Improved Dual System ABE in Prime-Order Groups via Predicate Encodings

Jie Chen[1,⋆], Romain Gay[2,⋆⋆], and Hoeteck Wee[2,⋆⋆⋆]

[1] East China Normal University, Shanghai, China
s080001@e.ntu.edu.sg

[2] ENS, Paris, France
rgay,wee@di.ens.fr

Abstract. We present a modular framework for the design of efficient adaptively secure attribute-based encryption (ABE) schemes for a large class of predicates under the standard k-Lin assumption in prime-order groups; this is the first uniform treatment of dual system ABE across different predicates and across both composite and prime-order groups. Via this framework, we obtain concrete efficiency improvements for several ABE schemes. Our framework has three novel components over prior works: (i) new techniques for simulating composite-order groups in prime-order ones, (ii) a refinement of prior encodings framework for dual system ABE in composite-order groups, (iii) an extension to weakly attribute-hiding predicate encryption (which includes anonymous identity-based encryption as a special case).

1 Introduction

Attribute-based encryption (ABE) [27, 15] is a new paradigm for public-key encryption that enables fine-grained access control for encrypted data. In ABE, ciphertexts are associated with descriptive values x in addition to a plaintext, secret keys are associated with values y, and a secret key decrypts the ciphertext if and only if $\mathsf{P}(x, y) = 1$ for some boolean predicate P. Here, y together with P may express an arbitrarily complex access policy, which is in stark contrast to traditional public-key encryption, where access is all or nothing. The simplest example of ABE is that of identity-based encryption (IBE) [28, 5, 12] where P corresponds to equality. The security requirement for ABE enforces resilience

[⋆] Shanghai Key Laboratory of Multidimensional Information Processing and Shanghai Key Lab of Trustworthy Computing. Supported by the National Natural Science Foundation of China (Grant Nos. 61472142, 61321064, 61172085), Science and Technology Commission of Shanghai Municipality (Grant Nos. 14YF1404200, 13JC1403500).

[⋆⋆] Supported in part by ANR-14-CE28-0003 (Project EnBiD).

[⋆⋆⋆] CNRS, INRIA and Columbia University. Supported in part by ANR-14-CE28-0003 (Project EnBiD), NSF Award CNS-1445424, the Alexander von Humboldt Foundation and a Google Faculty Research Award.

to collusion attacks, namely any group of users holding secret keys for different values learns nothing about the plaintext if none of them is individually authorized to decrypt the ciphertext. This should hold even if the adversary *adaptively* decides which secret keys to ask for.

ABE in Prime-Order Groups. The goal of this work is to obtain efficient adaptively secure ABE for a large class of predicates. We now have a fairly good understanding of how to obtain such schemes in composite-order bilinear groups, thanks to Waters' powerful *dual system encryption methodology* [30] and recent unifying frameworks in [2, 31] for the design of dual system ABE schemes. However, these latter frameworks only work in composite-order bilinear groups, for which group operations and especially pairing computations are prohibitively slow. In practice, prime-order bilinear groups are preferable [16] as they admit not only more efficient but also more compact instantiations. To mitigate the gap between ease of theoretical design and practical efficiency, a series of works studied techniques for converting cryptosystems relying on composite-order groups to cryptosystems based on prime-order groups [23, 24, 14, 20, 11, 10], largely in the context of dual system ABE. In addition, we have direct constructions of dual system prime-order hierarchical identity-based encryption (HIBE) schemes in [18, 3] that bypass a conversion from composite-order groups, but the techniques in these constructions do not seem to naturally extend beyond (H)IBE. Furthermore, the prior constructions rely on fairly distinct techniques, and efficiency improvements in one construction do not necessarily translate to a different construction and a different predicate. In short, prior works fall short of providing a unifying and modular framework for the design of efficient dual system ABE schemes in prime-order groups that work for a large class of predicates (c.f. Fig. 1).

1.1 Our contributions

We present a modular framework for the design of efficient dual system ABE schemes for a large class of predicates under the standard k-Lin assumption in prime-order groups; this is the first uniform treatment of dual system ABE across different predicates *and* across both composite and prime-order groups. Via this framework, we obtain concrete efficiency improvements for several ABE schemes. Our framework has three novel components over prior works: (i) new techniques for simulating composite-order groups in prime-order ones, (ii) a refinement of the encodings framework for dual system ABE for composite-order groups in [2, 31], (iii) an extension to weakly attribute-hiding predicate encryption [19, 6] (which includes anonymous IBE as a special case). The last two components answer the open problems left in [2, 31].

New techniques for simulating composite-order groups. The starting point of our construction is simply a simpler choice of basis. Fix a bilinear group (G_1, G_2, G_T) with $e : G_1 \times G_2 \to G_T$ of prime order p. We pick random

	compact HIBE	boolean formula	k-Lin	anonymous IBE	weakly-AH ZIPE
DPVS [23, 24, 20, 11]	no	yes	yes	yes	yes
sparse DPVS [26]	yes	?	?	yes	yes
QANIZK [18]	yes	?	yes	yes	?
dual system groups [10]	yes	?	yes	?	?
MAC-to-(H)IBE [3]	yes	?	yes	yes	?
this work	yes	yes	yes	yes	yes

Fig. 1. Summary of previous approaches for building efficient dual system (H)IBE and ABE in prime-order groups. The first column refers to HIBE with constant-size ciphertexts; the second refers to KP/CP-ABE for boolean formula. The third column refers to instantiations from the general k-Lin assumption. The last two columns address extensions to stronger notions of security like anonymity and weakly attribute-hiding (AH) inner product encryption (ZIPE). Additional discussion is provided in Section 1.2.

matrices $(\mathbf{A}, \mathbf{B}) \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}$, along with random vectors $\mathbf{a}^{\perp}, \mathbf{b}^{\perp} \in \mathbb{Z}_p^{k+1}$ so that $\mathbf{a}^{\perp\top}\mathbf{A} = \mathbf{b}^{\perp\top}\mathbf{B} = \mathbf{0}$, and we assume $\mathbf{a}^{\perp\top}\mathbf{b}^{\perp} \neq 0$. Observe that[1]

$$([\mathbf{A}]_1, [\mathbf{b}^{\perp}]_1) := (g_1^{\mathbf{A}}, g_1^{\mathbf{b}^{\perp}}) \in G_1^{(k+1)\times k} \times G_1^{k+1}$$

forms a basis for G_1^{k+1}. Similarly,

$$([\mathbf{B}]_2, [\mathbf{a}^{\perp}]_2) := (g_2^{\mathbf{B}}, g_2^{\mathbf{a}^{\perp}}) \in G_2^{(k+1)\times k} \times G_2^{k+1}$$

forms a basis for G_2^{k+1}. In the context of dual system encryption, we use $[\mathbf{A}]_1$ as a basis for normal components in the ciphertext space, and $[\mathbf{b}^{\perp}]_1$ as a basis for semi-functional components. Similarly, we use $[\mathbf{B}]_2$ as a basis for normal components in the secret key space, and $[\mathbf{a}^{\perp}]_2$ as a basis for semi-functional components. Indistinguishability for elements with and without random semi-functional components follow readily from the k-Lin assumption. Moreover, we have an orthogonality property given by $\mathbf{a}^{\perp\top}\mathbf{A} = \mathbf{b}^{\perp\top}\mathbf{B} = \mathbf{0}$, which tells us that the normal and semi-functional components in different spaces cancel out.

We can then randomize this basis by choosing $\mathbf{W} \in \mathbb{Z}_p^{(k+1)\times(k+1)}$ uniformly at random and using $([\mathbf{W}^{\top}\mathbf{A}]_1, [\mathbf{W}^{\top}\mathbf{b}^{\perp}]_1)$ for G_1^{k+1} and $([\mathbf{W}\mathbf{B}]_2, [\mathbf{W}\mathbf{a}^{\perp}]_2)$ for G_2^{k+1}. For decryption correctness, we will exploit the following "associative" property when the new basis interacts with the original one, namely:

$$e([\mathbf{A}]_1, [\mathbf{W}\mathbf{B}]_2) = e([\mathbf{W}^{\top}\mathbf{A}]_1, [\mathbf{B}]_2) \tag{1}$$

where we define the pairing operation on matrices via

$$e([\mathbf{M}]_1, [\mathbf{M}']_2) := e(g_1, g_2)^{\mathbf{M}^{\top}\mathbf{M}'}.$$

[1] Following [13], we use the implicit representation notation for group elements, as explained in Section 4.1.

Observe that $\mathbf{W}$ has one unit of residual entropy given $(\mathbf{W}^{\top}\mathbf{A}, \mathbf{W}\mathbf{B})$. This will be crucial for carrying out the information-theoretic argument in the proof of ABE security via the dual system encryption methodology [30, 2, 31].

We note that prior transformations in prime-order groups in [14, 23, 24, 20] try to simulate all of the structure in composite-order groups (e.g. orthogonality). We simulate less structure (associativity, c.f. Eqn. (1)), thus leading to better concrete efficiency. However, when combined with the existing encodings framework for dual system ABE schemes in composite-order groups, we cannot even guarantee ABE decryption correctness. We compensate for less structure while simulating composite-order groups by imposing more structure to the encodings, which we can achieve without increasing the size of the encodings. We will exploit the additional structure in the encodings for correctness and for security. We now proceed to describe our encodings framework for ABE.

Modular approach for ABE. We begin with the observation that the prior composite-order ABE schemes in [31, 2] (generalizing [21, 22]) may be modified so that master public key, secret key and ciphertext are of the form:

$$\begin{aligned} \mathsf{mpk} &:= \big(\, g_1,\ g_1^{\mathbf{w}},\ e(g_1,g_1)^{\alpha} \,\big) \\ \mathsf{sk}_y &:= \big(\, g_1^{r},\ g_1^{\mathsf{kE}(y,\alpha)+r\cdot\mathsf{rE}(y,\mathbf{w})} \,\big) \\ \mathsf{ct}_x &:= \big(\, g_1^{s},\ g_1^{s\cdot\mathsf{sE}(x,\mathbf{w})},\ e(g_1,g_1)^{\alpha s}\cdot m \big) \end{aligned} \tag{2}$$

Here, g_1 is a generator of order p_1 where the underlying composite group order is the product of three primes p_1, p_2, p_3 (for simplicity we consider the case of a symmetric bilinear group); $\mathbf{w}$ is a vector of length n; and $\mathsf{kE}, \mathsf{rE}, \mathsf{sE}$ are a triple of deterministic "encoding" functions that depend on the underlying predicate P (we refer to these functions as key encoding, receiver encoding and sender encoding respectively.) Syntactically, this is already a refinement of the prior frameworks in [31, 2] which associates a single function with sk_y given by

$$(y, \alpha, \mathbf{w}, r) \mapsto \big(r,\ \mathsf{kE}(y,\alpha) + r\cdot\mathsf{rE}(y,\mathbf{w})\big) \tag{3}$$

in the exponent. The prior frameworks allow for instance for kE to be randomized. With the refinement in place, we can now specify the restricted α-reconstruction property used for correctness:

> **(restricted α-reconstruction.)** For every x, y for which $\mathsf{P}(x,y) = 1$, there is a linear map L_{xy} such that for all α, r,
>
> $$L_{xy}\Big(\ \mathsf{kE}(y,\alpha) + r\cdot\mathsf{rE}(y,\mathbf{w}),\ r\cdot\mathsf{sE}(x,\mathbf{w})\ \Big) = \alpha.$$

This means that we can recover $e(g_1,g_1)^{\alpha s}$ given

$$e(g_1^{s}, g_1^{\mathsf{kE}(y,\alpha)+r\cdot\mathsf{rE}(y,\mathbf{w})}) \quad \text{and} \quad e(g_1^{s\cdot\mathsf{sE}(x,\mathbf{w})}, g_1^{r}),$$

upon which we can decrypt the ciphertext. Observe that we only need to pair the first component g_1^s of ct_x with the second component of sk_y and the second component of ct_x with the first component g_1^r of sk_y. Correctness now relies on a so-called *associativity property* [10], namely that for all i and all w_i:

$$e(g_1^s, g_1^{w_i r}) = e(g_1^{w_i s}, g_1^r) \tag{4}$$

To translate the scheme to prime-order groups, we carry out the following substitution:

$$\begin{aligned} w_i &\mapsto \mathbf{W}_i \in \mathbb{Z}_p^{(k+1)\times(k+1)}, & s &\mapsto \mathbf{s} \in \mathbb{Z}_p^k, & r &\mapsto \mathbf{r} \in \mathbb{Z}_p^k \\ & & g_1^s &\mapsto [\mathbf{A}\mathbf{s}]_1, & g_1^r &\mapsto [\mathbf{B}\mathbf{r}]_2 \\ & & g_1^{w_i s} &\mapsto [\mathbf{W}_i^\top \mathbf{A}\mathbf{s}]_1, & g_1^{w_i r} &\mapsto [\mathbf{W}_i \mathbf{B}\mathbf{r}]_2 \end{aligned}$$

Using (1), we have

$$e([\mathbf{A}\mathbf{s}]_1, [\mathbf{W}_i\mathbf{B}\mathbf{r}]_2) = e([\mathbf{W}_i^\top \mathbf{A}\mathbf{s}]_1, [\mathbf{B}\mathbf{r}]_2)$$

which is exactly what we used in composite-order groups in (4). In fact, a stronger "pairwise associativity" property holds in composite-order groups, namely for all i, j and all w_i, w_j:

$$e(g_1^{w_j s}, g_1^{w_i r}) = e(g_1^{w_i s}, g_1^{w_j r})$$

which is not satisfied by our prime-order techniques since $\mathbf{W}_i$ and $\mathbf{W}_j$ do not commute. Restricted α-reconstruction means that we do not need to pair $g_1^{w_j s}$ with $g_1^{w_i r}$ during decryption, and thus the associativity property already suffices for decryption correctness. For maximal modularity, we describe our compiler using the framework of dual system groups introduced in [10], which allows us to simultaneously capture prime-order and composite-order groups.

Next, we specify the privacy property which we use in the proof of ABE security:

> **(α-privacy.)** For every x, y for which $\mathsf{P}(x, y) = 0$, α is perfectly hidden given
>
> $$\mathsf{sE}(x, \mathbf{w}),\ \ \mathsf{kE}(y, \alpha) + \mathsf{rE}(y, \mathbf{w})$$
>
> where $\mathbf{w} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^n$.

We stress that the privacy requirement only needs to hold in a private-key setting where the adversary does not see $\mathbf{w}$ and in a one-time setting where the adversary only gets a single copy of $\mathsf{sE}(x, \mathbf{w}), \mathsf{kE}(y, \alpha) + \mathsf{rE}(y, \mathbf{w})$. As pointed out in [31], the dual system encryption methodology can be used to boost security in a private-key, one-time, non-adaptive setting as given by α-privacy to a full-fledged public-key, many-time, adaptive setting as is required for ABE security. One novelty in this work over [31, 2] lies in carrying this out over prime-order bilinear groups. In the proof, we exploit the fact that the key sk_y leaks no information about $\mathbf{w}$ when $r = 0$ (c.f. Eqn. (2)). This way, we can ensure that in each step in the proof of security, at most one secret key leaks information about $\mathbf{w}$ in the semi-functional space. This is important since α-privacy only holds when $\mathbf{w}$ is used once. We also introduce new attribute-hiding privacy requirements for encodings in this work (c.f. Section 7.2).

New encodings. For many predicates, the prior encodings in [31, 2] satisfy the new refinement trivially. In addition, we introduce a number of new encodings:

- For KP-ABE for boolean formula, the prior encoding corresponding to the secret key in [31, 2] is given by

  $$(r, \alpha_1 + rw_1, \ldots, \alpha_\ell + rw_\ell)$$

  where $(\alpha_1, \ldots, \alpha_\ell)$ are random shares of α using a linear secret-sharing scheme and fresh randomness for each secret key. This does not satisfy the syntactic refinement captured in Eqn. (3). In our scheme, we use

  $$(r, \alpha'_1 + r(w_1 + v_1), \ldots, \alpha'_\ell + r(w_\ell + v_\ell))$$

  where $(\alpha'_1, \ldots, \alpha'_\ell)$ are *deterministically* derived from α using the secret-sharing scheme with randomness fixed to 0 and $(v_1, \ldots, v_\ell)$ are *random* shares of 0. In the ensuing KP-ABE scheme, we use the same $v_1, \ldots, v_\ell$ across all secret keys whereas prior constructions use fresh randomness for secret-sharing for each key. In addition, we obtain an analogous construction for CP-ABE. Here, we avoid having to consider randomized sender encodings as in [31, 2]. The final encodings have the same sizes as the prior ones, while satisfying the new refinement requirement. Moreover, by using associativity (c.f. Eqn. (4)), we reduce the number of pairings for the decryption to a constant and avoid exponentiations in the target group at the cost of cheaper exponentiations in the source groups.
- We extend the encodings for KP-ABE and CP-ABE to arithmetic branching programs, based on the selectively secure KP-ABE for arithmetic branching programs in [17]. Combined with our generic framework, we obtain the first adaptively secure KP-ABE and CP-ABE for arithmetic branching programs.
- We also present a new encoding for broadcast encryption with n users where both the receiver and sender encoding have sublinear $O(\sqrt{n})$ length and a simple encoding for large universe fuzzy IBE.

Achieving weak attribute-hiding. In a weakly attribute-hiding scheme, we need to guarantee the privacy of the ciphertext attribute x against collusions that are not authorized to decrypt the challenge ciphertext. To achieve this property, we require additional properties from the underlying encoding and the underlying group structure (extending ideas from [25, 1, 3]). We use the fact that for any vector $\mathbf{c} \in \mathbb{Z}_p^{k+1}$ outside the span of $\mathbf{A}$, the vector $\mathbf{W}^\top \mathbf{c}$ is uniformly random given $\mathbf{W}^\top \mathbf{A}$, where $\mathbf{W}$ is a uniformly random matrix. We can then use $\mathbf{W}^\top \mathbf{c}$ to information-theoretically blind the attribute in the challenge ciphertext. For this to work, we need to make sure that the semi-functional secret keys do not leak any *additional* information about $\mathbf{WB}$.

New ABE schemes. We describe several concrete new ABE schemes obtained via our new framework (c.f. Fig. 2). Specifically, we obtain:

functionality	improvements
KP-ABE boolean formula	50% savings in SK size, faster Dec
CP-ABE boolean formula	50% savings in CT size, faster Dec
KP-ABE arithmetic formula	first adaptively secure scheme
CP-ABE arithmetic formula	first adaptively secure scheme
NIPE	25-50% savings in SK and CT size and in Dec time
weakly attribute-hiding ZIPE	25% savings in SK and CT size and in Dec time

Fig. 2. Summary of efficiency improvements in our new ABE schemes. Here, SK, CT, and Dec stand for secret key, ciphertext, and decryption respectively.

- ABE schemes for the inner product and non-zero inner product predicates with a 25% improvement in secret key and ciphertext sizes and decryption time, improving upon previous constructions in [26];
- a key-policy ABE scheme for boolean formula with a 50% improvement in secret key size and faster decryption and an analogous result for ciphertext-policy ABE, improving upon previous constructions in [25, 20];
- the first adaptively secure key-policy and ciphertext-policy ABE schemes for arithmetic formula and branching programs without an exponential security loss, improving upon previous constructions in [17, 8].

Along the way, we also generalize several previous constructions for $k = 2$ to general k with $k = 1$ being particularly relevant for practical efficiency. More generally, the parameters of our schemes under k-Lin is $k + 1$ times those of the best composite-order schemes based on subgroup assumptions: this achieves a "seemingly best-possible" composite-to-prime-order transformation where each composite element is simulated using $k + 1$ prime-order elements.

Finally, our prime-order ABE schemes are simpler to describe than prior schemes as they share the same structure as existing composite-order schemes. In particular, we obtain the following anonymous IBE scheme:

$$\begin{aligned}
\mathsf{mpk} &= [\mathbf{A}, \mathbf{W}_0^{\top}\mathbf{A}, \mathbf{W}_1^{\top}\mathbf{A}]_1, [\mathbf{k}^{\top}\mathbf{A}]_T \\
\mathsf{sk}_{\mathsf{id}} &= [\mathbf{Br}, \mathbf{k} + (\mathbf{W}_0 + \mathsf{id}\cdot\mathbf{W}_1)\mathbf{Br}]_2 \ \in G_2^{2(k+1)} \\
\mathsf{ct}_{\mathsf{id}} &= [\mathbf{As}, (\mathbf{W}_0 + \mathsf{id}\cdot\mathbf{W}_1)^{\top}\mathbf{As}]_1, [\mathbf{k}^{\top}\mathbf{As}]_T \cdot m \ \in G_1^{2(k+1)} \times G_T
\end{aligned}$$

where $\mathbf{A}, \mathbf{B} \in \mathbb{Z}_p^{(k+1)\times k}$, $\mathbf{W}_0, \mathbf{W}_1 \in \mathbb{Z}_p^{(k+1)\times(k+1)}$, $\mathbf{s}, \mathbf{r} \in \mathbb{Z}_p^k$, $\mathbf{k} \in \mathbb{Z}_p^{k+1}$. This scheme extends naturally to a non-anonymous BBG-style compact HIBE [7] (this is not the case for the prime-order IBE schemes in [20, 11]).

1.2 Discussion

Comparison with prior works. A summary of the prior approaches for obtaining efficient adaptively secure efficient dual system (H)IBE and ABE is presented in Fig. 1. The most general technique we have for simulating composite-order groups in prime-order ones are those based on "dual pairing vector spaces" (DPVS) [23, 24, 20, 11]. However, these techniques do not preserve the asymptotic efficiency of the underlying schemes; in particular, applying them to the composite-order compact HIBE schemes in [21] blows up the ciphertext size from constant to linear. The sparse DPVS technique [26, 29] uses subgroups of sparse matrices with mostly zero entries to overcome this limitation; however, they substantially limit the generality of the DPVS technique: the structure of these matrices now depend on the predicate and the composite-order scheme (to preserve efficiency), and the analysis for correctness, efficiency and security are more involved. The constructions in [10] fail to extend to boolean formula due to the need for additional randomness for secret-sharing, and also do not extend to yield anonymous IBE. The direct constructions in [18, 3] that bypass a conversion from composite-order groups do not seem to naturally extend beyond (H)IBE: the former uses tag-based languages where tags correspond to identities, and the latter relies on the notion of message authentication codes where messages correspond to identities. In particular, we do not know analogues of these constructions for either the inner product predicate or CP/KP-ABE for boolean formula.

As noted earlier, another novel contribution in this work over prior unifying frameworks in [2, 31] (generalizing [21, 22]) for composite-order groups lies in realizing the weakly-attribute guarantee. This is particularly challenging in composite-order groups for two reasons: (i) there is an explicit anonymity attack on the Lewko-Waters IBE [21] in composite-order group and (ii) the attribute in the semi-functional ciphertext is leaked in the G_{p_1}-component. Interestingly, we are still able to show that our prime-order analog of the Lewko-Waters IBE is anonymous.

Organization. We recall the definition of an attribute-based encryption scheme in Section 2. We recall the notion of dual system groups in Section 3 and describe our instantiations in Section 4. We describe our notion of predicate encodings in Section 5. We present our generic ABE construction in Section 6. We handle weakly attribute-hiding predicate encryption in Section 7. We defer instantiations of predicate encodings and all other details to the full version of this paper.

2 Preliminaries

Notation. We denote by $s \leftarrow_{\mathrm{R}} S$ the fact that s is picked uniformly at random from a finite set S. By PPT, we denote a probabilistic polynomial-time algorithm. Throughout this paper, we use 1^λ as the security parameter. We use $\cdot$ to denote multiplication as well as component-wise multiplication.

2.1 Attribute-Based Encryption

An attribute-based encryption (ABE) scheme for a predicate $\mathsf{P}(\,\cdot\,,\,\cdot\,)$ consists of four algorithms $(\mathsf{Setup}, \mathsf{Enc}, \mathsf{KeyGen}, \mathsf{Dec})$:

$\mathsf{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M}) \to (\mathsf{mpk}, \mathsf{msk})$. The setup algorithm gets as input the security parameter λ, the attribute universe $\mathcal{X}$, the predicate universe $\mathcal{Y}$, the message space $\mathcal{M}$ and outputs the public parameter mpk, and the master key msk.

$\mathsf{Enc}(\mathsf{mpk}, x, m) \to \mathsf{ct}_x$. The encryption algorithm gets as input mpk, an attribute $x \in \mathcal{X}$ and a message $m \in \mathcal{M}$. It outputs a ciphertext ct_x. Note that x is public given ct_x.

$\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, y) \to \mathsf{sk}_y$. The key generation algorithm gets as input msk and a value $y \in \mathcal{Y}$. It outputs a secret key sk_y. Note that y is public given sk_y.

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

Correctness. We require that for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 1$ and all $m \in \mathcal{M}$,

$$\Pr[\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}_y, \mathsf{Enc}(\mathsf{mpk}, x, m)) = m] = 1,$$

where the probability is taken over $(\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M})$, $\mathsf{sk}_y \leftarrow \mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, y)$, and the coins of Enc.

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

$$\mathsf{Adv}_{\mathcal{A}}^{\mathrm{ABE}}(\lambda) := \Pr\left[b = b' : \begin{array}{c} (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M}); \\ (x^*, m_0, m_1) \leftarrow \mathcal{A}^{\mathsf{KeyGen}(\mathsf{msk}, \cdot)}(\mathsf{mpk}); \\ b \leftarrow_{\mathrm{R}} \{0,1\}; \mathsf{ct}_{x^*} \leftarrow \mathsf{Enc}(\mathsf{mpk}, x^*, m_b); \\ b' \leftarrow \mathcal{A}^{\mathsf{KeyGen}(\mathsf{msk}, \cdot)}(\mathsf{ct}_{x^*}) \end{array} \right] - \frac{1}{2}$$

with the restriction that all queries y that $\mathcal{A}$ makes to $\mathsf{KeyGen}(\mathsf{msk}, \cdot)$ satisfies $\mathsf{P}(x^*, y) = 0$ (that is, sk_y does not decrypt ct_{x^*}). An ABE scheme is *adaptively secure* if for all PPT adversaries $\mathcal{A}$, the advantage $\mathsf{Adv}_{\mathcal{A}}^{\mathrm{ABE}}(\lambda)$ is a negligible function in λ.

3 Dual System Groups

This section is largely adapted from [10].

3.1 Overview

Dual system groups contain a triple of abelian groups $(\mathbb{G}, \mathbb{H}, \mathbb{G}_T)$ and a non-degenerate bilinear map $e : \mathbb{G} \times \mathbb{H} \to \mathbb{G}_T$. For concreteness, we may think of $(\mathbb{G}, \mathbb{H}, \mathbb{G}_T)$ as composite-order bilinear groups. Dual system groups take as input a parameter 1^n (think of n as the universe size in KP-ABE) and satisfy the following properties:

(subgroup indistinguishability.) There are two computationally indistinguishable ways to sample correlated $(n+1)$-tuples from $\mathbb{G}^{n+1}$: the "normal" distribution, and a higher-entropy distribution with "semi-functional components". We sample the normal distribution using SampG and the semi-functional components using $\widehat{\mathsf{SampG}}$. An analogous property holds for $\mathbb{H}^{n+1}$, with algorithms SampH and $\widehat{\mathsf{SampH}}$ respectively, with an important distinction in the auxiliary input provided to the distinguisher. For concreteness, think in terms of symmetric bilinear groups of composite order N where

$$\begin{aligned} \mathsf{SampG} &\to (g_1^s, g_1^{s\mathbf{w}}) \in G_N^{n+1} \quad \text{and} \quad \mathsf{SampH} \to (g_1^r, g_1^{r\mathbf{w}}) \in G_N^{n+1} \\ \widehat{\mathsf{SampG}} &\to (g_2^s, g_2^{s\mathbf{w}}) \in G_N^{n+1} \quad \text{and} \quad \widehat{\mathsf{SampH}} \to (g_2^r, g_2^{r\mathbf{w}}) \in G_N^{n+1} \end{aligned}$$

Here, N is the product of three primes p_1, p_2, p_3; g_1, g_2 are generators of order p_1, p_2; and $g_1^{\mathbf{w}} \in G_N^n$ is part of the public parameters.

(associativity.) For all $(g_0, g_1, \ldots, g_n) \in \mathbb{G}^{n+1}$ and all $(h_0, h_1, \ldots, h_n) \in \mathbb{H}^{n+1}$ drawn from the respective normal distributions according to SampG and SampH, we have that for all $i = 1, \ldots, n$,

$$e(g_0, h_i) = e(g_i, h_0).$$

We require this property for correctness (c.f. Eqn. (4)).

(right subgroup $\mathbb{H}$.) There is some distinguished element $h^* \in \mathbb{H}$, which generates the semi-functional components in $\mathbb{H}$. It is convenient to think of h^* as being orthogonal to the normal distribution over $\mathbb{G}$ (c.f. orthogonality). On the other hand, we require that h^* is *not* orthogonal to the semi-functional components in $\mathbb{G}$ (c.f. non-degeneracy), so that we get a random value when we decrypt a semi-functional ciphertext with a semi-functional key.

(parameter-hiding.) Both normal distributions can be efficiently sampled given the public parameters; on the other hand, given only the public parameters, the higher-entropy distributions contain n "units" of information-theoretic entropy (in the semi-functional component), one unit for each of the n elements in the $(n+1)$-tuple apart from the first. In the formal statement, the hidden entropy is captured by n random exponents $(u_1, \ldots, u_n)$ shared across $\mathbb{G}$ and $\mathbb{H}$. It is crucial here that we use the same u_i in $\mathbb{G}$ and in $\mathbb{H}$, so that decryption succeeds with nominally semi-functional objects.

3.2 Definitions

Syntax. Dual system groups consist of six randomized algorithms given by (SampP, SampGT, SampG, SampH) along with ($\widehat{\mathsf{SampG}}, \widehat{\mathsf{SampH}}$):

$\mathsf{SampP}(1^\lambda, 1^n)$: On input $(1^\lambda, 1^n)$, output public and secret parameters $(\mathsf{pp}, \mathsf{sp})$, where:

- pp contains a prime p of length $\Omega(\lambda)$, a triple of abelian groups $(\mathbb{G}, \mathbb{H}, \mathbb{G}_T)$, a non-degenerate bilinear map $e : \mathbb{G} \times \mathbb{H} \to \mathbb{G}_T$, a linear map μ defined on $\mathbb{H}$, along with some additional parameters used by $\mathsf{SampG}, \mathsf{SampH}$;
- the groups $(\mathbb{G}, \mathbb{H}, \mathbb{G}_T)$ are $\mathbb{Z}_p$-modules where $\mathbb{Z}_p$ acts on $\mathbb{G}, \mathbb{H}, \mathbb{G}_T$ via exponentiation;
- given pp, we can uniformly sample from $\mathbb{H}$;
- sp contains $h^* \in \mathbb{H}$ (where $h^* \neq 1$), along with some additional parameters used by $\widehat{\mathsf{SampG}}, \widehat{\mathsf{SampH}}$;

$\mathsf{SampGT} : \mathrm{Im}(\mu) \to \mathbb{G}_\mathrm{T}$. (As a concrete example, suppose $\mu : \mathbb{H} \to \mathbb{G}_T$ and $\mathrm{Im}(\mu) = \mathbb{G}_\mathrm{T}$.)

$\mathsf{SampG}(\mathsf{pp})$: Output $\mathbf{g} \in \mathbb{G}^{n+1}$.

$\mathsf{SampH}(\mathsf{pp})$: Output $\mathbf{h} \in \mathbb{H}^{n+1}$.

$\widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})$: Output $\hat{\mathbf{g}} \in \mathbb{G}^{n+1}$.

$\widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp})$: Output $\hat{\mathbf{h}} \in \mathbb{H}^{n+1}$.

The first four algorithms are used in the actual scheme, whereas the last two algorithms are used only in the proofs of security. We define SampG_0 to denote the first group element in the output of SampG, and we define $\widehat{\mathsf{SampG}}_0, \widehat{\mathsf{SampH}}_0$ analogously.

Remark 1. Given a $\mathbb{Z}_p$-linear function $L : \mathbb{Z}_p^n \to \mathbb{Z}_p$ given by $(w_1, \ldots, w_n) \mapsto a_1 w_1 + \cdots + a_n w_n$ (where $a_1, \ldots, a_n \in \mathbb{Z}_p$ are fixed constants), L acts on $\mathbb{Z}_p$-modules $\mathbb{G}^n, \mathbb{H}^n, \mathbb{G}_T^n$ in the natural way. For instance, $L : \mathbb{G}^n \to \mathbb{G}$ is given by $(g_1, \ldots, g_n) \mapsto g_1^{a_1} \cdots g_n^{a_n}$. This extends also to general $\mathbb{Z}_p$-linear functions $L : \mathbb{Z}_p^n \to \mathbb{Z}_p^m$ coordinate-wise.

Correctness. The requirements for correctness are as follows:

(projective.) For all $h \in \mathbb{H}$ and all coin tosses s, we have $\mathsf{SampGT}(\mu(h); s) = e(\mathsf{SampG}_0(\mathsf{pp}; s), h)$.

(associative.) For all $(g_0, g_1, \ldots, g_n) \leftarrow \mathsf{SampG}(\mathsf{pp})$ and $(h_0, h_1, \ldots, h_n) \leftarrow \mathsf{SampH}(\mathsf{pp})$ and for all $i = 1, \ldots, n$, we have $e(g_0, h_i) = e(g_i, h_0)$.

($\mathbb{H}$-subgroup.) The output of $\mathsf{SampH}(\mathsf{pp})$ is the uniform distribution over a subgroup of $\mathbb{H}^{n+1}$.

Security. The requirements for security are as follows:

(orthogonality.) $\mu(h^*) = 1$.

(non-degeneracy.) For all $\hat{h}_0 \leftarrow \widehat{\mathsf{SampH}}_0(\mathsf{pp}, \mathsf{sp})$, h^* lies in the group generated by $\hat{h}_0$. For all $\hat{g}_0 \leftarrow \widehat{\mathsf{SampG}}_0(\mathsf{pp}, \mathsf{sp})$, we have $e(\hat{g}_0, h^*)^\alpha$ is identically distributed to the uniform distribution over $\mathbb{G}_T$, where $\alpha \leftarrow_{\mathrm{R}} \mathbb{Z}_p$.

(left subgroup indistinguishability.) For any adversary $\mathcal{A}$, we define the advantage function:

$$\mathsf{Adv}^{\mathrm{LS}}_{\mathcal{A}}(\lambda) := \left|\Pr[\ \mathcal{A}(\mathsf{pp}, \boxed{\mathbf{g}}) = 1\ ] - \Pr[\ \mathcal{A}(\mathsf{pp}, \boxed{\mathbf{g} \cdot \hat{\mathbf{g}}}) = 1\ ]\right|$$

where $(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{SampP}(1^\lambda, 1^n)$, $\mathbf{g} \leftarrow \mathsf{SampG}(\mathsf{pp})$, $\hat{\mathbf{g}} \leftarrow \widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})$.

(right subgroup indistinguishability.) For any adversary $\mathcal{A}$, we define the advantage function:

$$\mathsf{Adv}^{\mathrm{RS}}_{\mathcal{A}}(\lambda) := \left|\Pr[\ \mathcal{A}(\mathsf{pp}, h^*, \mathbf{g} \cdot \hat{\mathbf{g}}, \boxed{\mathbf{h}}) = 1\ ] - \Pr[\ \mathcal{A}(\mathsf{pp}, h^*, \mathbf{g} \cdot \hat{\mathbf{g}}, \boxed{\mathbf{h} \cdot \hat{\mathbf{h}}}) = 1\ ]\right|$$

where $(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{SampP}(1^\lambda, 1^n)$, $\mathbf{g} \leftarrow \mathsf{SampG}(\mathsf{pp})$, $\hat{\mathbf{g}} \leftarrow \widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})$, $\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp})$, $\hat{\mathbf{h}} \leftarrow \widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp})$.

(parameter-hiding.) The following distributions are identically distributed

$$\{\mathsf{pp}, h^*, \boxed{\hat{\mathbf{g}}, \hat{\mathbf{h}}}\} \quad \text{and} \quad \{\mathsf{pp}, h^*, \boxed{\hat{\mathbf{g}} \cdot \hat{\mathbf{g}}', \hat{\mathbf{h}} \cdot \hat{\mathbf{h}}'}\}$$

where

$$\begin{aligned} &(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{SampP}(1^\lambda, 1^n); && u_1, \ldots, u_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p; \\ &\hat{\mathbf{g}} = (\hat{g}_0, \ldots) \leftarrow \widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp}); && \hat{\mathbf{h}} = (\hat{h}_0, \ldots) \leftarrow \widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp}); \\ &\hat{\mathbf{g}}' := (1, \hat{g}_0^{u_1}, \ldots, \hat{g}_0^{u_n}) \in \mathbb{G}^{n+1}; && \hat{\mathbf{h}}' := (1, \hat{h}_0^{u_1}, \ldots, \hat{h}_0^{u_n}) \in \mathbb{H}^{n+1}. \end{aligned}$$

4 Instantiations of DSG from k-Lin

We present a new instantiation of dual system groups under the k-Lin assumption, inspired by the constructions in [3, 10].

Overview. The prior construction of DSG [10] (building upon [24, 25, 20, 11]) starts with a random $\mathbf{B} \leftarrow_{\mathrm{R}} \mathrm{GL}_{k+1}(\mathbb{Z}_p)$ and defines $\mathbf{B}^* := (\mathbf{B}^\top)^{-1}$ so that $\mathbf{B}^\top \mathbf{B}^*$ is the identity matrix; then uses $\mathbf{B}$ for $\mathsf{SampG}, \widehat{\mathsf{SampG}}$ and $\mathbf{B}^*$ for $\mathsf{SampH}, \widehat{\mathsf{SampH}}$. In our construction, we may start with any pair of matrices $\mathbf{A}, \mathbf{B}$ in $\mathbb{Z}_p^{(k+1)\times k}$ of full rank:

- In addition, we pick $\mathbf{a}^\perp, \mathbf{b}^\perp$ so that $\mathbf{a}^{\perp\top}\mathbf{A} = \mathbf{b}^{\perp\top}\mathbf{B} = \mathbf{0}$ and $\mathbf{a}^{\perp\top}\mathbf{b}^\perp \neq 0$; we then use $(\mathbf{A}, \mathbf{b}^\perp)$ for $\mathsf{SampG}, \widehat{\mathsf{SampG}}$ and $(\mathbf{B}, \mathbf{a}^\perp)$ for $\mathsf{SampH}, \widehat{\mathsf{SampH}}$.

- We achieve randomization as follows: again, pick a random $\mathbf{W} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times(k+1)}$ and replace $(\mathbf{A}, \mathbf{B})$ with $(\mathbf{W}^\top \mathbf{A}, \mathbf{W}\mathbf{B})$. The associativity property follows from the equation:
$$(\mathbf{W}^\top \mathbf{A})^\top \mathbf{B} = \mathbf{A}^\top (\mathbf{W}\mathbf{B})$$

Interestingly, the prior construction in [10] randomizes by multiplying a random $\mathbf{W}$ on the right, whereas our construction multiplies a random $\mathbf{W}$ on the left. Together with the fact that we no longer require the fact that $\mathbf{B}^\top \mathbf{B}^*$ is the identity, we substantially simplify the proof of subgroup indistinguishability.

4.1 Cryptographic assumptions

We follow the notation and algebraic framework for Diffie-Hellman-like assumptions in [13].

Prime-order bilinear groups. A generator $\mathcal{G}$ takes as input a security parameter λ and outputs a description $(p, G_1, G_2, G_T, g_1, g_2, e)$, where p is a prime of $\Theta(\lambda)$ bits; G_1, G_2 and G_T are cyclic groups of order p; g_1, g_2 are generators of G_1 and G_2 respectively; and $e : G_1 \times G_2 \to G_T$ is a non-degenerate bilinear map. Given $a \in \mathbb{Z}_p$, we use $[a]_1$ to denote g_1^a, $[a]_2$ to denote g_2^a, $[a]_T$ to denote $e(g_1, g_2)^a$. This extends to vectors and matrices in the obvious way. We define $e([\mathbf{A}]_1, [\mathbf{B}]_2) := [\mathbf{A}^\top \mathbf{B}]_T$.

Linear assumption. Let $\mathcal{D}_k$ be an efficiently samplable distribution of matrices $(\mathbf{A}, \mathbf{a}^\perp)$ over $\mathbb{Z}_p^{(k+1)\times k} \times \mathbb{Z}_p^{k+1}$ so that $\mathbf{A}^\top \mathbf{a}^\perp = \mathbf{0}$ and $\mathbf{a}^\perp \neq \mathbf{0}$. In particular, we consider the distribution generated as follows: pick $a_1, \ldots, a_k \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*$ and set

$$\mathbf{A} := \begin{pmatrix} a_1 & & & \\ & a_2 & & \\ & & \ddots & \\ & & & a_k \\ 1 & 1 & \cdots & 1 \end{pmatrix} \in \mathbb{Z}_p^{(k+1)\times k} \quad \text{and} \quad \mathbf{a}^\perp := \begin{pmatrix} a_1^{-1} \\ a_2^{-1} \\ \vdots \\ a_k^{-1} \\ -1 \end{pmatrix} \in \mathbb{Z}_p^{(k+1)}.$$

This distribution captures the k-linear assumption, which stipulates that

$$([\mathbf{A}], [\mathbf{As}]) \approx_c ([\mathbf{A}], [\mathbf{z}])$$

where $\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k, \mathbf{z} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$ in both G_1 and G_2.

Assumption 1 (k-Lin: the k-linear assumption in G_1) *For any adversary $\mathcal{A}$, we define the advantage function:*

$$\begin{aligned}\mathsf{Adv}_{\mathcal{A}}^{k\text{-Lin}} :=\Big| &\Pr[\mathcal{A}((p, G_1, G_2, G_T, g_1, g_2, e); [\mathbf{A}]_1, [\mathbf{As}]_1) = 1] \\ &- \Pr[\mathcal{A}((p, G_1, G_2, G_T, g_1, g_2, e); [\mathbf{A}]_1, [\mathbf{z}]_1) = 1]\Big|\end{aligned}$$

where $(p, G_1, G_2, G_T, g_1, g_2, e) \leftarrow \mathcal{G}(1^\lambda)$, $(\mathbf{A}, \mathbf{a}^\perp) \leftarrow \mathcal{D}_k$, $\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$, $\mathbf{z} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$.

We will slightly abuse notation and also use $\mathsf{Adv}_{\mathcal{A}}^{k\text{-}\mathbf{Lin}}$ to denote the corresponding advantage function for G_2.

Basis lemma. The following structural lemma tells us that if we pick random $(\mathbf{A}, \mathbf{a}^{\perp}), (\mathbf{B}, \mathbf{b}^{\perp}) \leftarrow \mathcal{D}_k$, then with overwhelming probability, both $(\mathbf{A}, \mathbf{b}^{\perp})$ and $(\mathbf{B}, \mathbf{a}^{\perp})$ form a basis for $\mathbb{Z}_p^{k+1}$ and $\mathbf{a}^{\perp}, \mathbf{b}^{\perp}$ are not orthogonal. We will assume henceforth that this property always holds.

Lemma 1 (basis lemma). *With probability* $1 - 1/p$ *over* $(\mathbf{A}, \mathbf{a}^{\perp}), (\mathbf{B}, \mathbf{b}^{\perp}) \leftarrow \mathcal{D}_k$*, it holds that:*

$$\big(\mathbf{a}^{\perp} \notin \mathit{span}(\mathbf{B})\big) \wedge \big(\mathbf{b}^{\perp} \notin \mathit{span}(\mathbf{A})\big) \wedge \big(\mathbf{a}^{\perp^{\top}} \mathbf{b}^{\perp} \neq 0\big).$$

Proof. It is easy to see that if $\mathbf{a}^{\perp^{\top}} \mathbf{b}^{\perp} \neq 0$, then

$$\big(\mathbf{a}^{\perp} \notin \mathrm{span}(\mathbf{B})\big) \quad \text{and} \quad \big(\mathbf{b}^{\perp} \notin \mathrm{span}(\mathbf{A})\big)$$

since every vector in $\mathrm{span}(\mathbf{A})$ is orthogonal to $\mathbf{a}^{\perp}$ and every vector in $\mathrm{span}(\mathbf{B})$ is orthogonal to $\mathbf{b}^{\perp}$. Observe that $\mathbf{a}^{\perp^{\top}} \mathbf{b}^{\perp} = 1 + \sum_{i=1}^{d} (a_i b_i)^{-1}$ and

$$\Pr\Big[1 + \sum_{i=1}^{d} (a_i b_i)^{-1} \neq 0 : a_1, b_1, \ldots, a_k, b_k \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*\Big] = 1 - 1/p.$$

The lemma then follows readily. □

Remark 2. Observe that Lemma 1 is not particular to the k-Lin distribution, since a similar proof works for any example of matrix distribution $\mathcal{D}_k$ presented in [13], namely $\mathcal{U}_{k+1,k}$, k-Casc, k-SCasc and k-ILin [13, Section 3.4].

4.2 Construction

Our construction is as follows:

$\mathsf{SampP}(1^{\lambda}, 1^n)$: On input $(1^{\lambda}, 1^n)$, do:

- run $(p, G_1, G_2, G_T, g_1, g_2, e) \leftarrow \mathcal{G}(1^{\lambda})$, where $\mathcal{G}(1^{\lambda})$ is an asymmetric prime-order group generator;
- define $(\mathbb{G}, \mathbb{H}, \mathbb{G}_T, e) := (G_1^{k+1}, G_2^{k+1}, G_T, e)$;
- sample $(\mathbf{A}, \mathbf{a}^{\perp}), (\mathbf{B}, \mathbf{b}^{\perp}) \leftarrow \mathcal{D}_k$, along with $\mathbf{W}_1, \ldots, \mathbf{W}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times(k+1)}$;
- define $\mu : G_2^{k+1} \to G_T^k$ by $\mu([\mathbf{k}]_2) = [\mathbf{A}^{\top}\mathbf{k}]_T$;
- set $h^* := \big[\mathbf{a}^{\perp}\big]_2$.

Output

$$\mathsf{pp} := \left((p, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, e); \begin{array}{c} [\mathbf{A}]_1, [\mathbf{W}_1^{\top}\mathbf{A}]_1, \ldots, [\mathbf{W}_n^{\top}\mathbf{A}]_1 \\ [\mathbf{B}]_2, [\mathbf{W}_1\mathbf{B}]_2, \ldots, [\mathbf{W}_n\mathbf{B}]_2 \end{array} \right);$$

$$\mathsf{sp} := \big(\; \mathbf{a}^{\perp}, \mathbf{b}^{\perp}, \mathbf{W}_1, \ldots, \mathbf{W}_n \;\big).$$

$\mathsf{SampGT}([\mathbf{p}]_T)$: Pick $\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$ and output $[\mathbf{s}^\top \mathbf{p}]_T \in G_T$.

$\mathsf{SampG}(\mathsf{pp})$: Pick $\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$ and output

$$\Big([\mathbf{A}\mathbf{s}]_1 \,, [\mathbf{W}_1^\top \mathbf{A}\mathbf{s}]_1 \,, \ldots , [\mathbf{W}_n^\top \mathbf{A}\mathbf{s}]_1 \Big) \in (G_1^{k+1})^{n+1}.$$

$\mathsf{SampH}(\mathsf{pp})$: Pick $\mathbf{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$ and output

$$\Big([\mathbf{B}\mathbf{r}]_2 \,, [\mathbf{W}_1 \mathbf{B}\mathbf{r}]_2 \,, \ldots , [\mathbf{W}_n \mathbf{B}\mathbf{r}]_2 \Big) \in (G_2^{k+1})^{n+1}.$$

$\widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})$: Pick $\hat{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*$ and output

$$\Big(\big[\mathbf{b}^\perp \hat{s}\big]_1 \,, \big[\mathbf{W}_1^\top \mathbf{b}^\perp \hat{s}\big]_1 \,, \ldots , \big[\mathbf{W}_n^\top \mathbf{b}^\perp \hat{s}\big]_1 \Big) \in (G_1^{k+1})^{n+1}.$$

$\widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp})$: Pick $\hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*$ and output

$$\Big(\big[\mathbf{a}^\perp \hat{r}\big]_2 \,, \big[\mathbf{W}_1 \mathbf{a}^\perp \hat{r}\big]_2 \,, \ldots , \big[\mathbf{W}_n \mathbf{a}^\perp \hat{r}\big]_2 \Big) \in (G_2^{k+1})^{n+1}.$$

Correctness. We check correctness properties as follows:

(projective.) This follows readily from the fact that for all $\mathbf{k} \in \mathbb{Z}_p^{k+1}, \mathbf{s} \in \mathbb{Z}_p^k$:

$$(\mathbf{A}\mathbf{s})^\top \mathbf{k} = (\mathbf{A}^\top \mathbf{k})^\top \mathbf{s}.$$

(associative.) This follows readily from the fact that for all $\mathbf{s} \in \mathbb{Z}_p^k, \mathbf{r} \in \mathbb{Z}_p^k, \mathbf{W}_i \in \mathbb{Z}_p^{(k+1)\times(k+1)}$:

$$(\mathbf{W}_i^\top \mathbf{A}\mathbf{s})^\top (\mathbf{B}\mathbf{r}) = (\mathbf{A}\mathbf{s})^\top (\mathbf{W}_i \mathbf{B}\mathbf{r}).$$

($\mathbb{H}$-subgroup.) This follows readily from the fact that $\mathbb{Z}_p^k$ is an additive group.

Security. We check security properties as follows:

(orthogonality.) This follows readily from $\mathbf{A}^\top \mathbf{a}^\perp = \mathbf{0}$.

(non-degeneracy.) This follows readily from $\mathbf{b}^{\perp\top} \mathbf{a}^\perp \neq 0$.

We establish left subgroup indistinguishability, right subgroup indistinguishability, and parameter-hiding in the next three lemmas. The left and right subgroup indistinguishability relies on the k-Lin assumption in prime-order groups, whereas parameter-hiding is unconditional.

Lemma 2 (left subgroup indistinguishability from k-Lin). *For any adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that:*

$$\mathsf{Adv}_{\mathcal{A}}^{\mathrm{LS}}(\lambda) \leq \mathsf{Adv}_{\mathcal{B}}^{k\text{-Lin}} + 2/p$$

and $\mathsf{Time}(\mathcal{B}) \approx \mathsf{Time}(\mathcal{A}) + k^2 \cdot \mathrm{poly}(\lambda, n)$ *where* $\mathrm{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$*.*

The proof is a simpler case of the proof of Lemma 3, we omit it here.

Lemma 3 (right subgroup indistinguishability from k-Lin). *For any adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that:*

$$\mathsf{Adv}^{\mathrm{RS}}_{\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{k\text{-Lin}}_{\mathcal{B}} + 2/p$$

and $\mathsf{Time}(\mathcal{B}) \approx \mathsf{Time}(\mathcal{A}) + k^2 \cdot \mathrm{poly}(\lambda, n)$ *where* $\mathrm{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$.

We may rewrite the corresponding advantage function as:

$$\mathsf{Adv}^{\mathrm{RS}}_{\mathcal{A}}(\lambda) := \left|\Pr[\, \mathcal{A}(\mathsf{pp}, h^*, \mathbf{g} \cdot \hat{\mathbf{g}}, \mathbf{h}) = 1\,] - \Pr[\, \mathcal{A}(\mathsf{pp}, h^*, \mathbf{g} \cdot \hat{\mathbf{g}}, \mathbf{h} \cdot \hat{\mathbf{h}}) = 1\,]\right|$$

where

$$\begin{aligned}
(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{SampP}(1^\lambda, 1^n); \quad \mathbf{s}, \mathbf{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k;\ \hat{s}, \hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*; \quad h^* := \left[\mathbf{a}^\perp\right]_2;\\
\mathbf{g} \cdot \hat{\mathbf{g}} := \Big(\left[\mathbf{As} + \mathbf{b}^\perp \hat{s}\right]_1, \left[\mathbf{W}_1^\top(\mathbf{As} + \mathbf{b}^\perp \hat{s})\right]_1, \ldots, \left[\mathbf{W}_n^\top(\mathbf{As} + \mathbf{b}^\perp \hat{s})\right]_1\Big);\\
\mathbf{h} := \Big([\mathbf{Br}]_2, [\mathbf{W}_1\mathbf{Br}]_2, \ldots, [\mathbf{W}_n\mathbf{Br}]_2\Big);\\
\mathbf{h} \cdot \hat{\mathbf{h}} := \Big(\left[\mathbf{Br} + \mathbf{a}^\perp \hat{r}\right]_2, \left[\mathbf{W}_1(\mathbf{Br} + \mathbf{a}^\perp \hat{r})\right]_2, \ldots, \left[\mathbf{W}_n(\mathbf{Br} + \mathbf{a}^\perp \hat{r})\right]_2\Big).
\end{aligned}$$

Proof. The adversary $\mathcal{B}$ samples $(\mathbf{A}, \mathbf{a}^\perp) \leftarrow \mathcal{D}_k$ along with $\mathbf{W}_1, \ldots, \mathbf{W}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times(k+1)}$. Recall that $(\mathbf{B}, \mathbf{a}^\perp)$ is a basis for $\mathbb{Z}_p^{k+1}$, so $\{\mathbf{Br} + \mathbf{a}^\perp \hat{r} : \mathbf{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k, \hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*\}$ is statistically close to the uniform distribution. Adversary $\mathcal{B}$ then gets as input

$$\big(\ (p, G_1, G_2, G_T, g_1, g_2, e), [\mathbf{B}]_2, \left[\mathbf{Br} + \mathbf{a}^\perp \hat{r}\right]_2\ \big)$$

where either $\hat{r} = 0$ or $\hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*$, and proceeds as follows:

Simulating pp, h^*. Output

$$\begin{array}{l} [\mathbf{A}]_1,\ [\mathbf{W}_1^\top \mathbf{A}]_1,\ \ldots,\ [\mathbf{W}_n^\top \mathbf{A}]_1 \\ [\mathbf{B}]_2,\ [\mathbf{W}_1\mathbf{B}]_2,\ \ldots,\ [\mathbf{W}_n\mathbf{B}]_2 \end{array} \quad \text{and} \quad \left[\mathbf{a}^\perp\right]_2$$

Simulating $\left[\mathbf{As} + \mathbf{b}^\perp \hat{s}\right]_1, \left[\mathbf{W}_i^\top(\mathbf{As} + \mathbf{b}^\perp \hat{s})\right]_1$. Note that $\mathcal{B}$ does not know $\mathbf{b}^\perp$. Instead, $\mathcal{B}$ samples $\tilde{\mathbf{s}} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$ and outputs

$$[\tilde{\mathbf{s}}]_1,\ [\mathbf{W}_1^\top \tilde{\mathbf{s}}]_1,\ \ldots,\ [\mathbf{W}_n^\top \tilde{\mathbf{s}}]_1\,.$$

Observe that $\mathbf{As} + \mathbf{b}^\perp \hat{s}$ is statistically close to the uniform vector $\tilde{\mathbf{s}}$ as long as $\mathbf{b}^\perp \notin \mathrm{span}(\mathbf{A})$ and $\hat{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p$.

Simulating the challenge. Upon receiving a k-Lin challenge, $\mathcal{B}$ outputs

$$\left[\mathbf{Br} + \mathbf{a}^\perp \hat{r}\right]_2,\ \left[\mathbf{W}_1(\mathbf{Br} + \mathbf{a}^\perp \hat{r})\right]_2,\ \ldots,\ \left[\mathbf{W}_n(\mathbf{Br} + \mathbf{a}^\perp \hat{r})\right]_2$$

where either $\hat{r} = 0$ or $\hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p$.

Observe that:

- if $\hat{r} = 0$, then we can write the output challenge as

$$[\mathbf{Br}]_2\,,\ [\mathbf{W}_1\mathbf{Br}]_2\,,\ \ldots,\ [\mathbf{W}_n\mathbf{Br}]_2\,.$$

  which equals $\mathbf{h}$; we obtain the left distribution in the statement of the lemma;
- if $\hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p$, then we can write the output challenge as

$$\left[\mathbf{Br} + \mathbf{a}^{\perp}\hat{r}\right]_2\,,\ \left[\mathbf{W}_1(\mathbf{Br} + \mathbf{a}^{\perp}\hat{r})\right]_2\,,\ \ldots,\ \left[\mathbf{W}_n(\mathbf{Br} + \mathbf{a}^{\perp}\hat{r})\right]_2\,.$$

  which equals $\mathbf{h} \cdot \hat{\mathbf{h}}$; we obtain the right distribution in the statement of the lemma.

Typically, we sample $\hat{s}, \hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*$ for $\widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})$ and $\widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp})$; this yields a $2/p$ negligible difference in the advantage. The lemma then follows readily. □

Lemma 4 (parameter-hiding). *The following distributions are identically distributed*

$$\left\{ \mathit{pp}, \left[\mathbf{a}^{\perp}\right]_2, \begin{matrix} \left[\mathbf{b}^{\perp}\hat{s}\right]_1, \left[\mathbf{W}_1^{\top}\mathbf{b}^{\perp}\hat{s}\right]_1, \ldots, \left[\mathbf{W}_n^{\top}\mathbf{b}^{\perp}\hat{s}\right]_1 \\ \left[\mathbf{a}^{\perp}\hat{r}\right]_2, \left[\mathbf{W}_1\mathbf{a}^{\perp}\hat{r}\right]_2, \ldots, \left[\mathbf{W}_n\mathbf{a}^{\perp}\hat{r}\right]_2 \end{matrix} \right\} \quad \textit{and}$$

$$\left\{ \mathit{pp}, \left[\mathbf{a}^{\perp}\right]_2, \begin{matrix} \left[\mathbf{b}^{\perp}\hat{s}\right]_1, \left[(\mathbf{W}_1^{\top}\mathbf{b}^{\perp} + u_1\mathbf{b}^{\perp})\hat{s}\right]_1, \ldots, \left[(\mathbf{W}_n^{\top}\mathbf{b}^{\perp} + u_n\mathbf{b}^{\perp})\hat{s}\right]_1 \\ \left[\mathbf{a}^{\perp}\hat{r}\right]_2, \left[(\mathbf{W}_1\mathbf{a}^{\perp} + u_1\mathbf{a}^{\perp})\hat{r}\right]_2, \ldots, \left[(\mathbf{W}_n\mathbf{a}^{\perp} + u_n\mathbf{a}^{\perp})\hat{r}\right]_2 \end{matrix} \right\}$$

where $(\mathit{pp}, \mathit{sp}) \leftarrow \mathsf{SampP}(1^{\lambda}, 1^n)$, $\hat{s}, \hat{r} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^*$ *and* $u_1, \ldots, u_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p$.

Proof. Fix $g_1, g_2, (\mathbf{A}, \mathbf{a}^{\perp}), (\mathbf{B}, \mathbf{b}^{\perp}), \hat{s}, \hat{r}$; that is, we prove that the statement holds for all $g_1, g_2, (\mathbf{A}, \mathbf{a}^{\perp}), (\mathbf{B}, \mathbf{b}^{\perp}), \hat{s}, \hat{r}$. Set $\mathbf{V} := \mathbf{a}^{\perp}\mathbf{b}^{\perp\top} \in \mathbb{Z}_p^{(k+1)\times(k+1)}$ which satisfies the following properties:

$$\mathbf{V}^{\top}\mathbf{A} = \mathbf{0} \quad \text{and} \quad \mathbf{V}\mathbf{B} = \mathbf{0} \tag{5}$$

$$\mathbf{V}\mathbf{a}^{\perp} = (\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})\mathbf{a}^{\perp} \quad \text{and} \quad \mathbf{V}^{\top}\mathbf{b}^{\perp} = (\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})\mathbf{b}^{\perp} \tag{6}$$

Eqn. (6) basically says that $\mathbf{a}^{\perp}$ and $\mathbf{b}^{\perp}$ are the respective eigenvectors of $\mathbf{V}$ and $\mathbf{V}^{\top}$. Now, consider the following "change of variables" in the first distribution, namely, replace

$$\mathbf{W}_i \ \text{ with } \ \mathbf{W}_i + u_i(\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})^{-1}\mathbf{V}, \ \ i = 1, \ldots, n.$$

Clearly, this does not change the first distribution. Now, observe that

$$\left[(\mathbf{W}_i + u_i(\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})^{-1}\mathbf{V})^{\top}\mathbf{A}\right]_1 = [\mathbf{W}_i^{\top}\mathbf{A}]_1\,;$$

$$\left[(\mathbf{W}_i + u_i(\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})^{-1}\mathbf{V})\mathbf{B}\right]_2 = [\mathbf{W}_i\mathbf{B}]_2$$

where we use (5) in the last equalities. That is, pp remains unchanged. In addition, we have

$$\Big[(\mathbf{W}_i + u_i(\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})^{-1}\mathbf{V})^{\top}\mathbf{b}^{\perp}\Big]_1 = \big[\mathbf{W}_i^{\top}\mathbf{b}^{\perp} + u_i\mathbf{b}^{\perp}\big]_1\,;$$
$$\Big[(\mathbf{W}_i + u_i(\mathbf{a}^{\perp\top}\mathbf{b}^{\perp})^{-1}\mathbf{V})\mathbf{a}^{\perp}\Big]_2 = \big[\mathbf{W}_i\mathbf{a}^{\perp} + u_i\mathbf{a}^{\perp}\big]_2$$

where we use (6) in the last equalities. Indeed, this is exactly the second distribution. □

5 Predicate Encodings

In this section, we describe a refinement of the predicate encodings from [31, 2] which we use in this work. We refer to Section 1.1 for an overview of the refinement.

Predicate encodings. Fix a predicate $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$. A $\mathbb{Z}_p$*-bilinear predicate encoding* for P is a tuple of deterministic algorithms $(\mathsf{sE}, \mathsf{rE}, \mathsf{kE}, \mathsf{sD}, \mathsf{rD})$ satisfying the following properties:

(linearity.) For all $(x,y) \in \mathcal{X} \times \mathcal{Y}$, the functions $\mathsf{sE}(x,\cdot)$, $\mathsf{rE}(y,\cdot)$, $\mathsf{kE}(y,\cdot)$, $\mathsf{sD}(x,y,\cdot)$, $\mathsf{rD}(x,y,\cdot)$ are $\mathbb{Z}_p$-linear.

(restricted α-reconstruction.) For all $(x,y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x,y) = 1$ and for all $\mathbf{w} \in \mathcal{W}$:

$$\mathsf{sD}(x,y,\mathsf{sE}(x,\mathbf{w})) = \mathsf{rD}(x,y,\mathsf{rE}(y,\mathbf{w})) \quad \text{and} \quad \mathsf{rD}(x,y,\mathsf{kE}(y,\alpha)) = \alpha.$$

(α-privacy.) For all $(x,y) \in \mathcal{X}\times\mathcal{Y}$ such that $\mathsf{P}(x,y) = 0$, and for all $\alpha \in \mathbb{Z}_p$, the joint distribution $\{\mathsf{sE}(x,\mathbf{w}), \mathsf{kE}(y,\alpha) + \mathsf{rE}(y,\mathbf{w})\}$ *perfectly* hides α. That is, for all $\alpha \in \mathbb{Z}_p$, the following joint distributions[2] are *identically* distributed:

$$\big\{x, y, \alpha, \mathsf{sE}(x,\mathbf{w}), \mathsf{kE}(y,\alpha) + \mathsf{rE}(y,\mathbf{w})\big\} \quad \text{and} \quad \big\{x, y, \alpha, \mathsf{sE}(x,\mathbf{w}), \mathsf{rE}(y,\mathbf{w})\big\}$$

where the randomness is taken over $\mathbf{w} \leftarrow_{\mathrm{R}} \mathcal{W}$.

Remark 3. Given a predicate encoding as defined above, we can construct an encoding $(\mathsf{rE}', \mathsf{sE}')$ which achieves the notion in [31, 2] by considering:

$$\mathsf{sE}' = \mathsf{sE} \quad \text{and} \quad \mathsf{rE}'(y,\alpha,\mathbf{w},r) = \big(r,\, \mathsf{kE}(y,\alpha) + r\cdot\mathsf{rE}(y,\mathbf{w})\big).$$

Note that rE' leaks no information about $\mathbf{w}$ when $r = 0$ which trivially yields the $\mathbf{w}$-hiding property in [31] (aka parameter-hiding in [2]). Here, we use the fact that kE does not depend on $\mathbf{w}$.

[2] Note that since $\mathsf{kE}(y,\cdot)$ is $\mathbb{Z}_p$-linear, we have $\mathsf{kE}(y,0) + \mathsf{rE}(y,\mathbf{w}) = \mathsf{rE}(y,\mathbf{w})$.

Example: equality. Fix a prime integer p. Consider the equality predicate where $\mathcal{X} = \mathcal{Y} = \mathbb{Z}_p$ and $\mathsf{P}(x, y) = 1$ iff $x = y$. The following is a predicate encoding for equality used in [4, 21]:

$$\begin{array}{lll} \mathsf{sE}(x, (w_1, w_2)) := w_1 + w_2 x & \mathsf{rE}(y, (w_1, w_2)) := w_1 + w_2 y & \mathsf{kE}(y, \alpha) := \alpha \\ \mathsf{sD}(x, y, c) = c & \mathsf{rD}(x, y, k) = k & \end{array}$$

When $x = y$, $w_1 + w_2 x = w_1 + w_2 y$ and we can reconstruct α. For α-privacy, we exploit the fact that $(w_1 + w_2 x, w_1 + w_2 y)$ are pairwise independent when $x \neq y$.

6 ABE from Dual System Groups and Predicate Encodings

Starting from a predicate encoding for P, we construct an ABE for P using dual system groups. We refer to Section 1.1 for an overview of the scheme, which is of the form:

$$\begin{aligned} \mathsf{mpk} &:= \big(\, g_1,\ g_1^{\mathbf{w}},\ e(g_1, g_1)^{\alpha} \,\big) \\ \mathsf{sk}_y &:= \big(\, g_1^r,\ g_1^{\mathsf{kE}(y,\alpha) + r \cdot \mathsf{rE}(y, \mathbf{w})} \,\big) \\ \mathsf{ct}_x &:= \big(\, g_1^s,\ g_1^{s \cdot \mathsf{sE}(x, \mathbf{w})},\ e(g_1, g_1)^{\alpha s} \cdot m \big) \end{aligned}$$

We will generate mpk using $\mathsf{SampP}(1^\lambda, 1^n)$, where $\mathbf{w} \in \mathbb{Z}_p^n$. We will use $\mathsf{SampG}(\mathsf{pp})$ to generate the terms $(g_1^s, g_1^{s\mathbf{w}})$ in the ciphertext, from which we can compute $(g_1^s, g_1^{s \cdot \mathsf{sE}(x, \mathbf{w})})$ by linearity of $\mathsf{sE}(x, \cdot)$. Similarly, we use $\mathsf{SampH}(\mathsf{pp})$ to generate the terms $(g_1^r, g_1^{r\mathbf{w}})$ in the secret key, from which we can compute $(g_1^r, g_1^{r \cdot \mathsf{rE}(y, \mathbf{w})})$. We replace g_1^α with $\mathsf{msk} \leftarrow_{\mathrm{R}} \mathbb{H}$.

6.1 Construction

$\mathsf{Setup}(1^\lambda, 1^n)$: On input $(1^\lambda, 1^n)$, first sample

$$(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{SampP}(1^\lambda, 1^n).$$

Pick $\mathsf{msk} \leftarrow_{\mathrm{R}} \mathbb{H}$ and output the master public and secret key pair

$$\mathsf{mpk} := (\ \mathsf{pp},\ \mu(\mathsf{msk})\ ) \quad \text{and} \quad \mathsf{msk}.$$

$\mathsf{Enc}(\mathsf{mpk}, x, m)$: On input $x \in \mathcal{X}$ and $m \in \mathbb{G}_T$, sample

$$(g_0, g_1, \ldots, g_n) \leftarrow \mathsf{SampG}(\mathsf{pp}; s),\ g_T' \leftarrow \mathsf{SampGT}(\mu(\mathsf{msk}); s)$$

and output[3]

$$\mathsf{ct}_x := (\ C_0 := g_0,\ \mathbf{C}_1 := \mathsf{sE}(x, (g_1, \ldots, g_n)),\ C' := g_T' \cdot m\ ).$$

[3] See Remark 1 for an explanation of the function $\mathsf{sE}(x, (g_1, \ldots, g_n))$.

$\mathsf{KeyGen}(\mathsf{mpk}, \mathsf{msk}, y)$: On input $y \in \mathcal{Y}$, sample

$$(h_0, h_1, \dots, h_n) \leftarrow \mathsf{SampH}(\mathsf{pp})$$

and output

$$\mathsf{sk}_y := (\ K_0 := h_0,\ \mathbf{K}_1 := \mathsf{kE}(y, \mathsf{msk}) \cdot \mathsf{rE}(y, (h_1, \dots, h_n))\ ) \,.$$

$\mathsf{Dec}(\mathsf{mpk}, \mathsf{sk}_y, \mathsf{ct}_x)$: Compute

$$e(g_0, \mathsf{msk}) \leftarrow e(C_0, \mathsf{rD}(x, y, \mathbf{K}_1))/e(\mathsf{sD}(x, y, \mathbf{C}_1), K_0)$$

and recover the message as

$$m \leftarrow C' \cdot e(g_0, \mathsf{msk})^{-1} \in \mathbb{G}_T.$$

Correctness. For all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 1$, we have

$$\begin{aligned}
&e(C_0, \mathsf{rD}(x, y, \mathbf{K}_1)) \\
&= e(g_0, \mathsf{rD}(x, y, \mathsf{rE}(y, (h_1, \dots, h_n)))) \cdot e(g_0, \mathsf{rD}(x, y, \mathsf{kE}(y, \mathsf{msk}))) \\
&= e(g_0, \mathsf{rD}(x, y, \mathsf{rE}(y, (h_1, \dots, h_n)))) \cdot e(g_0, \mathsf{msk}) \\
&= \mathsf{rD}(x, y, \mathsf{rE}(y, (e(g_0, h_1), \dots, e(g_0, h_n)))) \cdot e(g_0, \mathsf{msk}) \\
&= \mathsf{rD}(x, y, \mathsf{rE}(y, (e(g_1, h_0), \dots, e(g_n, h_0)))) \cdot e(g_0, \mathsf{msk}) \\
&= \mathsf{sD}(x, y, \mathsf{sE}(x, (e(g_1, h_0), \dots, e(g_n, h_0)))) \cdot e(g_0, \mathsf{msk}) \\
&= e(\mathsf{sD}(x, y, \mathsf{sE}(x(g_1, \dots, g_n))), h_0) \cdot e(g_0, \mathsf{msk}) \\
&= e(\mathsf{sD}(x, y, \mathbf{C}_1), K_0) \cdot e(g_0, \mathsf{msk})
\end{aligned}$$

In line 2, we use linearity of $\mathsf{rD}(x, y, \cdot)$ and $e(g_0, \cdot)$. In line 3 and line 6, we use α-reconstruction. In line 4 and line 7, we use the fact that the functions $e(g_0, \cdot)$, $e(\cdot, h_0)$ and $\mathsf{sD}(x, y, \mathsf{sE}(y, \cdot))$ commute with linear functions. That is, given a $\mathbb{Z}_p$-linear function $L : \mathbb{Z}_p^n \to \mathbb{Z}_p$ given by $(w_1, \dots, w_n) \mapsto a_1 w_1 + \dots + a_n w_n$, we have:

$$\begin{aligned}
e(g_0, L(h_1, \dots, h_n)) &= e(g_0, h_1^{a_1} \cdots h_n^{a_n}) \\
&= e(g_0, h_1)^{a_1} \cdots e(g_0, h_n)^{a_n} \\
&= L(e(g_0, h_1), \dots, e(g_0, h_n))
\end{aligned}$$

In line 5, we use associativity in DSG. Finally, by *projective*, $g_T' = e(g_0, \mathsf{msk})$. Correctness follows readily.

6.2 Proof of Security

We prove the following theorem:

game	ciphertext $(C_0, \mathbf{C}_1, C')$	secret key $(K_0, \mathbf{K}_1)$	justification
0	$(1, \mathbf{1}, 1)$	$(1, (h^*)^{\mathsf{kE}(y,0)} \cdot \mathbf{1})$	$\mathbf{1} = (h^*)^{\mathsf{kE}(y,0)}$
1	$\boxed{(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))}$	$(1, (h^*)^{\mathsf{kE}(y,0)} \cdot \mathbf{1})$	left subgroup ind.
2.i.1	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))$	$(\boxed{\hat{h}_0}, (h^*)^{\mathsf{kE}(y,0)} \cdot \boxed{\mathsf{rE}(y, \hat{\mathbf{h}})})$	right subgroup ind.
2.i.2	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))$	$(\hat{h}_0, \boxed{(h^*)^{\mathsf{kE}(y,\alpha)}} \cdot \mathsf{rE}(y, \hat{\mathbf{h}}))$	α-privacy
2.i.3	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))$	$(\boxed{1}, (h^*)^{\mathsf{kE}(y,\alpha)} \cdot \boxed{\mathbf{1}})$	right subgroup ind.
3	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), \boxed{\text{random}})$	$(1, (h^*)^{\mathsf{kE}(y,\alpha)} \cdot \mathbf{1})$	

Fig. 3. Sequence of games in the "semi-functional" space. We omitted the normal components: those sampled using $\mathsf{SampG}, \mathsf{SampH}$, and we omitted $e(g_0, \mathsf{msk}) \cdot m$ in C' and $\mathsf{kE}(y, \mathsf{msk})$ in sk_y. We drew a box to highlight the differences between each game and the preceding one, and games $2.i.$x refer to the i'th secret key. The semi-functional components of the keys transition from $(h^*)^{\mathsf{kE}(y,0)}$ to $(h^*)^{\mathsf{kE}(y,\alpha)}$. For the final transition, we use the fact that $e(\hat{g}_0, \mathsf{msk})$ is statistically random given $\mathsf{msk} \cdot (h^*)^{\alpha}$.

Theorem 1. *Under the left and right subgroup indistinguishability (described in Section 3), the ABE scheme described in Section 6.1 is adaptively secure (in the sense of Definition 2.1). More precisely, for any adversary $\mathcal{A}$ that makes at most q key queries against the ABE scheme, there exist adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that:*

$$\mathsf{Adv}^{\text{ABE}}_{\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\text{LS}}_{\mathcal{B}_1}(\lambda) + q \cdot \mathsf{Adv}^{\text{RS}}_{\mathcal{B}_2}(\lambda) + q \cdot \mathsf{Adv}^{\text{RS}}_{\mathcal{B}_3}(\lambda)$$

and

$$\max\{\mathsf{Time}(\mathcal{B}_1), \mathsf{Time}(\mathcal{B}_2), \mathsf{Time}(\mathcal{B}_3)\} \approx \mathsf{Time}(\mathcal{A}) + q \cdot \mathrm{poly}(\lambda, n)$$

where $\mathrm{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$*.*

The proof follows via a series of games, analogous to that in [10, 31, 30, 21], and outlined in Fig. 3. We first define two auxiliary algorithms and then the semi-functional distributions, upon which we can describe the games.

Auxiliary algorithms. We consider the following algorithms:

$\widehat{\mathsf{Enc}}(\mathsf{pp}, x, m; \mathsf{msk}, \mathbf{t})$: On input $x \in \mathcal{X}$, $m \in \mathbb{G}_T$, and $\mathbf{t} := (T_0, T_1, \ldots, T_n) \in \mathbb{G}^{n+1}$, output

$$\mathsf{ct}_x := (\ T_0,\ \mathsf{sE}(x, (T_1, \ldots, T_n)),\ e(T_0, \mathsf{msk}) \cdot m\ )\,.$$

$\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk}', y; \mathbf{t})$: On input $\mathsf{msk}' \in \mathbb{H}$, $y \in \mathcal{Y}$, and $\mathbf{t} := (T_0, T_1, \ldots, T_n) \in \mathbb{H}^{n+1}$, output

$$\mathsf{sk}_y := \big(\ T_0,\ \mathsf{kE}(y, \mathsf{msk}') \cdot \mathsf{rE}(y, (T_1, \ldots, T_n))\ \big)\,.$$

In all the proofs and figures that follow, we denote $\mathsf{sE}(x, (T_1, \ldots, T_n))$ by $\mathsf{sE}(x, \mathbf{t})$ for notational convenience, and we define $\mathsf{rE}(y, \mathbf{t})$ analogously.

Auxiliary distributions.

Semi-functional master secret key.

$$\widehat{\mathsf{msk}} := \mathsf{msk} \cdot (h^*)^{\alpha},$$

where $\boxed{\alpha \leftarrow_{\mathrm{R}} \mathbb{Z}_p}$.

Semi-functional ciphertext.

$$\widehat{\mathsf{Enc}}(\mathsf{pp}, x, m; \mathsf{msk}, \boxed{\mathbf{g} \cdot \hat{\mathbf{g}}}),$$

where $\boxed{\mathbf{g} \leftarrow \mathsf{SampG}(\mathsf{pp}) \text{ and } \hat{\mathbf{g}} \leftarrow \widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})}$.

Pseudo-normal secret key.

$$\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk}, y; \boxed{\mathbf{h} \cdot \hat{\mathbf{h}}}),$$

where fresh $\boxed{\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp}) \text{ and } \hat{\mathbf{h}} \leftarrow \widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp})}$ are chosen for each secret key.

Pseudo-SF secret key.

$$\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{\widehat{\mathsf{msk}}}, y; \mathbf{h} \cdot \hat{\mathbf{h}}),$$

where fresh $\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp})$ and $\hat{\mathbf{h}} \leftarrow \widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp})$ are chosen for each secret key.

Semi-functional secret key.

$$\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \widehat{\mathsf{msk}}, y; \boxed{\mathbf{h}}),$$

where a fresh $\boxed{\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp})}$ is chosen for each secret key. We note that the semi-functional key generation algorithm is identical to the normal key generation except that it replaces msk with $\widehat{\mathsf{msk}}$ as input.

Game sequence. We present a series of games. We write $\mathsf{Adv}_{\mathrm{xxx}}(\lambda)$ to denote the advantage of $\mathcal{A}$ in $\mathsf{Game}_{\mathrm{xxx}}$.

- Game_0: is the real security game (c.f. Section 2.1).
- Game_1: is the same as Game_0 except that the challenge ciphertext is semi-functional.
- $\mathsf{Game}_{2,i,1}$ for i from 1 to q, $\mathsf{Game}_{2,i,1}$ is the same as Game_1 except that the first $i-1$ keys are semi-functional, the last $q-i$ keys are normal while the i'th key is pseudo-normal.

- $\mathsf{Game}_{2,i,2}$ for i from 1 to q, $\mathsf{Game}_{2,i,2}$ is the same as Game_1 except that the first $i-1$ keys are semi-functional, the last $q-i$ keys are normal while the i'th key is pseudo-SF.
- $\mathsf{Game}_{2,i,3}$ for i from 1 to q, $\mathsf{Game}_{2,i,3}$ is the same as Game_1 except that the first i keys are semi-functional, the last $q-i$ keys are normal.
- Game_3: is the same as $\mathsf{Game}_{2,q,3}$, except that the challenge ciphertext is a semi-functional encryption of a random message in $\mathbb{G}_T$.

In Game_3, the view of the adversary is statistically independent of the challenge bit b. Hence, $\mathsf{Adv}_3(\lambda) = 0$. We complete the proof by establishing the following sequence of lemmas. We omit the proofs of lemmas 5, 6, 8, 9 as they are the same as those of lemmas 1, 2, 5, 6 in [10, Section 4].

Lemma 5 (normal to SF ciphertext: Game_0 to Game_1). *For any adversary $\mathcal{A}$ that makes at most q key queries, there exists an adversary $\mathcal{B}_1$ such that*

$$|\mathsf{Adv}_0(\lambda) - \mathsf{Adv}_1(\lambda)| \leq \mathsf{Adv}^{\text{LS}}_{\mathcal{B}_1}(\lambda)$$

and $\mathsf{Time}(\mathcal{B}_1) \approx \mathsf{Time}(\mathcal{A}) + q \cdot \text{poly}(\lambda, n)$ *where* $\text{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$*.*

Lemma 6 (normal to pseudo-normal keys: $\mathsf{Game}_{2,i-1,3}$ to $\mathsf{Game}_{2,i,1}$). *For $i = 1, \ldots, q$, for any adversary $\mathcal{A}$ that makes at most q key queries, there exists an adversary $\mathcal{B}_2$ such that*

$$|\mathsf{Adv}_{2,i-1,3}(\lambda) - \mathsf{Adv}_{2,i,1}(\lambda)| \leq \mathsf{Adv}^{\text{RS}}_{\mathcal{B}_2}(\lambda)$$

and $\mathsf{Time}(\mathcal{B}_2) \approx \mathsf{Time}(\mathcal{A}) + q \cdot \text{poly}(\lambda, n)$ *where* $\text{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$*. (We note that* $\mathsf{Game}_{2,0,3}$ *is identical to* Game_1*.)*

Lemma 7 (pseudo-normal to pseudo-SF keys: $\mathsf{Game}_{2,i,1}$ to $\mathsf{Game}_{2,i,2}$). *For $i = 1, \ldots, q$, we have*

$$|\mathsf{Adv}_{2,i,1}(\lambda) - \mathsf{Adv}_{2,i,2}(\lambda)| = 0.$$

Proof. Observe that the only difference between $\mathsf{Game}_{2,i,1}$ and $\mathsf{Game}_{2,i,2}$ lies in that we replace msk in $\mathsf{Game}_{2,i,1}$ with $\widehat{\mathsf{msk}}$ in $\mathsf{Game}_{2,i,2}$ as input for the i'th secret key query, where $\mathsf{msk} \leftarrow_{\text{R}} \mathbb{H}$, $\alpha \leftarrow_{\text{R}} \mathbb{Z}_p$ and $\widehat{\mathsf{msk}} := \mathsf{msk} \cdot (h^*)^\alpha$. Thus, it suffices to establish the following:

> *Claim.* For all α, all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, where $\mathsf{P}(x, y) = 0$, the following distributions are identically distributed:
>
> $$\{\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha, \widehat{\mathsf{Enc}}(\mathsf{pp}, x, m_\beta; \mathsf{msk}, \mathbf{g} \cdot \hat{\mathbf{g}}), \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{\mathsf{msk}}, y; \mathbf{h} \cdot \hat{\mathbf{h}})\} \quad \text{and}$$
> $$\{\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha, \widehat{\mathsf{Enc}}(\mathsf{pp}, x, m_\beta; \mathsf{msk}, \mathbf{g} \cdot \hat{\mathbf{g}}), \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{\mathsf{msk} \cdot (h^*)^\alpha}, y; \mathbf{h} \cdot \hat{\mathbf{h}})\}.$$

We defer the proof of the claim for now, and first explain how the lemma follows from the claim. Given $(\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha)$, we can output $\mathsf{mpk} := (\mathsf{pp}, \mu(\mathsf{msk}))$ and generate the first $i-1$ semi-functional secret keys, and the remaining $q-i$ normal secret keys using

$$\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk} \cdot (h^*)^\alpha, y; \mathsf{SampH}(\mathsf{pp})) \quad \text{and} \quad \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk}, y; \mathsf{SampH}(\mathsf{pp}))$$

respectively.

This would in turn imply that $\mathsf{Game}_{2,i,1}$ and $\mathsf{Game}_{2,i,2}$ are statistically indistinguishable. We note that this holds even if the adversary chooses y adaptively after seeing the challenge ciphertext ct_{x^*}, or if the challenge x^* is chosen after the adversary sees sk_y. □

Proof (of claim). By linearity, we have:

$$\begin{aligned}
\widehat{\mathsf{Enc}}(\mathsf{pp}, x, m_\beta; \mathsf{msk}, \mathbf{g} \cdot \hat{\mathbf{g}}) &= \widehat{\mathsf{Enc}}(\mathsf{pp}, x, m_\beta; \mathsf{msk}, \mathbf{g}) \cdot \widehat{\mathsf{Enc}}(\mathsf{pp}, x, 1; \mathsf{msk}, \hat{\mathbf{g}}) \\
\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk}, y; \mathbf{h} \cdot \hat{\mathbf{h}}) &= \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk}, y; \mathbf{h}) \cdot \widehat{\mathsf{KeyGen}}(\mathsf{pp}, 1, y; \hat{\mathbf{h}}) \\
\widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk} \cdot (h^*)^\alpha, y; \mathbf{h} \cdot \hat{\mathbf{h}}) &= \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \mathsf{msk}, y; \mathbf{h}) \cdot \widehat{\mathsf{KeyGen}}(\mathsf{pp}, (h^*)^\alpha, y; \hat{\mathbf{h}})
\end{aligned}$$

Therefore, it suffices to show that:

$$\begin{aligned}
&\{\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha, \widehat{\mathsf{Enc}}(\mathsf{pp}, x, 1; \mathsf{msk}, \hat{\mathbf{g}}), \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{1}, y; \hat{\mathbf{h}})\} \quad \text{and} \\
&\{\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha, \widehat{\mathsf{Enc}}(\mathsf{pp}, x, 1; \mathsf{msk}, \hat{\mathbf{g}}), \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{\ (h^*)^\alpha\ }, y; \hat{\mathbf{h}})\}
\end{aligned}$$

are identically distributed.

By parameter-hiding, we may replace $(\mathsf{pp}, h^*, \boxed{\hat{\mathbf{g}}, \hat{\mathbf{h}}})$ with $(\mathsf{pp}, h^*, \boxed{\hat{\mathbf{g}} \cdot \hat{\mathbf{g}}', \hat{\mathbf{h}} \cdot \hat{\mathbf{h}}'})$, which means it suffices to show that:

$$\begin{aligned}
&\{\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha, \widehat{\mathsf{Enc}}(\mathsf{pp}, x, 1; \mathsf{msk}, \hat{\mathbf{g}} \cdot \hat{\mathbf{g}}'), \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{1}, y; \hat{\mathbf{h}} \cdot \hat{\mathbf{h}}')\} \quad \text{and} \\
&\{\mathsf{pp}, \mathsf{msk}, (h^*)^\alpha, \widehat{\mathsf{Enc}}(\mathsf{pp}, x, 1; \mathsf{msk}, \hat{\mathbf{g}} \cdot \hat{\mathbf{g}}'), \widehat{\mathsf{KeyGen}}(\mathsf{pp}, \boxed{\ (h^*)^\alpha\ }, y; \hat{\mathbf{h}} \cdot \hat{\mathbf{h}}')\}
\end{aligned}$$

are identically distributed. At this point, we expand the expressions for $\widehat{\mathsf{Enc}}$ and $\widehat{\mathsf{KeyGen}}$:

$$\begin{aligned}
\widehat{\mathsf{Enc}}(\mathsf{pp}, x, 1; \mathsf{msk}, \hat{\mathbf{g}} \cdot \hat{\mathbf{g}}') &= (\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}) \cdot \mathsf{sE}(x, \hat{\mathbf{g}}'),\ e(\hat{g}_0, \mathsf{msk})) \\
&= (\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}) \cdot \hat{g}_0^{\mathsf{sE}(x,\mathbf{u})},\ e(\hat{g}_0, \mathsf{msk}))
\end{aligned}$$

where $\mathbf{u}$ denotes the vector $\mathbf{u} := (u_1, \ldots, u_n)$ and thus $\mathsf{sE}(x, \hat{\mathbf{g}}') = \mathsf{sE}(x, \hat{g}_0^{\mathbf{u}}) = \hat{g}_0^{\mathsf{sE}(x,\mathbf{u})}$;

$$\begin{aligned}
\widehat{\mathsf{KeyGen}}(\mathsf{pp}, 1, y; \hat{\mathbf{h}} \cdot \hat{\mathbf{h}}') &= (\hat{h}_0, \mathsf{rE}(y, \hat{\mathbf{h}}) \cdot \hat{h}_0^{\mathsf{rE}(y,\mathbf{u})}) \\
\widehat{\mathsf{KeyGen}}(\mathsf{pp}, (h^*)^\alpha, y; \hat{\mathbf{h}} \cdot \hat{\mathbf{h}}') &= (\hat{h}_0, \mathsf{kE}(y, (h^*)^\alpha) \cdot \mathsf{rE}(y, \hat{\mathbf{h}}) \cdot \hat{h}_0^{\mathsf{rE}(y,\mathbf{u})})
\end{aligned}$$

Since h^* lies in the group generated by $\hat{h}_0$, we have $\mathsf{kE}(y, (h^*)^\alpha) = \mathsf{kE}(y, (h_0)^{\alpha'}) = \hat{h}_0^{\mathsf{kE}(y,\alpha')}$ for some $\alpha' \in \mathbb{Z}_p$; the claim then follows readily from α'-privacy, that is, $\mathsf{rE}(y, \mathbf{u})$ and $\mathsf{kE}(y, \alpha') + \mathsf{rE}(y, \mathbf{u})$ are identically distributed. □

Lemma 8 (pseudo-SF to SF keys: $\mathsf{Game}_{2,i,2}$ to $\mathsf{Game}_{2,i,3}$). *For $i = 1, \ldots, q$, for any adversary $\mathcal{A}$ that makes at most q key queries, there exists an adversary $\mathcal{B}_3$ such that*

$$|\mathsf{Adv}_{2,i,2}(\lambda) - \mathsf{Adv}_{2,i,3}(\lambda)| \leq \mathsf{Adv}^{\text{RS}}_{\mathcal{B}_3}(\lambda)$$

and $\mathsf{Time}(\mathcal{B}_3) \approx \mathsf{Time}(\mathcal{A}) + q \cdot \mathrm{poly}(\lambda, n)$ *where* $\mathrm{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$.

Lemma 9 (final transition: $\mathsf{Game}_{2,q,3}$ to Game_3). *For any adversary $\mathcal{A}$, we have*

$$|\mathsf{Adv}_{2,q,3}(\lambda) - \mathsf{Adv}_3(\lambda)| = 0.$$

7 Extension to Weakly Attribute-Hiding

We present an extension of our framework to weakly attribute-hiding predicate encryption [19, 6]. A predicate encryption scheme has the same syntax as an ABE in Section 2.1 except the attribute x on the ciphertext is not public; for security, we require in addition that x remains hidden from the adversary.

7.1 Security definition

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

$$\mathsf{Adv}^{\text{PE}}_{\mathcal{A}}(\lambda) := \Pr\left[b = b' : \begin{array}{c} (\mathsf{mpk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, \mathcal{X}, \mathcal{Y}, \mathcal{M}); \\ (x_0^*, x_1^*, m_0, m_1) \leftarrow \mathcal{A}^{\mathsf{KeyGen}(\mathsf{msk}, \cdot)}(\mathsf{mpk}); \\ b \leftarrow_{\text{R}} \{0,1\}; \mathsf{ct}_{x_b^*} \leftarrow \mathsf{Enc}(\mathsf{mpk}, x_b^*, m_b); \\ b' \leftarrow \mathcal{A}^{\mathsf{KeyGen}(\mathsf{msk}, \cdot)}(\mathsf{ct}_{x_b^*}) \end{array} \right] - \frac{1}{2}$$

with the restriction that all queries y that $\mathcal{A}$ makes to $\mathsf{KeyGen}(\mathsf{msk}, \cdot)$ satisfies $\mathsf{P}(x_0^*, y) = \mathsf{P}(x_1^*, y) = 0$ (that is, sk_y does not decrypt the challenge ciphertext). A predicate encryption scheme is *adaptively secure and weakly attribute-hiding* if for all PPT adversaries $\mathcal{A}$, the advantage $\mathsf{Adv}^{\text{PE}}_{\mathcal{A}}(\lambda)$ is a negligible function in λ.[4]

7.2 Attribute-Hiding Encodings

We say that a $\mathbb{Z}_p$-bilinear predicate encoding (c.f. Section 5) for $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$ is *attribute-hiding* if it satisfies the following additional properties:

(x-oblivious α-reconstruction.) $\mathsf{sD}(x, y, \cdot)$ and $\mathsf{rD}(x, y, \cdot)$ are independent of x.

[4] In a fully attribute-hiding scheme, the adversary is also allowed key queries y for which $\mathsf{P}(x_0^*, y) = \mathsf{P}(x_1^*, y) = 1$, in which case the challenge messages m_0, m_1 must be equal.

(attribute-hiding.) For all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 0$, the joint distribution of $\{\mathsf{sE}(x, \mathbf{w}), \mathsf{rE}(y, \mathbf{w})\}$ is uniformly random. That is, the following distributions are *identically* distributed:

$$\big\{x, y, \mathsf{sE}(x, \mathbf{w}), \mathsf{rE}(y, \mathbf{w})\big\} \quad \text{and} \quad \big\{x, y, \mathbf{v}\big\}$$

where the randomness is taken over $\mathbf{w} \leftarrow_{\mathrm{R}} \mathcal{W}$ and $\mathbf{v} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{|\mathsf{sE}(\cdot)|+|\mathsf{rE}(\cdot)|}$.

7.3 Attribute-Hiding Dual System Groups

Recall from the introduction in Section 1.1 that to realize weakly attribute-hiding predicate encryption, we will use the fact that for any vector $\mathbf{c} \in \mathbb{Z}_p^{k+1}$ outside the span of $\mathbf{A}$, the vector $\mathbf{W}^\top \mathbf{c} \in \mathbb{Z}_p^{k+1}$ is uniformly random given $\mathbf{W}^\top \mathbf{A} \in \mathbb{Z}_p^{(k+1)\times k}$, provided $\mathbf{WB}$ remains hidden. We can then use $\mathbf{W}^\top \mathbf{c}$ to completely blind the attribute in the challenge ciphertext. We also need to make sure that the semi-functional secret keys do not leak any *additional* information about $\mathbf{WB}$. The former is captured by $\mathbb{G}$-uniformity, and the latter by $\mathbb{H}$-hiding. In particular, the secret keys in the predicate encryption scheme satisfy the following properties:

- the distribution of normal secret keys is completely determined given $\mathbf{B}, \mathbf{W}_1\mathbf{B}, \ldots, \mathbf{W}_n\mathbf{B}$ and leaks no *additional* information about $\mathbf{W}_1, \ldots, \mathbf{W}_n$;
- the distribution of semi-functional secret keys is completely determined given $\mathbf{A}, \mathbf{W}_1^\top\mathbf{A}, \ldots, \mathbf{W}_n^\top\mathbf{A}$ and leaks no *additional* information about $\mathbf{W}_1, \ldots, \mathbf{W}_n$.

Additional properties. We assume that pp in dual system groups has a $\mathsf{pp}_{\mathbb{G}}$-component which is sufficient to run SampG. We then require dual system groups to satisfy the following additional properties.

($\mathbb{H}$-hiding) There is an (inefficient) randomized procedure SampH^* that given $\mathsf{pp}_{\mathbb{G}}$ and h^*, outputs a distribution identical to that of

$$\mathbf{h} \cdot (h^*)^{(0,\mathbf{v})}$$

where $\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp})$, $\mathbf{v} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^n$.

($\mathbb{G}$-uniformity) The following distributions are identically distributed

$$\Big\{\ \mathsf{pp}_{\mathbb{G}}, h^*, \boxed{\mathbf{g} \cdot \hat{\mathbf{g}}}\ \Big\} \quad \text{and} \quad \Big\{\ \mathsf{pp}_{\mathbb{G}}, h^*, \boxed{\mathbf{g}'}\ \Big\}$$

where $(\mathsf{pp}, \mathsf{sp}) \leftarrow \mathsf{SampP}(1^\lambda, 1^n)$, $\mathbf{g} = (g_0, \ldots) \leftarrow \mathsf{SampG}(\mathsf{pp})$, $\hat{\mathbf{g}} = (\hat{g}_0, \ldots) \leftarrow \widehat{\mathsf{SampG}}(\mathsf{pp}, \mathsf{sp})$, $\mathbf{g}' \leftarrow_{\mathrm{R}} \{g_0\hat{g}_0\} \times \mathbb{G}^n$.

In the full version of this paper, we show that our instantiations satisfy the additional attribute-hiding requirements when $\mathsf{pp}_{\mathbb{G}}$ is defined to be:

$$\mathsf{pp}_{\mathbb{G}} := \big(\ (p, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, e); [\mathbf{A}]_1, [\mathbf{W}_1^\top\mathbf{A}]_1, \ldots, [\mathbf{W}_n^\top\mathbf{A}]_1,\ [\mathbf{B}]_2\ \big).$$

7.4 Weakly Attribute-Hiding PE

Starting from an attribute-hiding encoding and an attribute-hiding dual system group, we can construct a predicate encryption scheme as described in Section 6.1, with the following modification: we put $\mathsf{pp}_{\mathbb{G}}$ instead of pp in mpk (which suffices for SampG and Enc). We show that the ensuing scheme is weakly attribute-hiding:

Theorem 2. *Under the left and right subgroup indistinguishability (described in Section 3), the predicate encryption scheme described above is adaptively secure and weakly attribute-hiding (in the sense of Definition 7.1). More precisely, for any adversary $\mathcal{A}$ that makes at most q key queries against the predicate encryption scheme, there exist adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that:*

$$\mathsf{Adv}^{\mathrm{PE}}_{\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\mathrm{LS}}_{\mathcal{B}_1}(\lambda) + q \cdot \mathsf{Adv}^{\mathrm{RS}}_{\mathcal{B}_2}(\lambda) + q \cdot \mathsf{Adv}^{\mathrm{RS}}_{\mathcal{B}_3}(\lambda)$$

and

$$\max\{\mathsf{Time}(\mathcal{B}_1), \mathsf{Time}(\mathcal{B}_2), \mathsf{Time}(\mathcal{B}_3)\} \approx \mathsf{Time}(\mathcal{A}) + q \cdot \mathrm{poly}(\lambda, n)$$

where $\mathrm{poly}(\lambda, n)$ *is independent of* $\mathsf{Time}(\mathcal{A})$.

The proof follows via a series of games, outlined in Fig. 4.

Auxiliary distributions. The auxiliary algorithms and distributions are the same as in Section 6.2 with the following modifications: (1) pseudo-SF and semi-functional secret keys have additional h^*-components, (2) $\widehat{\mathsf{Enc}}$ and $\widehat{\mathsf{KeyGen}}$ get as input $\mathsf{pp}_{\mathbb{G}}$ instead of pp (neither algorithm needs to run SampH).

Pseudo-SF secret key.

$$\widehat{\mathsf{KeyGen}}(\mathsf{pp}_{\mathbb{G}}, \boxed{\widehat{\mathsf{msk}}}, y; \boxed{\mathbf{h} \cdot \hat{\mathbf{h}} \cdot (h^*)^{(0,\mathbf{v})}}),$$

where fresh $\boxed{\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp}),\ \hat{\mathbf{h}} \leftarrow \widehat{\mathsf{SampH}}(\mathsf{pp}, \mathsf{sp}),\text{ and } \mathbf{v} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^n}$ are chosen for each secret key.

Semi-functional secret key.

$$\widehat{\mathsf{KeyGen}}(\mathsf{pp}_{\mathbb{G}}, \widehat{\mathsf{msk}}, y; \boxed{\mathbf{h} \cdot (h^*)^{(0,\mathbf{v})}}),$$

where a fresh $\boxed{\mathbf{h} \leftarrow \mathsf{SampH}(\mathsf{pp}) \text{ and } \mathbf{v} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^n}$ are chosen for each secret key.

game	ciphertext $(C_0, \mathbf{C}_1, C')$	secret key $(K_0, \mathbf{K}_1)$	justification
0	$(1, \mathbf{1}, 1)$	$(1, (h^*)^{\mathsf{kE}(y,0)} \cdot \mathbf{1})$	$\mathbf{1} = (h^*)^{\mathsf{kE}(y,0)}$
1	$\boxed{(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))}$	$(1, (h^*)^{\mathsf{kE}(y,0)} \cdot \mathbf{1})$	left subgroup ind.
2.i.1	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))$	$(\boxed{\hat{h}_0}, (h^*)^{\mathsf{kE}(y,0)} \cdot \boxed{\mathsf{rE}(y, \hat{\mathbf{h}})})$	right subgroup ind.
2.i.2	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))$	$(\hat{h}_0, \boxed{(h^*)^{\mathsf{kE}(y,\alpha)+\mathsf{rE}(y,\mathbf{v}_i)}} \cdot \mathsf{rE}(y, \hat{\mathbf{h}}))$	AH encoding
2.i.3	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), e(\hat{g}_0, \mathsf{msk}))$	$(\boxed{1}, (h^*)^{\mathsf{kE}(y,\alpha)+\mathsf{rE}(y,\mathbf{v}_i)} \cdot \boxed{\mathbf{1}})$	right subgroup ind.
3	$(\hat{g}_0, \mathsf{sE}(x, \hat{\mathbf{g}}), \boxed{\text{random}})$	$(1, (h^*)^{\mathsf{kE}(y,\alpha)+\mathsf{rE}(y,\mathbf{v}_i)} \cdot \mathbf{1})$	
4	$(\hat{g}_0, \boxed{\text{random}}, \text{random})$	$(1, (h^*)^{\mathsf{kE}(y,\alpha)+\mathsf{rE}(y,\mathbf{v}_i)} \cdot \mathbf{1})$	AH encoding
			$\mathbb{H}$-hiding
			$\mathbb{G}$-uniformity

Fig. 4. Sequence of games in the "semi-functional" space for weakly attribute-hiding PE. We omitted the normal components: those sampled using $\mathsf{SampG}, \mathsf{SampH}$, and we omitted $e(g_0, \mathsf{msk}) \cdot m$ in C' and $\mathsf{kE}(y, \mathsf{msk})$ in sk_y. We drew a box to highlight the differences between each game and the preceding one, and games $2.i.$x refer to the i'th secret key. The semi-functional components of the keys transition from $(h^*)^{\mathsf{kE}(y,0)}$ to $(h^*)^{\mathsf{kE}(y,\alpha)+\mathsf{rE}(y,\mathbf{v}_i)}$, with a fresh $\mathbf{v}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^n$ for the i'th key. In the penultimate transition, we use the fact that $e(\hat{g}_0, \mathsf{msk})$ is statistically random given $\mathsf{msk} \cdot (h^*)^{\alpha}$. In the final transition, we use the fact that $\mathbf{C}_1$ (including normal components) is statistically random.

Game sequence. We proceed exactly as in Section 6.2 with the same auxiliary algorithms but with the following modifications: (1) the distributions of pseudo-SF and semi-functional secret keys have additional h^*-components, (2) the challenge ciphertext uses the attribute x_b^* as defined in the security experiment, and (3) we append an extra game Game_4 where we switch x_b^* to random at the end:

- Game_0: is the real security game (c.f. Section 7.1).
- The descriptions of Game_1, $\mathsf{Game}_{2,i,1}$, $\mathsf{Game}_{2,i,2}$, $\mathsf{Game}_{2,i,3}$, and Game_3 are identical to those in Section 6.2, we omit them here.
- Game_4: is the same as Game_3, except we replace x_b^* in the challenge ciphertext with a random attribute $x^* \leftarrow_{\mathrm{R}} \mathcal{X}$.

In Game_4, the view of the adversary is statistically independent of the challenge bit b. Hence, $\mathsf{Adv}_4(\lambda) = 0$. We defer the proofs to the full version of this paper.

Acknowledgments. We thank Eike Kiltz and Jiaxin Pan for insightful discussions, and the anonymous reviewers for helpful feedback on the write-up.

References

[1] S. Agrawal, D. M. Freeman, and V. Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *ASIACRYPT*, pages 21–40, 2011.

[2] N. Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In *EUROCRYPT*, pages 557–577, 2014.

[3] O. Blazy, E. Kiltz, and J. Pan. (Hierarchical) identity-based encryption from affine message authentication. In *CRYPTO*, pages 408–425, 2014.

[4] D. Boneh and X. Boyen. Efficient selective-ID secure identity-based encryption without random oracles. In *EUROCRYPT*, pages 223–238, 2004.

[5] D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.

[6] D. Boneh and B. Waters. Conjunctive, subset, and range queries on encrypted data. In *TCC*, pages 535–554, 2007. Also Cryptology ePrint Archive, Report 2006/287.

[7] D. Boneh, X. Boyen, and E.-J. Goh. Hierarchical identity based encryption with constant size ciphertext. In *EUROCRYPT*, pages 440–456, 2005.

[8] D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, pages 533–556, 2014.

[9] J. Chen and H. Wee. Fully, (almost) tightly secure IBE and dual system groups. In *CRYPTO (2)*, pages 435–460, 2013.

[10] J. Chen and H. Wee. Dual system groups and its applications — compact HIBE and more. IACR Cryptology ePrint Archive, Report 2014/265, 2014. Preliminary version in [9].

[11] J. Chen, H. W. Lim, S. Ling, H. Wang, and H. Wee. Shorter IBE and signatures via asymmetric pairings. In *Pairing*, pages 122–140, 2012.

[12] C. Cocks. An identity based encryption scheme based on quadratic residues. In *IMA Int. Conf.*, pages 360–363, 2001.

[13] A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. L. Villar. An algebraic framework for diffie-hellman assumptions. In *CRYPTO (2)*, pages 129–147, 2013.

[14] D. M. Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In *EUROCRYPT*, pages 44–61, 2010.

[15] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM Conference on Computer and Communications Security*, pages 89–98, 2006.

[16] A. Guillevic. Comparing the pairing efficiency over composite-order and prime-order elliptic curves. In *ACNS*, pages 357–372, 2013.

[17] Y. Ishai and H. Wee. Partial garbling schemes and their applications. In *ICALP*, pages 650–662, 2014.

[18] C. S. Jutla and A. Roy. Shorter quasi-adaptive NIZK proofs for linear subspaces. In *ASIACRYPT (1)*, pages 1–20, 2013.

[19] J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.

[20] A. Lewko. Tools for simulating features of composite order bilinear groups in the prime order setting. In *EUROCRYPT*, pages 318–335, 2012.

[21] A. B. Lewko and B. Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In *TCC*, pages 455–479, 2010.

[22] A. B. Lewko, T. Okamoto, A. Sahai, K. Takashima, and B. Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *EUROCRYPT*, pages 62–91, 2010.

[23] T. Okamoto and K. Takashima. Homomorphic encryption and signatures from vector decomposition. In *Pairing*, pages 57–74, 2008.

[24] T. Okamoto and K. Takashima. Hierarchical predicate encryption for inner-products. In *ASIACRYPT*, pages 214–231, 2009.

[25] T. Okamoto and K. Takashima. Fully secure functional encryption with general relations from the decisional linear assumption. In *CRYPTO*, pages 191–208, 2010.

[26] T. Okamoto and K. Takashima. Achieving short ciphertexts or short secret-keys for adaptively secure general inner-product encryption. In *CANS*, pages 138–159, 2011. Also, Cryptology ePrint Archive, Report 2011/648.

[27] A. Sahai and B. Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.

[28] A. Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.

[29] K. Takashima. Expressive attribute-based encryption with constant-size ciphertexts from the decisional linear assumption. In *Security and Cryptography for Networks (SCN)*, pages 298–317, 2014.

[30] B. Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *CRYPTO*, pages 619–636, 2009.

[31] H. Wee. Dual system encryption via predicate encodings. In *TCC*, pages 616–637, 2014.