
$d_i^{(0)} := \alpha\mathbf{s}_i[1],\ d_i^{(1)} := \omega_i \xleftarrow{\$} \mathbb{Z}_p$
$\tilde{d}_i^{(0)} := \tilde{\alpha}\tilde{\mathbf{s}}_i[1],\ \tilde{d}_i^{(1)} := \tilde{\omega}_i \leftarrow \mathsf{Var}$
$L_1$.append$(c_i^1(\mathbf{s}_i \otimes \mathbf{b}))$, $L_2$.append$(c_i^2(\mathbf{s}_i))$, $L_T$.append$(d_i^{(\beta)})$
$L^{\sim}_1$.append$(c_i^1(\tilde{\mathbf{s}}_i \otimes \tilde{\mathbf{b}}))$, $L^{\sim}_2$.append$(c_i^2(\tilde{\mathbf{s}}_i))$, $L^{\sim}_T$.append$(\tilde{d}_i^{(\beta)})$
$X$.append$(x)$, $i := i + 1$
return $|L_1|, |L_2|, |L_T|$

$O_{\mathrm{sk}}(y \in \mathcal{Y})$
$(k_j^1(\tilde{\alpha}, \tilde{\mathbf{r}}_j, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}_j), k_j^2(\tilde{\mathbf{r}}_j)) \leftarrow \mathsf{KeyGen}_0(y)$
$\mathbf{r}_j \xleftarrow{\$} \mathbb{Z}_p^m, \tilde{\mathbf{r}}_j \leftarrow \mathsf{Var}^m$
$L_1$.append$(k_j^1(\alpha, \mathbf{r}_j, \mathbf{b} \otimes \mathbf{r}_j))$, $L_2$.append$(k_j^2(\mathbf{r}_j))$
$L^{\sim}_1$.append$(k_j^1(\tilde{\alpha}, \tilde{\mathbf{r}}_j, \tilde{\mathbf{b}} \otimes \tilde{\mathbf{r}}_j))$, $L^{\sim}_2$.append$(k_j^2(\tilde{\mathbf{r}}_j))$
$Y$.append$(y)$, $j := j + 1$
return $|L_1|, |L_2|$

$\mathsf{H}(u)$
$L_1$.append$(\mathbf{b}[u])$ , $L^{\sim}_1$.append$(\tilde{\mathbf{b}}[u])$
$\mathcal{H} := \mathcal{H} \cup \{u\}$
return $|L_1|$
```

Figure 2: Games G_μ for $\mu \in [0, Q_{\mathrm{eq}}]$ for the proof of Theorem 5.1. Note that the games only differ in oracle O_{eq} (which depends on μ). Here, G_0 corresponds to the GGM experiment that makes only use of components in light gray frames, whereas $\mathrm{G}_{Q_{\mathrm{eq}}}$ makes only use of components in dark gray frames. W.l.o.g. we assume that no query to $O_{\mathrm{add}}, O_{\mathrm{pair}}, O_{\mathrm{eq}}$ contains indices $i', j' \in \mathbb{N}$ which exceed the size of the involved lists.

and b' such that $\mathsf{H}(|\mathcal{U}|+1) = g_1^{b'}$. If $\mathcal{S}$ satisfies $(\mathbf{M}, \pi)$, then there exist constants $(\gamma_i)_{i\in[n_1]}$ such that $\sum_{i\in I} \gamma_i \mathbf{M}_i = (1, 0, \ldots, 0)$ and decryption computes

(1) $e(g_1^\alpha \cdot \mathsf{H}(|\mathcal{U}|+1)^r, g_2^{s_1}) = [\alpha s_1 + b' r s_1]_T$

(2) $\prod_{j\in[\tau]} e(\prod_{i\in I,\rho(i)=j} \mathsf{H}(\pi(i))^{\gamma_i r}, g_2^{\mathbf{s}'[j]})$
$= [r \sum_{j\in[\tau]} \sum_{i\in I,\rho(i)=j} \gamma_i \mathbf{b}[\pi(i)]\mathbf{s}'[j]]_T$

(3) $e(\prod_{i\in I}(\mathsf{H}(|\mathcal{U}|+1)^{\gamma_i \mathbf{M}_i (s_1\|\mathbf{v})^\top} \cdot \mathsf{H}(\pi(i))^{\gamma_i \mathbf{s}'[\rho(i)]}), g_2^r)$
$$= [br' \underbrace{\sum_{i\in I} \gamma_i \mathbf{M}_i((s_1\|\mathbf{v})^\top}_{=s_1}]_T \cdot \underbrace{[r \sum_{i\in I} \gamma_i \mathbf{b}[\pi(i)]\mathbf{s}'[\rho(i)]]_T}_{=(2)}$$

Note that by definition of ρ, (2) and the second term of (3) are the same. Thus computing $(1) \cdot (2)/(3)$ yields $d = [\alpha s_1]_T$.

Symbolic Security. We need to show that for all $(\mathbf{M}, \pi) \in \mathcal{X}$, $\mathcal{S} \in \mathcal{Y}$ such that $\mathsf{P}((\mathbf{M}, \pi), \mathcal{S}) = 0$, it holds that

$$\mathrm{span}\big(\tilde{\alpha} \otimes (\tilde{s}_1\|\tilde{\mathbf{s}}')\big) \cap \mathrm{span}\big((\mathbf{M}_i(\tilde{s}_1, \tilde{\mathbf{v}})^\top \tilde{b}' + \tilde{\mathbf{s}}'[\rho(i)]\tilde{\mathbf{b}}[\pi(i)])_{i\in[n_1]} \otimes \tilde{r} \| \\ (\tilde{\alpha} + \tilde{r}\tilde{b}' \| (\tilde{r}\tilde{\mathbf{b}}[u])_{u\in\mathcal{S}}) \otimes (\tilde{s}_1\|\tilde{\mathbf{s}}')\big) = \{0\},$$

where we define $\tilde{b}' := \tilde{\mathbf{b}}[|\mathcal{U}|+1]$.

We prove this property by contradiction. So assume there exists $\mathbf{e}^* \in \mathbb{Z}_p^2$, $\mathbf{e}$, $\mathbf{e}'^{(1)}$, $\mathbf{e}'^{(2)}$, $\mathbf{e}'^{(3)}$ such that $\mathbf{e}^* \neq \mathbf{0}$ and

$$(\tilde{\alpha} \otimes (\tilde{s}_1\|\tilde{\mathbf{s}}')) \cdot \mathbf{e}^{*\top} = (\mathbf{M}_i(\tilde{s}_1\|\tilde{\mathbf{v}})^\top \tilde{b}'\tilde{r} + \tilde{\mathbf{s}}'[\rho(i)]\tilde{\mathbf{b}}[\pi(i)]\tilde{r})_{i\in[n_1]} \cdot \mathbf{e}^\top \\ + ((\tilde{\alpha} + \tilde{r}\tilde{b}') \otimes (\tilde{s}_1\|\tilde{\mathbf{s}}')) \cdot \mathbf{e}'^{(1)\top} \\ + (\{\tilde{r}\tilde{\mathbf{b}}[u]\tilde{s}_1\}_{u\in\mathcal{S}}) \cdot \mathbf{e}'^{(2)\top} + (\{\tilde{r}\tilde{\mathbf{b}}[u] \otimes \tilde{\mathbf{s}}'\}_{u\in\mathcal{S}}) \cdot \mathbf{e}'^{(3)\top}$$

Now we use the fact that $\mathsf{P}((\mathbf{M}, \pi), \mathcal{S}) = 0$. Recall that this means that there exists a vector $\mathbf{w} \in \mathbb{Z}_p^{n_2}$ such that $\langle \mathbf{w}, \mathbf{M}_i \rangle = 0$ for all $\pi(i) \in \mathcal{S}$ and that $\mathbf{w}[1] = 1$. So evaluating on $(\tilde{s}_1\|\tilde{\mathbf{v}}) = \mathbf{w}$ gives us

$$(\tilde{\alpha} \otimes (1\|\tilde{\mathbf{s}}')) \cdot \mathbf{e}^{*\top} = (\tilde{\mathbf{s}}'[\rho(i)]\tilde{\mathbf{b}}[\pi(i)]\tilde{r})_{i\in[n_1],\pi(i)\in\mathcal{S}} \cdot \overline{\mathbf{e}}^\top \\ + (\mathbf{M}_i \mathbf{w}^\top \tilde{b}'\tilde{r} + \tilde{\mathbf{s}}'[\rho(i)]\tilde{\mathbf{b}}[\pi(i)]\tilde{r})_{i\in[n_1],\pi(i)\notin\mathcal{S}} \cdot \underline{\mathbf{e}}^\top \\ + ((\tilde{\alpha} + \tilde{r}\tilde{b}') \otimes (1\|\tilde{\mathbf{s}}')) \cdot \mathbf{e}'^{(1)\top} \\ + (\tilde{r}\tilde{\mathbf{b}}[u])_{u\in\mathcal{S}} \cdot \mathbf{e}'^{(2)\top} + (\tilde{r}\tilde{\mathbf{b}}[u] \otimes \tilde{\mathbf{s}}')_{u\in\mathcal{S}} \cdot \mathbf{e}'^{(3)\top}$$

where we split $\mathbf{e}$ into two vectors $\overline{\mathbf{e}} \in \mathbb{Z}_p^{|\mathcal{S}|}$ and $\underline{\mathbf{e}} \in \mathbb{Z}_p^{n_1-|\mathcal{S}|}$, capturing those rows of $\mathbf{M}$ that belong to $u \in \mathcal{S}$ and those that do not belong to an attribute in $\mathcal{S}$. Note that the monomials $\{\tilde{r}\tilde{\mathbf{b}}[u] \otimes \tilde{\mathbf{s}}'\}_{u\in\mathcal{S}}$ only appear in the first and the last term. By definition of ρ, the monomials $\tilde{\mathbf{s}}'[\rho(i)]\tilde{\mathbf{b}}[\pi(i)]\tilde{r}$ in the first term are mutually distinct. Thus, we must have $\overline{\mathbf{e}}[i] = -\mathbf{e}'^{(3)}[j]$ for all $i \in [n_1]$ such that $\pi(i) \in \mathcal{S}$ and unique indices j, while all other entries in $\mathbf{e}'^{(3)}$ must be 0. Further looking at monomials on the RHS, $\tilde{\mathbf{s}}'[\rho(i)]\tilde{\mathbf{b}}[\pi(i)]\tilde{r}$ for $\pi(i) \notin \mathcal{S}$ and $\tilde{r}\tilde{\mathbf{b}}[u]$ for $u \in \mathcal{S}$ are also mutually distinct and only appear in one of the terms, thus $\underline{\mathbf{e}}$ as well as $\mathbf{e}'^{(2)}$ must be $\mathbf{0}$. Therefore, the following equation must be satisfied

$$(\tilde{\alpha} \otimes (1\|\tilde{\mathbf{s}}')) \cdot \mathbf{e}^{*\top} = ((\tilde{\alpha} + \tilde{r}\tilde{b}') \otimes (1\|\tilde{\mathbf{s}}')) \cdot \mathbf{e}'^{(1)\top},$$

which leads to a contradiction that $\mathbf{e}^* \neq \mathbf{0}$ since $\tilde{r}\tilde{b}'$ only appears on the RHS.

Corollary 6.1. *Let $\lambda \in \mathbb{N}$ be the security parameter and $\mathcal{A}$ be an adversary that on input $(1^\lambda, p)$ makes Q_{op} group operation queries*

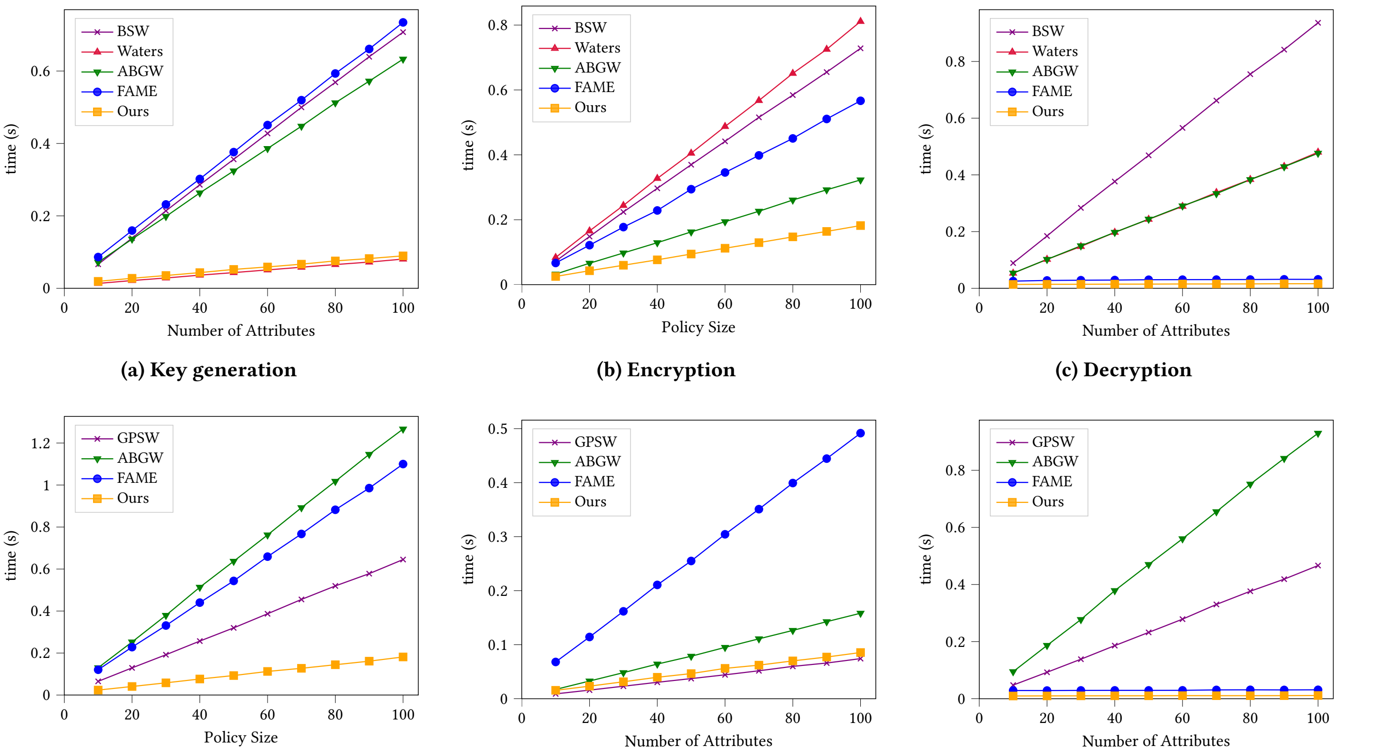

(a) Key generation **(b) Encryption** **(c) Decryption**

(d) Key generation **(e) Encryption** **(f) Decryption**

Figure 3: Running times for CP-ABE (top) and KP-ABE (bottom) schemes. We use one-use formulas (i.e., $\tau = 1$). In particular, for 100 attributes, CP-ABE decryption takes 0.016s in FABEO and 0.032s in FAME, and KP-ABE decryption takes 0.011s in FABEO and 0.031s in FAME.

to oracles O_{add} and O_{pair}, as well as Q_{ct}, Q_{sk} queries to oracles O_{ct}, O_{sk}, and Q_{H} queries to the random oracle H. CP-ABE *is adaptively secure in the GGM such that*

$$\mathrm{Adv}^{\mathrm{GGM}}_{\text{CP-ABE},\mathcal{A}}(\lambda) \le \frac{3 \cdot (Q_{\mathrm{H}} + (n_1+3) \cdot Q_{\mathrm{ct}} + (|\mathcal{S}|+2) \cdot Q_{\mathrm{sk}} + Q_{\mathrm{op}})^2}{p},$$

where $|\mathcal{S}|$ is the maximum size of the attribute sets queried to O_{sk} and n_1 is the maximum number of rows of $\mathbf{M}$ *queried to O_{ct}.*

6.2 KP-ABE

Our KP-ABE scheme is shown in Figure 1. The proof of correctness and symbolic security are similar to that of the CP-ABE scheme and can be found in the full version of the paper.

Corollary 6.2. *Let $\lambda \in \mathbb{N}$ be the security parameter and $\mathcal{A}$ be an adversary that on input $(1^\lambda, p)$ makes Q_{op} group operation queries to oracles O_{add} and O_{pair}, as well as Q_{ct}, Q_{sk} queries to oracles O_{ct}, O_{sk}, and Q_{H} queries to the random oracle* H. KP-ABE *is adaptively secure in the GGM such that*

$$\mathrm{Adv}^{\mathrm{GGM}}_{\text{KP-ABE},\mathcal{A}}(\lambda) \le \frac{3 \cdot (Q_{\mathrm{H}} + (|\mathcal{S}|+2) \cdot Q_{\mathrm{ct}} + (n_1+1) \cdot Q_{\mathrm{sk}} + Q_{\mathrm{op}})^2}{p},$$

where $|\mathcal{S}|$ is the maximum size of the attribute sets queried to O_{ct} and n_1 is the maximum number of rows of $\mathbf{M}$ *queried to O_{sk}.*

7 IMPLEMENTATION AND EVALUATION

We use two metrics to compare our scheme with prior work, the first is in terms of efficiency and the second is in terms of tightness.

7.1 Efficiency

We implemented several ABE schemes in Python 2.7.12 using the Charm 0.43 framework [4] and the MNT224 curve for pairings.[8] We ran the schemes on a Lenovo Thinkpad Yoga X1 laptop with a 1.80GHz Intel Core i7-10510U CPU and 16GB RAM. Our implementation extends the code of Agrawal and Chase [1] and we provide the implementation on GitHub [44]. In particular, we compare the CP-ABE and KP-ABE schemes described in Table 1.

All schemes are implemented in the asymmetric setting. Agrawal and Chase already transferred the original constructions of BSW, Waters and GPSW that use symmetric bilinear maps to the asymmetric setting [2, Appendices D-F]. Apart from our schemes, we additionally implement the unbounded CP-ABE and KP-ABE of ABGW.

In our experiment, we use access policies of the form "Attr_1 **and** Attr_2 **and** ... **and** Attr_N" for $N \in \{10, 20, \ldots, 100\}$ without re-use (i.e., $\tau = 1$). This way, $|\mathcal{S}| = n_1 = n_2 := N$ and all attributes

[8] The implementations in FAME and ABGW also use the Charm framework. Unfortunately, the PBC library used in Charm does not support BLS12-381.

	Key generation					*Encryption*					*Decryption*		
	$\mathbb{G}_1$			$\mathbb{G}_2$		$\mathbb{G}_1$			$\mathbb{G}_2$		$\mathbb{G}_1$	$\mathbb{G}_T$	
Schemes	Mul	Exp	Hash	Mul	Exp	Mul	Exp	Hash	Mul	Exp	Mul	Mul	Pair
BSW	$m+1$	$m+2$	m	-	m	-	n_1	n_1	-	n_1+1	-	$2I+1$	$2I+1$
Waters	1	$m+1$	-	-	1	n_1	$2n_1$	-	-	n_1+1	I	$I+2$	$I+2$
FAME	$6\hat{\tau}m+9$	$9\hat{\tau}m+9$	$6\hat{\tau}m+6$	-	3	$6n_1n_2+3n_1$	$6n_1$	$6n_1+6n_2$	-	3	$6I+3$	6	6
ABGW	-	-	-	-	$2m+1$	$2n_1$	$5n_1$	-	-	-	$2I$	$I+2$	$I+2$
Ours	1	$m+2$	$m+1$	-	1	n_1	$2n_1$	n_1+1	-	$\tau+1$	$2I$	$\tau+2$	$\tau+2$
GPSW	-	-	-	-	n_1	-	m	-	-	-	-	I	I
FAME	$9n_1n_2+3n_1$	$9n_1+3n_2$	$6n_1+6n_2$	-	3	$3\hat{\tau}m$	$6\hat{\tau}m$	$6\hat{\tau}m$	-	3	$6I$	6	6
ABGW	-	-	-	-	$2n_1$	$2m$	$3m+1$	-	-	-	-	$2I$	$2I$
Ours	n_1	$2n_1$	n_1	-	τ	-	m	m	-	1	$2I$	$\tau+1$	$\tau+1$

Table 2: Number of group operations in $\mathbb{G}_1$ and $\mathbb{G}_2$ for key generation and encryption of CP-ABE (top) and KP-ABE (bottom) schemes. m denotes the number of attributes in the set $\mathcal{S}$, n_1 and n_2 are the number of rows and columns of the MSP matrix and τ denotes the maximum number of multi-use. I denotes the number of attributes used in decryption (counted with multiplicity). Note that $\tau \leq I$. The experiments and most comparison in the text consider $\hat{\tau} = \tau = 1$.

are used in decryption. As [2], we first convert the policies into a Boolean formula and then to an MSP using the Lewko-Waters' method [39]. This way, the matrix $\mathbf{M}$ has only entries in $\{0, 1, -1\}$ and the reconstruction coefficients are always 0 or 1, reducing the number of exponentiations.

In Figure 3, we show the average running times for the key generation, encryption and decryption algorithms. All our experiments compute the average time in 20 executions. It is worth noting that each algorithm of our two schemes performs better or comparatively the same as all the others. These results are supported by our theoretical overview in Table 2 which lists the number of multiplications and exponentiations for each group as well as the number of hashing and pairing operations. Recall also that exponentiation in $\mathbb{G}_2$ is much slower than in $\mathbb{G}_1$ and the pairing operation is comparatively expensive. Additionally, we provide the number of group elements of secret keys and ciphertexts in Table 3. Since in general elements in $\mathbb{G}_2$ are about 2 to 3 times the size of elements in $\mathbb{G}_1$, our keys and ciphertexts always achieve the same size or even improve considerably upon the other schemes.

One-use restriction. FAME has a one-use restriction described in [2, Section 4]. A common way to work around this problem is to make $\hat{\tau}$ copies of each attribute, for some $\hat{\tau}$ chosen at set-up[9]; this way, FAME can support $\hat{\tau}$-use MSPs. The downside of this transformation is that in the CP-ABE, the size of the keys grow by a factor of $\hat{\tau}$ though encryption and decryption time are not affected. Similarly, in the KP-ABE, the ciphertexts and encryption time grow by a factor of $\hat{\tau}$. We explicitly account for $\hat{\tau}$ when describing FAME in our comparison tables. For applications where τ may be large and fast decryption is paramount, we can apply the same transformation to our schemes so that decryption only requires 2-3 pairings. For this reason, the experiments and most comparison in the text consider $\hat{\tau} = \tau = 1$. A follow-up to FAME by Tomida, Kawahara and Nishimaki (TKN) [49] shows how to remove the one-use restriction using techniques from [37], paying a multiplicative factor τ in the number of pairings required for decryption, and a much larger security loss in the reduction to DLIN. The TKN scheme essentially coincides with FAME when $\tau = 1$, and for larger τ, remains at least 2-3 times less efficient than FABEO. All of our experiments are for $\tau = 1$, hence the omission of TKN.

7.2 Bit Security based on Tightness

Whereas considering multiple secret key queries in the security definition is considered standard in terms of ABE security, we additionally consider many ciphertext or challenge queries in our security proof. The two definitions are polynomially equivalent, but the non-trivial implication from one to many ciphertexts incurs a security loss linear in the number of ciphertext queries. On the contrary, if the security loss is only constant, we say that the bound is tight, as is the case for our bounds. The security loss plays an important role in choosing the system parameters of the scheme, e.g., the size of the underlying pairing group which provides a determined level of security, which is usually stated in bits. Further, we can define the success ratio of an adversary $\mathcal{A}$ by its advantage Adv and its running time t. For λ-bit security, we then require that $\mathsf{Adv}/t \leq 2^{-\lambda}$. From this value, we can then deduce whether a concrete instantiation provides the desired security level.

In Table 4, we compute the bit security of our scheme, as well as ABGW, FAME and BSW in different scenarios, that is we use different numbers of secret key and ciphertext queries. The running time t captures the *offline* time of an adversary, e.g. to perform group operations or also to evaluate a hash function (thus including random oracle queries). We assume t to be rather large, whereas secret key and ciphertext queries are considered *online* running time and therefore considerably lower. The advantage also depends on the order of the underlying group and for our comparison we assume $p = 2^{256}$. Since a discrete logarithm attack on the elliptic curve group yields a bound $O(t^2/p)$, this parameter choice is based on a security level of around 128 bit and this should be the target for the bit security of the ABE schemes as well.

We consider four different scenarios from small-scale to large-scale adversaries, based on the running time $t \in \{2^{40}, 2^{60}, 2^{80}, 2^{128}\}$.

[9] For FAME and more generally, “unbounded” ABE schemes, this parameter could also be chosen on on a per-key basis during key generation for CP-ABE, or a per-ciphertext basis during encryption for KP-ABE

	Key size		*Ciphertext size*	
Schemes	$\mathbb{G}_1$	$\mathbb{G}_2$	$\mathbb{G}_1$	$\mathbb{G}_2$
BSW	$m+1$	m	n_1	n_1+1
Waters	$m+1$	1	n_1	n_1+1
FAME	$3\hat{\tau}m+3$	3	$3n_1$	3
ABGW	-	$m+2$	$3n_1$	-
Ours	$m+1$	1	n_1	$\tau+1$
GPSW	-	n_1	m	-
FAME	$3n_1$	3	$3\hat{\tau}m$	3
ABGW	-	$2n_1$	$2m$	-
Ours	n_1	τ	m	1

Table 3: Key and ciphertext sizes of CP-ABE (top) and KP-ABE (bottom) schemes. The columns $\mathbb{G}_1$ and $\mathbb{G}_2$ denote the number of elements in the respective group (in general, $|\mathbb{G}_2| \geq 2|\mathbb{G}_1|$). m denotes the number of attributes in the set $\mathcal{S}$, n_1 and n_2 are the number of rows and columns of the MSP matrix and τ denotes the maximum number of multi-use. The experiments consider $\hat{\tau} = \tau = 1$.

Resources			*Bit Security*			
t	Q_{sk}	Q_{ct}	ABGW	FAME	BSW	Ours
2^{40}	2^{20}	2^{20}	2^{-176}	2^{-176}	2^{-196}	2^{-216}
2^{40}	2^{10}	2^{30}	2^{-176}	2^{-176}	2^{-196}	2^{-216}
2^{60}	2^{30}	2^{30}	2^{-136}	2^{-136}	2^{-166}	2^{-196}
2^{60}	2^{20}	2^{40}	2^{-136}	2^{-136}	2^{-156}	2^{-196}
2^{80}	2^{40}	2^{40}	2^{-96}	2^{-96}	2^{-136}	2^{-176}
2^{80}	2^{30}	2^{50}	2^{-96}	2^{-96}	2^{-126}	2^{-176}
2^{128}	2^{40}	2^{40}	2^{-48}	2^{-48}	2^{-88}	2^{-128}
2^{128}	2^{30}	2^{50}	2^{-48}	2^{-48}	2^{-78}	2^{-128}

Table 4: Bit security of ABE schemes depending on the adversary's running time t and number of secret key queries Q_{sk} and ciphertext queries Q_{ct}. Bit security is defined as Adv/t, where we use $p = 2^{256}$. The values coincide for CP-ABE and KP-ABE schemes. For ABGW and FAME we use $\mathsf{Adv} = O(Q_{\text{ct}}Q_{\text{sk}}t^2/p)$, for BSW we use $\mathsf{Adv} = O(Q_{\text{ct}}t^2/p)$ and for ours we use $\mathsf{Adv} = O(t^2/p)$.

For each scenario, we choose the number of secret key queries Q_{sk} and ciphertexts Q_{ct} accordingly, once for $Q_{\text{ct}} = Q_{\text{sk}}$ and once for $Q_{\text{sk}} < Q_{\text{ct}}$, since in practice an adversary may easily observe a large number of ciphertexts, rather than a large number of keys.

Evaluation. We omit Waters and GSPW here as those schemes are only selectively secure. The numbers in Table 4 are based on the security bounds stated in the corresponding papers as well as an additional hybrid argument on the number of ciphertexts as mentioned above. All schemes meet the target bound in a small to medium-scale scenario. When we increase t to 2^{80} or 2^{128}, both ABGW and FAME cannot meet the target anymore and therefore should not be used for applications in large-scale scenarios. BSW still achieves 78 resp. 88 bits, which may be sufficient for some applications. Due to the tight bound, our scheme meets the target of 128 bits in all scenarios.

8 EXTENSIONS

In this section, we briefly describe how we can extend our definition of PES-ABE to capture more schemes, e.g., ABE for deterministic finite automata (DFA).

8.1 A variant of PES-ABE

We want to capture PES-ABE schemes as in ABGW with Setup_0, Enc_0, KeyGen_0 as before, except:

$$\begin{aligned} \mathsf{mpk} &:= ([\mathbf{b}]_1, [\alpha]_T), \\ \mathsf{msk} &:= (\mathbf{b}, \alpha), \\ \mathsf{ct} &:= ([\mathbf{c}^1]_1, [\mathbf{c}^2]_1), \\ \mathsf{sk} &:= ([\mathbf{k}^1]_2, [\mathbf{k}^2]_2) \end{aligned}$$

For such schemes, we impose an additional constraint on k^1, k^2 as with ABGW, namely that $\mathbf{r} = (\mathbf{r}'\|\mathbf{r}'')$ and $k^1(\alpha, \mathbf{r}'', \mathbf{b} \otimes \mathbf{r}'), k^2(\mathbf{r}')$ (that is, we removed $\mathbf{r}', \mathbf{b} \otimes \mathbf{r}''$ from the input to k^1 and $\mathbf{r}''$ from the input to k^2).[10] This way, we can ensure that $\mathrm{span}(\mathbf{k}^1) \cap \mathrm{span}(\mathbf{k}^2) = \{0\}$, which we will need in the proof of strong symbolic security.

Strong Symbolic Security (Variant). For all $Q_{\text{ct}}, Q_{\text{sk}} \in \mathbb{N}$, $X \in \mathcal{X}^{Q_{\text{ct}}}$, $Y \in \mathcal{Y}^{Q_{\text{sk}}}$ such that $\mathsf{P}(X[i], Y[j]) = 0$ for all $i \in [Q_{\text{ct}}]$, $j \in [Q_{\text{sk}}]$, we have

$$\mathrm{span}(\tilde{\alpha} \otimes \mathbf{c}_X^2) \cap \mathrm{span}(\tilde{\alpha} \parallel (1\|\tilde{\mathbf{b}}\|\mathbf{c}_X^1\|\mathbf{c}_X^2) \otimes (1\|\mathbf{k}_Y^1\|\mathbf{k}_Y^2)) = \{0\}\,,$$

where $X, Y, \mathbf{c}_X^1, \mathbf{c}_X^2, \mathbf{k}_Y^1, \mathbf{k}_Y^2$ are as in Definition 4.2.

Claim 1. *If PES-ABE satisfies $(1,1)$-symbolic security (Definition 4.1), then it also satisfies the variant of strong symbolic security.*

The proof of this claim is deferred to the full version of the paper.

8.2 ABE for DFA

We consider the ABE scheme for DFAs in [28, equation (1)] (building on [52]). Recall that a DFA is specified by a tuple (Q, Σ, δ, F) where the state space is $[Q] := \{1, 2, \ldots, Q\}$; 1 is the unique start state; $F \subseteq [Q]$ is the set of accept states, and $\delta : [Q] \times \Sigma \to [Q]$ is the state transition function.

We provide a self-contained overview of our ABE scheme for DFA in the full version of the paper. In the following, we describe the underlying PES-ABE.

- $\underline{\mathsf{Setup}_0}$. Output $n := 3+|\Sigma|$, where we parse $\mathbf{b}$ as $(w_{\text{start}}, w_{\text{end}}, z, \{w_\sigma\}_{\sigma\in\Sigma})$.
- $\underline{\mathsf{Enc}_0(x)}$. Set $w = \ell + 1$, $w_1 = \ell + 2$, $w_2 = \ell + 1$, and output (c^1, c^2) where we parse $\mathbf{s}$ as $(s_\ell, s_0, s_1, \ldots, s_{\ell-1})$ and

$$\begin{aligned} c^1(\mathbf{s} \otimes \mathbf{b}) &:= (s_0 w_{\text{start}} \parallel \{s_{i-1}z + s_i w_{x_i}\}_{i\in[\ell]} \parallel s_\ell w_{\text{end}}), \\ c^2(\mathbf{s}) &:= (\mathbf{s}) \end{aligned}$$

[10] ABGW refers to $\mathbf{r}'$ as the non-lone variables and $\mathbf{r}''$ as the lone variables. Also, ABGW considers a more general setting for c^1, c^2 with $\mathbf{s} = (\mathbf{s}'\|\mathbf{s}'')$ and $c^2(\mathbf{s}'), c^1(\mathbf{s}'', \mathbf{b} \otimes \mathbf{s}')$. To the best of our knowledge, none of the existing ABE schemes exploit this generalization.

	Key size		*Ciphertext size*	
Schemes	$\mathbb{G}_1$	$\mathbb{G}_2$	$\mathbb{G}_1$	$\mathbb{G}_2$
Waters	-	$3Q\|\Sigma\| + 2\|F\| + 2$	$2\ell + 3$	-
Ours	-	$Q\|\Sigma\| + Q + \|F\| + 1$	$2\ell + 3$	-

Table 5: Key and ciphertext sizes of ABE schemes for DFA.

- $\underline{\mathsf{KeyGen}_0(Q, \Sigma, \delta, F)}$. Set $m = 2Q$, $m_1 = 1 + Q + Q \cdot |\Sigma| + |F|$, $m_2 = Q$, and output (k^1, k^2) where we parse $\mathbf{r} = (\mathbf{r}' \| \mathbf{r}'') := (\{r_u\}_{u \in [Q]} \parallel \{d_u\}_{u \in [Q]})$ and

$$\begin{aligned} k^1(\alpha, \mathbf{r}'', \mathbf{b} \otimes \mathbf{r}') := (&d_1 + w_{\mathrm{start}} r_1 \parallel \{-d_u + z r_u\}_{u \in [Q]} \parallel \\ &\{d_{\delta(u,\sigma)} + w_\sigma r_u\}_{u \in [Q], \sigma \in \Sigma} \parallel \\ &\{\alpha - d_u + w_{\mathrm{end}} r_u\}_{u \in F}) \\ k^2(\mathbf{r}) := (&\mathbf{r}') \end{aligned}$$

In applications, think of $\ell \gg |\Sigma|, Q$. We note that our scheme differs from Waters' scheme in that we reuse r_u for all the transitions starting from u instead of a fresh $r_{u,\sigma}$ for each (u, σ). This modification yields a smaller secret key (cf. Table 5).

ACKNOWLEDGEMENTS

Doreen Riepel was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972 and the European Union (ERC AdG REWORC - 101054911). Part of this work was done during an internship at NTT Research and a visit at UC Berkeley. We would also like to thank Kei Karasawa and his team for motivating discussions, as well as Sanjam Garg for hosting Doreen Riepel at UC Berkeley.

REFERENCES

[1] Shashank Agrawal and Melissa Chase. 2017. Attribute-based Encryption. https://github.com/sagrawal87/ABE.

[2] Shashank Agrawal and Melissa Chase. 2017. FAME: Fast Attribute-based Message Encryption. In *ACM CCS 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 665–682. https://doi.org/10.1145/3133956.3134014

[3] Shashank Agrawal and Melissa Chase. 2017. Simplifying Design and Analysis of Complex Predicate Encryption Schemes. In *EUROCRYPT 2017, Part I (LNCS)*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.), Vol. 10210. Springer, Heidelberg, 627–656. https://doi.org/10.1007/978-3-319-56620-7_22

[4] Joseph A. Akinyele, Christina Garman, Ian Miers, Matthew W. Pagano, Michael Rushanan, Matthew Green, and Aviel D. Rubin. 2013. Charm: a framework for rapidly prototyping cryptosystems. *Journal of Cryptographic Engineering* 3, 2 (June 2013), 111–128. https://doi.org/10.1007/s13389-013-0057-3

[5] Joseph A. Akinyele, Matthew W. Pagano, Matthew D. Green, Christoph U. Lehmann, Zachary N. J. Peterson, and Aviel D. Rubin. 2011. Securing electronic medical records using attribute-based encryption on mobile devices. In *SPSM'11, Proceedings of the 1st ACM Workshop Security and Privacy in Smartphones and Mobile Devices, Co-located with CCS 2011, October 17, 2011, Chicago, IL, USA*, Xuxian Jiang, Amiya Bhattacharya, Partha Dasgupta, and William Enck (Eds.). ACM, 75–86. https://doi.org/10.1145/2046614.2046628

[6] Miguel Ambrona, Gilles Barthe, Romain Gay, and Hoeteck Wee. 2017. Attribute-Based Encryption in the Generic Group Model: Automated Proofs and New Constructions. In *ACM CCS 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 647–664. https://doi.org/10.1145/3133956.3134088

[7] Nuttapong Attrapadung. 2014. Dual System Encryption via Doubly Selective Security: Framework, Fully Secure Functional Encryption for Regular Languages, and More. In *EUROCRYPT 2014 (LNCS)*, Phong Q. Nguyen and Elisabeth Oswald (Eds.), Vol. 8441. Springer, Heidelberg, 557–577. https://doi.org/10.1007/978-3-642-55220-5_31

[8] Nuttapong Attrapadung. 2016. Dual System Encryption Framework in Prime-Order Groups via Computational Pair Encodings. In *ASIACRYPT 2016, Part II (LNCS)*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.), Vol. 10032. Springer, Heidelberg, 591–623. https://doi.org/10.1007/978-3-662-53890-6_20

[9] Randy Baden, Adam Bender, Neil Spring, Bobby Bhattacharjee, and Daniel Starin. 2009. Persona: An Online Social Network with User-defined Privacy. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication (SIGCOMM '09)*. ACM, New York, NY, USA, 135–146. https://doi.org/10.1145/1592568.1592585

[10] Balthazar Bauer, Georg Fuchsbauer, and Antoine Plouviez. 2021. The One-More Discrete Logarithm Assumption in the Generic Group Model. In *Advances in Cryptology – ASIACRYPT 2021*, Mehdi Tibouchi and Huaxiong Wang (Eds.). Springer International Publishing, Cham, 587–617.

[11] Mihir Bellare and Phillip Rogaway. 1993. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *ACM CCS 93*, Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby (Eds.). ACM Press, 62–73. https://doi.org/10.1145/168588.168596

[12] John Bethencourt, Amit Sahai, and Brent Waters. 2007. Ciphertext-Policy Attribute-Based Encryption. In *2007 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 321–334. https://doi.org/10.1109/SP.2007.11

[13] Olivier Blazy, Eike Kiltz, and Jiaxin Pan. 2014. (Hierarchical) Identity-Based Encryption from Affine Message Authentication. In *CRYPTO 2014, Part I (LNCS)*, Juan A. Garay and Rosario Gennaro (Eds.), Vol. 8616. Springer, Heidelberg, 408–425. https://doi.org/10.1007/978-3-662-44371-2_23

[14] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. 2005. Hierarchical Identity Based Encryption with Constant Size Ciphertext. In *EUROCRYPT 2005 (LNCS)*, Ronald Cramer (Ed.), Vol. 3494. Springer, Heidelberg, 440–456. https://doi.org/10.1007/11426639_26

[15] Melissa Chase, Mary Maller, and Sarah Meiklejohn. 2016. Déjà Q All Over Again: Tighter and Broader Reductions of q-Type Assumptions. In *ASIACRYPT 2016, Part II (LNCS)*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.), Vol. 10032. Springer, Heidelberg, 655–681. https://doi.org/10.1007/978-3-662-53890-6_22

[16] Jie Chen, Romain Gay, and Hoeteck Wee. 2015. Improved Dual System ABE in Prime-Order Groups via Predicate Encodings. In *EUROCRYPT 2015, Part II (LNCS)*, Elisabeth Oswald and Marc Fischlin (Eds.), Vol. 9057. Springer, Heidelberg, 595–624. https://doi.org/10.1007/978-3-662-46803-6_20

[17] Jie Chen, Junqing Gong, and Jian Weng. 2017. Tightly Secure IBE Under Constant-Size Master Public Key. In *PKC 2017, Part I (LNCS)*, Serge Fehr (Ed.), Vol. 10174. Springer, Heidelberg, 207–231. https://doi.org/10.1007/978-3-662-54365-8_9

[18] Jie Chen and Hoeteck Wee. 2013. Fully, (Almost) Tightly Secure IBE and Dual System Groups. In *CRYPTO 2013, Part II (LNCS)*, Ran Canetti and Juan A. Garay (Eds.), Vol. 8043. Springer, Heidelberg, 435–460. https://doi.org/10.1007/978-3-642-40084-1_25

[19] Jung Hee Cheon. 2006. Security Analysis of the Strong Diffie-Hellman Problem. In *EUROCRYPT 2006 (LNCS)*, Serge Vaudenay (Ed.), Vol. 4004. Springer, Heidelberg, 1–11. https://doi.org/10.1007/11761679_1

[20] Arush Chhatrapati, Susan Hohenberger, James Trombo, and Satyanarayana Vusirikala. 2022. A Performance Evaluation of Pairing-Based Broadcast Encryption Systems. In *Applied Cryptography and Network Security*. Springer International Publishing.

[21] Antonio de la Piedra, Marloes Venema, and Greg Alpár. 2022. ABE Squared: Accurately Benchmarking Efficiency of Attribute-Based Encryption. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2022, 2 (2022), 192–239. https://doi.org/10.46586/tches.v2022.i2.192-239

[22] Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. 2021. More Efficient Digital Signatures with Tight Multi-user Security. In *PKC 2021, Part II (LNCS)*, Juan Garay (Ed.), Vol. 12711. Springer, Heidelberg, 1–31. https://doi.org/10.1007/978-3-030-75248-4_1

[23] Denis Diemert and Tibor Jager. 2021. On the Tight Security of TLS 1.3: Theoretically Sound Cryptographic Parameters for Real-World Deployments. *Journal of Cryptology* 34, 3 (July 2021), 30. https://doi.org/10.1007/s00145-021-09388-x

[24] Armando Faz-Hernández, Sam Scott, Nick Sullivan, Riad S. Wahby, and Christopher A. Wood. 2022. Hashing to Elliptic Curves. https://datatracker.ietf.org/doc/draft-irtf-cfrg-hash-to-curve/.

[25] Craig Gentry. 2006. Practical Identity-Based Encryption Without Random Oracles. In *EUROCRYPT 2006 (LNCS)*, Serge Vaudenay (Ed.), Vol. 4004. Springer, Heidelberg, 445–464. https://doi.org/10.1007/11761679_27

[26] Kristian Gjøsteen and Tibor Jager. 2018. Practical and Tightly-Secure Digital Signatures and Authenticated Key Exchange. In *CRYPTO 2018, Part II (LNCS)*, Hovav Shacham and Alexandra Boldyreva (Eds.), Vol. 10992. Springer, Heidelberg, 95–125. https://doi.org/10.1007/978-3-319-96881-0_4

[27] Junqing Gong, Xiaolei Dong, Jie Chen, and Zhenfu Cao. 2016. Efficient IBE with Tight Reduction to Standard Assumption in the Multi-challenge Setting. In *ASIACRYPT 2016, Part II (LNCS)*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.), Vol. 10032. Springer, Heidelberg, 624–654. https://doi.org/10.1007/978-3-662-53890-6_21

[28] Junqing Gong, Brent Waters, and Hoeteck Wee. 2019. ABE for DFA from k-Lin. In *CRYPTO 2019, Part II (LNCS)*, Alexandra Boldyreva and Daniele Micciancio (Eds.), Vol. 11693. Springer, Heidelberg, 732–764. https://doi.org/10.1007/978-3-030-26951-7_25

[29] Junqing Gong and Hoeteck Wee. 2020. Adaptively Secure ABE for DFA from k-Lin and More. In *EUROCRYPT 2020, Part III (LNCS)*, Anne Canteaut and Yuval Ishai (Eds.), Vol. 12107. Springer, Heidelberg, 278–308. https://doi.org/10.1007/978-3-030-45727-3_10

[30] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. 2006. Attribute-Based Encryption for Fine-Grained Access Control of Encrypted Data. In *ACM CCS 2006*, Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati (Eds.). ACM Press, 89–98. https://doi.org/10.1145/1180405.1180418 Available as Cryptology ePrint Archive Report 2006/309.

[31] Jens Groth. 2016. On the Size of Pairing-Based Non-interactive Arguments. In *EUROCRYPT 2016, Part II (LNCS)*, Marc Fischlin and Jean-Sébastien Coron (Eds.), Vol. 9666. Springer, Heidelberg, 305–326. https://doi.org/10.1007/978-3-662-49896-5_11

[32] Jens Groth and Victor Shoup. 2021. On the security of ECDSA with additive key derivation and presignatures. Cryptology ePrint Archive, Report 2021/1330. https://eprint.iacr.org/2021/1330.

[33] Viet Tung Hoang, Stefano Tessaro, and Aishwarya Thiruvengadam. 2018. The Multi-user Security of GCM, Revisited: Tight Bounds for Nonce Randomization. In *ACM CCS 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM Press, 1429–1440. https://doi.org/10.1145/3243734.3243816

[34] Dennis Hofheinz, Jessica Koch, and Christoph Striecks. 2015. Identity-Based Encryption with (Almost) Tight Security in the Multi-instance, Multi-ciphertext Setting. In *PKC 2015 (LNCS)*, Jonathan Katz (Ed.), Vol. 9020. Springer, Heidelberg, 799–822. https://doi.org/10.1007/978-3-662-46447-2_36

[35] Mihaela Ion, Jianqing Zhang, and Eve M. Schooler. 2013. Toward content-centric privacy in ICN: attribute-based encryption and routing. In *ICN'13, Proceedings of the 3rd, 2013 ACM SIGCOMM Workshop on Information-Centric Networking, August 12, 2013, Hong Kong, China*, Börje Ohlman, George C. Polyzos, and Lixia Zhang (Eds.). ACM, 39–40. https://doi.org/10.1145/2491224.2491237

[36] Taechan Kim and Razvan Barbulescu. 2016. Extended Tower Number Field Sieve: A New Complexity for the Medium Prime Case. In *CRYPTO 2016, Part I (LNCS)*, Matthew Robshaw and Jonathan Katz (Eds.), Vol. 9814. Springer, Heidelberg, 543–571. https://doi.org/10.1007/978-3-662-53018-4_20

[37] Lucas Kowalczyk and Hoeteck Wee. 2019. Compact Adaptively Secure ABE for NC^1 from k-Lin. In *EUROCRYPT 2019, Part I (LNCS)*, Yuval Ishai and Vincent Rijmen (Eds.), Vol. 11476. Springer, Heidelberg, 3–33. https://doi.org/10.1007/978-3-030-17653-2_1

[38] Roman Langrehr and Jiaxin Pan. 2020. Hierarchical Identity-Based Encryption with Tight Multi-challenge Security. In *PKC 2020, Part I (LNCS)*, Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas (Eds.), Vol. 12110. Springer, Heidelberg, 153–183. https://doi.org/10.1007/978-3-030-45374-9_6

[39] Allison B. Lewko and Brent Waters. 2011. Unbounded HIBE and Attribute-Based Encryption. In *EUROCRYPT 2011 (LNCS)*, Kenneth G. Paterson (Ed.), Vol. 6632. Springer, Heidelberg, 547–567. https://doi.org/10.1007/978-3-642-20465-4_30

[40] Allison B. Lewko and Brent Waters. 2012. New Proof Methods for Attribute-Based Encryption: Achieving Full Security through Selective Techniques. In *CRYPTO 2012 (LNCS)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, 180–198. https://doi.org/10.1007/978-3-642-32009-5_12

[41] Jin Li, Man Ho Au, Willy Susilo, Dongqing Xie, and Kui Ren. 2010. Attribute-based signature and its applications. In *ASIACCS 10*, Dengguo Feng, David A. Basin, and Peng Liu (Eds.). ACM Press, 60–69.

[42] Huijia Lin and Ji Luo. 2020. Compact Adaptively Secure ABE from k-Lin: Beyond NC^1 and Towards NL. In *EUROCRYPT 2020, Part III (LNCS)*, Anne Canteaut and Yuval Ishai (Eds.), Vol. 12107. Springer, Heidelberg, 247–277. https://doi.org/10.1007/978-3-030-45727-3_9

[43] Ueli M. Maurer. 2005. Abstract Models of Computation in Cryptography (Invited Paper). In *10th IMA International Conference on Cryptography and Coding (LNCS)*, Nigel P. Smart (Ed.), Vol. 3796. Springer, Heidelberg, 1–12.

[44] Doreen Riepel and Hoeteck Wee. 2022. https://github.com/DoreenRiepel/FABEO.

[45] Amit Sahai and Brent R. Waters. 2005. Fuzzy Identity-Based Encryption. In *EUROCRYPT 2005 (LNCS)*, Ronald Cramer (Ed.), Vol. 3494. Springer, Heidelberg, 457–473. https://doi.org/10.1007/11426639_27

[46] Yumi Sakemi, Tetsutaro Kobayashi, Tsunekazu Saito, and Riad Wahby. 2021. Pairing-Friendly Curves. Internet-Draft https://datatracker.ietf.org/doc/draft-irtf-cfrg-pairing-friendly-curves/.

[47] Victor Shoup. 1997. Lower Bounds for Discrete Logarithms and Related Problems. In *EUROCRYPT'97 (LNCS)*, Walter Fumy (Ed.), Vol. 1233. Springer, Heidelberg, 256–266. https://doi.org/10.1007/3-540-69053-0_18

[48] Nick Sullivan. 2017. Geo Key Manager: How It Works. https://blog.cloudflare.com/geo-key-manager-how-it-works/.

[49] Junichi Tomida, Yuto Kawahara, and Ryo Nishimaki. 2020. Fast, Compact, and Expressive Attribute-Based Encryption. In *PKC 2020, Part I (LNCS)*, Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas (Eds.), Vol. 12110. Springer, Heidelberg, 3–33. https://doi.org/10.1007/978-3-030-45374-9_1

[50] Brent Waters. 2009. Dual System Encryption: Realizing Fully Secure IBE and HIBE under Simple Assumptions. In *CRYPTO 2009 (LNCS)*, Shai Halevi (Ed.), Vol. 5677. Springer, Heidelberg, 619–636. https://doi.org/10.1007/978-3-642-03356-8_36

[51] Brent Waters. 2011. Ciphertext-Policy Attribute-Based Encryption: An Expressive, Efficient, and Provably Secure Realization. In *PKC 2011 (LNCS)*, Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi (Eds.), Vol. 6571. Springer, Heidelberg, 53–70. https://doi.org/10.1007/978-3-642-19379-8_4

[52] Brent Waters. 2012. Functional Encryption for Regular Languages. In *CRYPTO 2012 (LNCS)*, Reihaneh Safavi-Naini and Ran Canetti (Eds.), Vol. 7417. Springer, Heidelberg, 218–235. https://doi.org/10.1007/978-3-642-32009-5_14

[53] Hoeteck Wee. 2014. Dual System Encryption via Predicate Encodings. In *TCC 2014 (LNCS)*, Yehuda Lindell (Ed.), Vol. 8349. Springer, Heidelberg, 616–637. https://doi.org/10.1007/978-3-642-54242-8_26