

A Study of Pair Encodings: Predicate Encryption in Prime Order Groups

Shashank Agrawal*

Melissa Chase †

Abstract

Pair encodings and predicate encodings, recently introduced by Attrapadung (Eurocrypt 2014) and Wee (TCC 2014) respectively, greatly simplify the process of designing and analyzing predicate and attribute-based encryption schemes. However, they are still somewhat limited in that they are restricted to composite order groups, and the information theoretic properties are not sufficient to argue about many of the schemes. Here we focus on pair encodings, as the more general of the two. We first study the structure of these objects, then propose a new relaxed but still information theoretic security property. Next we show a generic construction for predicate encryption in prime order groups from our new property; it results in either semi-adaptive or full security depending on the encoding, and gives security under SXDH or DLIN. Finally, we demonstrate the range of our new property by using it to design the first semi-adaptively secure CP-ABE scheme with constant size ciphertexts.

Keywords. Predicate encryption, Attribute-based encryption, Pair encoding schemes, Dual system technique, Short ciphertexts

*University of Illinois Urbana-Champaign. Email: sagrawl2@illinois.edu. Part of this work was done when the author was at Microsoft Research.

†Microsoft Research. Email: melissac@microsoft.com.

1 Introduction

In traditional public key encryption systems, a message is encrypted under a particular public key, with the guarantee that it can only be decrypted by the party holding the corresponding secret key. Attribute based encryption (ABE), introduced in [SW05], instead allows us to use attributes to determine who has the power to decrypt. In these systems, there is a single entity which publishes system parameters and distributes the appropriate decryption keys to various parties. In key-policy ABE (KP-ABE) [GPSW06], a message is encrypted under a set of attributes describing that message, and each decryption key is associated with a policy describing which ciphertexts it can decrypt. Conversely, in ciphertext-policy ABE (CP-ABE) [BSW07] each user is given a decryption key that depends on his attributes, and ciphertexts are encrypted with policies describing which users can decrypt them. ABE has been proposed for a variety of applications, from social network privacy to pay-per-view broadcasting to health record access-control to cloud security (see e.g. [PTMW06, TBEM08, BBS⁺09, APG⁺11, SRGS12]).

Recently there has been a lot of progress in terms of both security and functionality. Using the dual system framework introduced by Waters [Wat09], several works [LOS⁺10, LW12] have designed ABE schemes that satisfy the natural security definition, avoiding the restrictions of selective security¹. Other works consider extra features like short ciphertexts whose length is independent of the size of the associated attribute set and policy [ALdP11, YAHK14], or “unbounded” schemes that place no bounds on the space of possible attributes or the number of attributes that can be tied to a ciphertext or key [LW11, OT12, RW13]. Predicate encryption [BSW11] generalizes the concept to require only that the ciphertext and key are associated with values x, y , and decryption succeeds iff some predicate $P(x, y)$ holds. Note that in this work we assume that x and y are revealed by the ciphertext and key respectively; we do not consider attribute-hiding [BW07, KSW08] or predicate-hiding [SSW09, BRS13].

As these schemes have progressed, however, constructions and proofs have become increasingly complex. Many of the proposed schemes require composite order pairings, in which the order of the pairing groups is a product of two or more primes; since these schemes require that factoring the group order is hard, this in practice means that these groups must be at least an order of magnitude larger than prime order groups of comparable security level, and according to [Gui13] composite order pairing computations are at least 2 orders of magnitude slower. This has prompted efforts to design schemes in prime order groups [OT10, OT12, Fre10, Lew12, HHH⁺14], but many of these schemes still have fairly high cost as compared to their selectively secure counterparts, and designing and analyzing security of such schemes can be quite challenging.

Two very recent works, by Wee [Wee14] and Attrapadung [Att14] make significant progress in simplifying the design and analysis of new constructions. These works introduce simple new objects, called predicate encodings and pair encodings respectively in the two works, which can be used to construct ABE and other predicate encryption schemes. Essentially, they consider one decryption key and one ciphertext, and focus on what happens in the exponent space. Both formalisms introduce simple information theoretic properties on these objects and show that if these properties are met, they can be extended into fully secure ABE/predicate encryption schemes. The major advantage of this approach is that instead of having to design and prove security of a complex scheme, now all one has to do is design and analyze an appropriate encoding, which is a much simpler task. This vastly simplifies the design of new schemes, and in fact, both works resulted in new constructions and more efficient variants of previously known schemes.

¹The original construction of Sahai and Waters [SW05], and much of the following work, considers what is referred to as the selective security model, in which the adversary must commit to the attributes/policy used in the challenge ciphertext before requesting any decryption keys.

Currently these works have two primary limitations. First, they both result in ABE schemes that rely on composite order pairings, which as explained above is very undesirable from an efficiency standpoint. The second drawback is that the strict information theoretic properties they require from the underlying objects mean that there are many constructions that they cannot capture in their model. Attrapadung [Att14] addresses this by introducing a computational security notion, which allows several more interesting constructions to be captured in the framework. However, this security notion is much harder to analyze - it involves not only the encodings in the exponent space, but also elements in the composite order group in which it is embedded, and the proofs that the encodings satisfy this notion are not only computational (rather than information theoretic) but are based on much stronger assumptions.

Still these encodings seem extremely promising as a way to simplify the design and analysis of predicate encryption schemes. In our work we further study these objects, with the aim of understanding them better and beginning to address these limitations. In particular we focus on the pair encodings from [Att14], as they seem to be able to capture more constructions.

Our Contributions. First, we study the *structure of pair encodings*. Attrapadung’s pair encodings have only limited structural requirements. This means that he is able to capture many existing constructions in his framework, although as mentioned above, in many cases the information theoretic security property he defines does not hold for these schemes. A better understanding of the natural structure of these schemes may help to design new schemes, by providing better intuition for what is important and simply by limiting the search space.

Here we consider two structural properties. First we assume a simple property that describes where the public parameters appear in the key and ciphertext. This seems to reflect some basic structure, as all the pair encodings in [Att14] have this property. Looking ahead, this property allows us to instantiate these schemes efficiently in prime order groups. We then show that this implies a second, seemingly unrelated property involving the use of random variables in the key and ciphertexts. We can use this second property to simplify our security definitions and analyses.

Using this understanding, we propose a *relaxation of the information theoretic security property* proposed in [Att14]. This property essentially allows us to consider the scheme at smaller granularity than an entire key or ciphertext. It is still information theoretic, and it does not depend on the group in which it will be used; this means it is still easy to analyze whether a given encoding satisfies this property. We consider two flavors of this property and show that the stronger of the two is implied by the security properties in [Att14]. However, we will see that our new property is indeed a relaxation in that it allows us to consider encodings that did not satisfy the original property. Thus, we make a first step towards addressing the limitations of the strict information theoretic property of previous work.

Next we present a *generic construction of predicate encryption* from pair encodings. Here we make use of the dual system groups introduced by [CW13]; although we must modify their properties slightly, we show that their instantiations are still sufficient². We show that pair encodings which satisfy the stronger flavor of our new property result in fully secure predicate encryption schemes, while pair encodings which satisfy the weaker flavor result in schemes which can still be shown to be semi-adaptively secure³. While full security is preferable, we will see that this second result allows us to design schemes in areas in which even selectively secure constructions are hard to construct.

This approach has two advantages. First, this means that we can transform any pair encoding

²Since we use these groups in a black box way, any improvement in the underlying instantiation will translate directly into an improvement in our generic construction. In particular we believe that the simplified new dual system groups proposed in [CGW15] satisfy our modified definitions as well, so they could be used to simplify our construction.

³Unlike selective security, in semi-adaptive security an adversary is not forced to commit to the challenge before seeing the public parameters.

scheme which satisfies the information theoretic security properties in [Att14] into a fully secure ABE or predicate encryption scheme *in a prime order group* based only on the SXDH or DLIN assumption. This results in schemes which are of practical efficiency, with strong security guarantees based on mild assumptions. Moreover, the advantage of this approach is that while proof of our generic construction is fairly involved, analyzing a given pair encoding scheme to verify the necessary property is still quite straightforward.

Finally, to demonstrate how our relaxed security property allows us to consider additional functionalities, we present a new pair encoding for *CP-ABE with constant-size ciphertext*. When used in our generic construction, this results in a CP-ABE with constant size-ciphertext which is semi-adaptively secure and can be instantiated under either SXDH or DLIN. To the best of our knowledge, prior to our work there were no known schemes for constant-size CP-ABE, even considering only selectively security and allowing for very strong assumptions.⁴ This shows then that our new techniques allow us to consider a strictly greater range of schemes; we hope that they will continue to prove useful and lead to other interesting constructions.

Other related work. As mentioned above, the original works of [Wee14] and [Att14] gave constructions only in composite order groups. In a recent work, however, Chen, Gay, and Wee [CGW15] proposed a transformation to go from pair encodings to prime order predicate encryption schemes, requiring the same strong information theoretic property on the underlying pair encoding as in [Wee14]. However, they also require strict restrictions on the structure of pair encodings, which are not satisfied by most of the encodings which had previously been proposed; essentially this requires that there be only one unit of randomness in each ciphertext or key. They show that the previous encodings which satisfy the information theoretic property from [Att14] (the basic KP- and CP-ABE schemes) have counterparts which satisfy these stricter requirements. This results in the most efficient known constructions for a number of problems. As mentioned above, our generic construction can be applied directly to the original pair encodings [Att14]; this will yield similar constructions, with slightly different tradeoffs (generally smaller public parameters but slower decryption). Interestingly, our relaxed perfect security property is designed to leverage exactly the kind of structure they prohibit, so perhaps it suggests another way forward for predicates that cannot be addressed under their model.

In concurrent work, Attrapadung [Att15] proposed a generic construction that compiles any secure (computational or information-theoretic) pair encoding scheme for a predicate R to a fully secure FE scheme for the same predicate in prime-order groups under Matrix Diffie-Hellman assumption [EHK⁺13] (of which DLIN is a special case) with an additional q-type assumption in the case of pair encodings that only satisfy the computational security definition from [Att14]. This then also gives prime order group constructions for any predicate encoding scheme satisfying the strong information theoretic property under DLIN, and for KP-ABE with short ciphertext (as well as unbounded KP-ABE and ABE for regular languages) under a q-type assumption. However, as compared to this work, our results have the following advantages: First, we use dual system groups in a black box way, which simplifies the transformation, unifies prime and composite order group constructions, and means that any new construction of dual system groups directly gives new constructions for ABE. Moreover, our relaxed perfect security property allows us to show semi-adaptive security for the short ciphertext schemes based only on SXDH or DLIN, without any q-type assumptions; in addition to giving us the new results on CP-ABE, we can also give a much simpler proof of semi-adaptive security for Attrapadung’s KP-ABE with short ciphertexts, and this proof does not require q-type assumptions. (See Appendix E.)

Finally, we mention the concurrent work of Attrapadung, Hanaoka, and Yamada [AHY15]. This

⁴Here we discount threshold access policies because when only threshold policies are considered, CP-ABE and KP-ABE are equivalent.

work presents various conversions among pair encoding schemes. Among other things, they show that if one starts with the KP-ABE scheme with constant-size ciphertexts recently proposed by Takashima [Tak14], then by applying the conversion one gets a CP-ABE scheme with constant-size ciphertexts, which is *selectively* secure under the DLIN assumption. On the other hand, we get a *semi-adaptive* scheme secure under any assumption which can be used to construct dual system groups (which includes SXDH, DLIN, etc). Moreover, since Takashima's construction does not use any abstractions, our construction is significantly more modular, easier to analyze and easier to extend. As we view the CP-ABE more as a test-case for the utility of our new definition and transformation, having an approach that can extend easily to other types of ABE schemes seems particularly valuable.

2 Preliminaries

We use \cong, \equiv and \approx to denote statistical, perfect and computational indistinguishability respectively. Security parameter is denoted by λ , and $\text{negl}(\lambda)$ denotes a negligible function in λ .

We normally use lower case letters in bold to denote vectors; but if a vector's elements are themselves vectors, we use upper case. For two vectors $\mathbf{u} = (u_1, \dots, u_n)$ and $\mathbf{v} = (v_1, \dots, v_n)$, we use $\mathbf{u} \cdot \mathbf{v}$ to denote the entry-wise product, i.e., (u_1v_1, \dots, u_nv_n) , and $\langle u, v \rangle$ to denote the inner-product, i.e., $\sum_{i=1}^n u_iv_i$. The \cdot operator naturally extends to vectors of vectors (or matrices): if $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m)$ and $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_m)$, then $\mathbf{U} \cdot \mathbf{V} = (\mathbf{u}_1 \cdot \mathbf{v}_1, \dots, \mathbf{u}_m \cdot \mathbf{v}_m)$. $g^{\mathbf{u}}$ should be interpreted as the vector $(g^{u_1}, \dots, g^{u_n})$. $g^{\mathbf{A}}$, where \mathbf{A} is a matrix, should be interpreted in an analogous way.

We use $\mathbf{u}_1, \dots, \mathbf{u}_m \leftarrow \text{SampAlg}(\cdot)$ to denote that the algorithm SampAlg is run m times with independent coin tosses to generate samples $\mathbf{u}_1, \dots, \mathbf{u}_m$. Since the output of this algorithm is a vector, we also use $(u_1, \dots, u_n) \leftarrow \text{SampAlg}(\cdot)$ to denote that a single sample with co-ordinates u_1, \dots, u_n is drawn from SampAlg (this should not be confused with the previous notation). Finally, $a \leftarrow_R S$ denotes drawing an element a uniformly at random from the set S .

Bilinear Pairings: Let \mathbb{G}, \mathbb{H} and \mathbb{G}_T be three multiplicative groups. A pairing $e : \mathbb{G} \times \mathbb{H} \rightarrow \mathbb{G}_T$ is bilinear if for all $g \in \mathbb{G}, h \in \mathbb{H}$ and $a, b \in \mathbb{Z}$, $e(g^a, h^b) = e(g, h)^{ab}$. This pairing is non-degenerate if whenever $e(g, h) = 1_{\mathbb{G}_T}$, then either $g = 1_{\mathbb{G}}$ or $h = 1_{\mathbb{H}}$ (where $1_{\mathbb{G}}$, for instance, denotes the identity element of \mathbb{G} .) We will only be interested in bilinear pairings that are efficiently computable.

The order of an element g of a group G is the smallest positive integer a such that $g^a = 1_G$. The exponent of a group is defined as the least common multiple of the orders of all elements of the group. One can show that if a non-degenerate bilinear pairing $e : \mathbb{G} \times \mathbb{H} \rightarrow \mathbb{G}_T$ can be defined over three groups \mathbb{G}, \mathbb{H} and \mathbb{G}_T , then they all have the same exponent. We use $\text{exp}(G)$ to denote the exponent of a group G .

Homomorphism: A homomorphism from a group $\langle G, \cdot \rangle$ to a group $\langle H, \oplus \rangle$ is a function $\psi : G \rightarrow H$ such that for all $g_1, g_2 \in G$, $\psi(g_1 \cdot g_2) = \psi(g_1) \oplus \psi(g_2)$. We define two sets with respect to a homomorphism: $\text{Image}(\psi) = \{\psi(g) \mid g \in G\}$ and $\text{Kernel}(\psi) = \{g \in G \mid \psi(g) = 1_H\}$.

2.1 Predicate Encryption (PE)

An encryption scheme for a predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$ over a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ consists of four PPT algorithms which satisfy a correctness condition defined below.

- $\text{Setup}(1^\lambda, \text{par}) \rightarrow (\text{MPK}, \text{MSK})$. The Setup algorithm takes as input the unary representation of the security parameter λ and some additional parameters par . It outputs a master public key

MPK and a master secret key MSK. The output of Setup defines a number $N \in \mathbb{N}$ (perhaps implicitly), and κ is set to (N, par) .

- $\text{Encrypt}(\text{MPK}, x, m) \rightarrow \text{CT}$. The encryption algorithm takes public parameters MPK, an $x \in \mathcal{X}_\kappa$ and an $m \in \mathcal{M}_\lambda$ as inputs, and outputs a ciphertext CT.
- $\text{KeyGen}(\text{MPK}, \text{MSK}, y) \rightarrow \text{SK}$. The key generation algorithm takes as input the public parameters MPK, the master secret key MSK and a $y \in \mathcal{Y}_\kappa$, and outputs a secret key SK.
- $\text{Decrypt}(\text{MPK}, \text{SK}, \text{CT}) \rightarrow m'$. The decryption algorithm takes as input the public parameters MPK, a secret key SK and a ciphertext CT, and outputs a message $m' \in \mathcal{M}_\lambda$.

Correctness: For all par , MPK and MSK output by $\text{Setup}(1^\lambda, \text{par})$, $m \in \mathcal{M}_\lambda$, $x \in \mathcal{X}_\kappa$ and $y \in \mathcal{Y}_\kappa$ such that $P_\kappa(x, y) = 1$, if

$$\text{CT} \leftarrow \text{Encrypt}(\text{MPK}, x, m) \quad \text{SK} \leftarrow \text{KeyGen}(\text{MPK}, \text{MSK}, y),$$

then

$$\Pr[\text{Decrypt}(\text{MPK}, \text{CT}, \text{SK}) \neq m] \leq \text{negl}(\lambda),$$

where the probability is over the random coin tosses of Encrypt, KeyGen and Decrypt.

Security: Let Π be an encryption scheme for a predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$ over a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$. Consider the following experiment $\text{Expt}_{\mathcal{A}, \Pi}^{(b)}$ (λ, par) between an adversary \mathcal{A} and a challenger Chl for $b \in \{0, 1\}$ when both are given input 1^λ and par :

1. **Setup:** Chl runs $\text{Setup}(1^\lambda, \text{par})$ to obtain MPK and MSK. It gives MPK to \mathcal{A} .
2. **Query:** \mathcal{A} issues a key query by sending $y \in \mathcal{Y}_\kappa$ to Chl, and obtains $\text{SK} \leftarrow \text{KeyGen}(\text{MPK}, \text{MSK}, y)$ in response. This step can be repeated any number of times \mathcal{A} desires.
3. **Challenge:** \mathcal{A} sends two messages $m_0, m_1 \in \mathcal{M}_\lambda$ and an $x \in \mathcal{X}_\kappa$ to Chl, and gets $\text{CT} \leftarrow \text{Encrypt}(\text{MPK}, x, m_b)$ as the challenge ciphertext.
4. **Query:** This step is identical to step 2.

At the end of the experiment, \mathcal{A} outputs a bit which is defined to be the output of the experiment. We call an adversary admissible if for every $y \in \mathcal{Y}_\kappa$ queried in steps 2 and 4, $P_\kappa(x, y) = 0$. This prevents \mathcal{A} from succeeding in the experiment simply by decrypting CT.

Definition 1. An encryption scheme Π is adaptively or fully secure for a predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$ if for every PPT admissible adversary \mathcal{A} and every par ,

$$|\Pr[\text{Expt}_{\mathcal{A}, \Pi}^{(0)}(\lambda, \text{par}) = 1] - \Pr[\text{Expt}_{\mathcal{A}, \Pi}^{(1)}(\lambda, \text{par}) = 1]| \leq \text{negl}(\lambda),$$

where the probabilities are taken over the coin tosses of \mathcal{A} and Chl. On the other hand, Π is semi-adaptively secure if the above condition is satisfied w.r.t. to a modified experiment where \mathcal{A} provides $x \in \mathcal{X}_\kappa$ to Chl right after the setup phase (instead of the challenge phase), i.e., before it starts querying [CW14b].

3 Pair encoding schemes

The notion of pair encoding schemes (PES) was introduced by Attrapadung [Att14]. Our definition of this scheme is slightly different from the one given by [Att14] in that we place a restriction on the structure. Though the latter definition is more general, we believe that our formulation mirrors the concrete design of such schemes more closely. In particular, all the constructions of pair encoding schemes given in [Att14] fit into our framework without any changes.

We first present the definition given by Attrapadung and discuss the restrictions we impose afterwards. A pair encoding scheme for a predicate family $P_\kappa : \mathcal{X}_\kappa \times \mathcal{Y}_\kappa \rightarrow \{0, 1\}$ indexed by $\kappa = (N, \text{par})$ consists of four polynomial-time *deterministic* algorithms which satisfy a correctness condition as defined below.

- $\text{Param}(\text{par}) \rightarrow n$. The Param algorithm takes the parameters par as input, and outputs a positive integer $n \in \mathbb{N}$ which specifies the number of common variables shared by the following two algorithms. Let $\mathbf{b} := (b_1, b_2, \dots, b_n)$ denote the common variables.
- $\text{EncC}(x, N) \rightarrow (\mathbf{c} := (c_1, c_2, \dots, c_{w_1}); w_2)$. The EncC algorithm takes an $N \in \mathbb{N}$ and an $x \in \mathcal{X}_{(N, \text{par})}$ as inputs, and outputs a sequence of w_1 polynomials c_1, c_2, \dots, c_{w_1} with coefficients in \mathbb{Z}_N and a $w_2 \in \mathbb{N}$. Every polynomial c_ℓ is a linear combination of monomials of the form $s, s_i, sb_j, s_i b_j$ in variables $s, s_1, s_2, \dots, s_{w_2}$ and b_1, \dots, b_n . More formally, for $\ell \in [1, w_1]$,

$$c_\ell := \zeta_\ell s + \sum_{i \in [1, w_2]} \eta_{\ell, i} s_i + \sum_{j \in [1, n]} \theta_{\ell, j} s b_j + \sum_{i \in [1, w_2], j \in [1, n]} \vartheta_{\ell, i, j} s_i b_j,$$

where $\zeta_\ell, \eta_{\ell, i}, \theta_{\ell, j}, \vartheta_{\ell, i, j} \in \mathbb{Z}_N$ are constants which define c_ℓ .

- $\text{EncK}(y, N) \rightarrow (\mathbf{k} := (k_1, k_2, \dots, k_{m_1}); m_2)$. The EncK algorithm takes an $N \in \mathbb{N}$ and a $y \in \mathcal{Y}_{(N, \text{par})}$ as inputs, and outputs a sequence of m_1 polynomials k_1, k_2, \dots, k_{m_1} with coefficients in \mathbb{Z}_N and an $m_2 \in \mathbb{N}$. Every polynomial k_t is a linear combination of monomials of the form $\alpha, r_{i'}, r_{i'} b_j$ in variables $\alpha, r_1, r_2, \dots, r_{m_2}$ and b_1, \dots, b_n . More formally, for $t \in [1, m_1]$,

$$k_t := \tau_t \alpha + \sum_{i' \in [1, m_2]} \nu_{t, i'} r_{i'} + \sum_{i' \in [1, m_2], j \in [1, n]} \phi_{t, i', j} r_{i'} b_j,$$

where $\tau_t, \nu_{t, i'}, \phi_{t, i', j} \in \mathbb{Z}_N$ are constants which define k_t .

- $\text{Pair}(x, y, N) \rightarrow \mathbf{E}$. The Pair algorithm takes an $N \in \mathbb{N}$, an $x \in \mathcal{X}_{(N, \text{par})}$ and a $y \in \mathcal{Y}_{(N, \text{par})}$ as inputs, and outputs a matrix $\mathbf{E} \in \mathbb{Z}_N^{m_1 \times w_1}$.

Correctness: A pair encoding scheme is correct if for every $\kappa = (N, \text{par})$, $x \in \mathcal{X}_\kappa$ and $y \in \mathcal{Y}_\kappa$ such that $P_\kappa(x, y) = 1$, the following holds symbolically

$$\mathbf{kEc}^T = \sum_{\substack{t \in [1, m_1], \\ \ell \in [1, w_1]}} E_{t, \ell} k_t c_\ell = \alpha s.$$

Structural restrictions. We impose an additional restriction on the form of \mathbf{E} . Essentially this says that if k_t has a monomial of the form $r_{i'} b_{j'}$ and a c_ℓ has a monomial of the form $s b_j$ or $s_i b_j$ then $E_{t, \ell}$ must be 0. One can easily verify that *every* pair encoding scheme given in [Att14] (as well as the new one we propose) satisfies this.

Moreover, we can show that given the constraint on \mathbf{E} , we can assume w.l.o.g. that the set of polynomials output by EncC and EncK have a fairly restricted structure. In simple words, if a polynomial contains the monomial $s_i b_j$ (resp. $r_{i'} b_j$), then there must exist a polynomial which *only* contains the monomial s_i (resp. $r_{i'}$). Further, there exists a polynomial with only the monomial s .

In Appendix A, we formally show that for any pair encoding which satisfies the restriction on \mathbf{E} , there is a corresponding one in which the polynomials output by EncC and EncK have the additional structure described before, and this correspondence preserves all of the security properties defined in [Att14]. For the rest of this work then, we will assume that all pair encodings satisfy the properties listed above.

3.1 Security

Attrapadung provided two security notions for pair encoding schemes: perfect and computational. As discussed in Section 1, in this paper we focus on perfect security, which is the information theoretic property, for which we propose a relaxation. First, we restate here the original security definition given by Attrapadung (which is referred to as *perfectly master-key hiding* in his paper).

Definition 2 (Perfect security [Att14]). *A pair encoding scheme $(\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$ for a predicate family P_κ is perfectly secure if for every $\kappa = (N, \text{par})$, $x \in \mathcal{X}_\kappa$ and $y \in \mathcal{Y}_\kappa$ such that $P_\kappa(x, y) = 0$,*

$$(\mathbf{c}(s, \mathbf{b}), \mathbf{k}(0, \mathbf{r}, \mathbf{b})) \equiv (\mathbf{c}(s, \mathbf{b}), \mathbf{k}(\alpha, \mathbf{r}, \mathbf{b})), \quad (1)$$

where $\mathbf{s} \leftarrow_R \mathbb{Z}_N^{w_2+1}$, $\mathbf{b} \leftarrow_R \mathbb{Z}_N^n$, $\mathbf{r} \leftarrow_R \mathbb{Z}_N^{m_2}$ and $\alpha \leftarrow_R \mathbb{Z}_N$.

We propose a new relaxed notion of perfect security that allows more flexibility in the design of pair encoding schemes. Very roughly, this property will allow us to add noise gradually to the parameters used in the key, as long as this noise is not detectable given the relevant part of the key and the ciphertext. The goal is to eventually add sufficient noise to completely hide the master secret. Towards this, we define a new *randomized* polynomial-time sampling algorithm for pair encoding schemes. While the algorithms above are used in the generic construction, the Samp algorithm described below will be used in the security proof.

- $\text{Samp}(d, x, y, N) \rightarrow (\mathbf{b}_d := (b_{d,1}, b_{d,2}, \dots, b_{d,n}))$. This algorithm takes a $d \in [1, m_2]$, an $N \in \mathbb{N}$, an $x \in \mathcal{X}_{(N, \text{par})}$, and a $y \in \mathcal{Y}_{(N, \text{par})}$ as inputs, and outputs a sequence of n numbers in \mathbb{Z}_N . We require that the probability of this algorithm producing $(u \cdot b_{d,1}, u \cdot b_{d,2}, \dots, u \cdot b_{d,n})$ as output is equal to the probability that it produces $(b_{d,1}, b_{d,2}, \dots, b_{d,n})$ as output, for any $u \in \mathbb{Z}_N^*$.

Jumping ahead, the dependence of Samp on its inputs will play a crucial role in the proof of security of our generic construction. We will see that if Samp doesn't depend on x , then we can prove our construction to be *fully* secure. But in case it does, we can only prove *semi-adaptive* security.

Recall that EncK on input y and N produces a sequence of polynomials $\mathbf{k}(\alpha, \mathbf{r}, \mathbf{b})$ with coefficients in \mathbb{Z}_N , where every polynomial is a linear combination of monomials of the form $\alpha, r_{i'}, r_{i'} b_j$ in variables $\alpha, r_1, r_2, \dots, r_{m_2}$ and b_1, \dots, b_n . In the following we use $\mathbf{k}_d(\alpha, r_d, \mathbf{b})$, for $d \in [1, m_2]$, to denote the polynomials in \mathbf{k} obtained by setting all the variables in $\{r_1, r_2, \dots, r_{m_2}\}$ except r_d to 0. We are now ready to define our new notion of perfect security.

Definition 3 (Relaxed perfect security). *A pair encoding scheme $\Gamma = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$ for a predicate family P_κ is relaxed perfectly secure if there exists a PPT algorithm Samp (as defined above) such that for every par , $x \in \mathcal{X}_\kappa$ and $y \in \mathcal{Y}_\kappa$ such that $P_\kappa(x, y) = 0$, and every $d \in [1, m_2]$:*

$$\{\mathbf{c}(s, \mathbf{b}), \mathbf{k}_d(0, r_d, \mathbf{b})\}_{N \in \mathbb{N}} \cong \{\mathbf{c}(s, \mathbf{b}), \mathbf{k}_d(0, r_d, \mathbf{b} + \mathbf{b}_d)\}_{N \in \mathbb{N}}, \quad (2)$$

where $\mathbf{s} \leftarrow_R \mathbb{Z}_N^{w_2+1}$, $\mathbf{b} \leftarrow_R \mathbb{Z}_N^n$, $r_d \leftarrow_R \mathbb{Z}_N$, $\mathbf{b}_d \leftarrow \text{Samp}(d, x, y, N)$. Furthermore,

$$\left\{ \mathbf{c}(\mathbf{s}, \mathbf{b}), \sum_{d \in [1, m_2]} \mathbf{k}_d(0, r_d, \mathbf{b} + \mathbf{b}_d) \right\}_{N \in \mathbb{N}} \cong \left\{ \mathbf{c}(\mathbf{s}, \mathbf{b}), \sum_{d \in [1, m_2]} \mathbf{k}_d(\alpha, r_d, \mathbf{b} + \mathbf{b}_d) \right\}_{N \in \mathbb{N}}, \quad (3)$$

where $\mathbf{s} \leftarrow_R \mathbb{Z}_N^{w_2+1}$, $\mathbf{b} \leftarrow_R \mathbb{Z}_N^n$, $r_1, r_2, \dots, r_{m_2} \leftarrow_R \mathbb{Z}_N$, $\alpha \leftarrow_R \mathbb{Z}_N$, $\mathbf{b}_d \leftarrow \text{Samp}(d, x, y, N)$ for $d \in [1, m_2]$, and \cong denotes statistical indistinguishability. We say Γ satisfies strong relaxed perfect security if Samp does not depend on x .

Note that in equations (2) and (3), we have distribution ensembles indexed by N , unlike the definition of perfect security where we are dealing with only one distribution. We require that the ensembles are statistically indistinguishable from each other, which means that for large enough values of N , the statistical distance between the distributions is negligible.

We now show that any pair encoding scheme that is perfectly secure under the original definition is also secure under the stronger flavor of the relaxed definition.

Lemma 1. *Let $\Gamma = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$ be a pair encoding scheme. If Γ is perfectly secure (Definition 2), then Γ is also relaxed perfectly secure (Definition 3). Moreover, we can define a Samp algorithm for Γ that does not depend on the input x .*

Proof. For any pair encoding scheme Γ , define Samp to output a vector of zeroes on any input. With this definition, (2) is trivially satisfied for every $d \in [1, m_2]$, and the two distributions in (3) reduce to

$$\left\{ \mathbf{c}(\mathbf{s}, \mathbf{b}), \sum_{d \in [1, m_2]} \mathbf{k}_d(0, r_d, \mathbf{b}) \right\} \quad \text{and} \quad \left\{ \mathbf{c}(\mathbf{s}, \mathbf{b}), \sum_{d \in [1, m_2]} \mathbf{k}_d(\alpha, r_d, \mathbf{b}) \right\}. \quad (4)$$

Since Γ is perfectly secure, we know that if $\mathbf{s} \leftarrow_R \mathbb{Z}_N^{w_2+1}$, $\mathbf{b} \leftarrow_R \mathbb{Z}_N^n$, $\mathbf{r} \leftarrow_R \mathbb{Z}_N^{m_2}$ and $\alpha \leftarrow_R \mathbb{Z}_N$, then

$$\{\mathbf{c}(\mathbf{s}, \mathbf{b}), \mathbf{k}(0, \mathbf{r}, \mathbf{b})\} \equiv \{\mathbf{c}(\mathbf{s}, \mathbf{b}), \mathbf{k}(\alpha, \mathbf{r}, \mathbf{b})\}.$$

We can replace $\mathbf{k}(\alpha, \mathbf{r}, \mathbf{b})$ with $\mathbf{k}(m_2\alpha, \mathbf{r}, \mathbf{b})$ in the above without changing the joint distribution. Now, observe that $\mathbf{k}(0, \mathbf{r}, \mathbf{b}) = \sum_{d \in [1, m_2]} \mathbf{k}_d(0, r_d, \mathbf{b})$ and $\mathbf{k}(m_2\alpha, \mathbf{r}, \mathbf{b}) = \sum_{d \in [1, m_2]} \mathbf{k}_d(\alpha, r_d, \mathbf{b})$ symbolically. Therefore, the two distributions in (4) are identical. \square

4 Dual System Groups

Our construction of predicate encryption schemes from pair encodings is based on dual system groups (DSG), introduced by Chen and Wee [CW14a] in a recent work. Our formulation of DSG, given below, can be seen as a generalization of theirs. However, as we will show, both their instantiations satisfy the new properties without making any changes.

A dual system group is parameterized by a security parameter λ and a number n . It consists of six PPT algorithms as described below.

4.1 Syntax

- $\text{SampP}(1^\lambda, 1^n)$: On input 1^λ and 1^n , SampP outputs public parameters PP and secret parameters SP , which have the following properties:

- PP contains a triple of groups $(\mathbb{G}, \mathbb{H}, \mathbb{G}_T)$ and a non-degenerate bilinear map $e : \mathbb{G} \times \mathbb{H} \rightarrow \mathbb{G}_T$, a homomorphism μ from \mathbb{H} to \mathbb{G}_T , along with some additional parameters used by SampG , SampH . Given PP, we know the exponent of group \mathbb{H} and how to sample uniformly from it. Let $N = \exp(\mathbb{H})$ (see Section 2). We require that N is a product of distinct primes of $\Theta(\lambda)$ bits.
 - SP contains $\tilde{h} \in \mathbb{H}$ (where $\tilde{h} \neq 1_{\mathbb{H}}$) along with additional parameters used by $\overline{\text{SampG}}$ and $\overline{\text{SampH}}$.
- SampGT takes an element in the image of μ and outputs another element from \mathbb{G}_T .
 - SampG and SampH take PP as input and output a vector of $n + 1$ elements from \mathbb{G} and \mathbb{H} respectively.
 - $\overline{\text{SampG}}$ and $\overline{\text{SampH}}$ take both PP and SP as inputs and output a vector of $n + 1$ elements from \mathbb{G} and \mathbb{H} respectively.

4.2 Properties

We require that all the properties below hold for every PP and SP output by SampP . Let SampG_0 be the algorithm that outputs only the first element of SampG . Analogously, SampH_0 , $\overline{\text{SampG}}_0$ and $\overline{\text{SampH}}_0$ can be defined. A dual system group is *correct* if it satisfies the following two properties⁵:

Projective: For all $h \in \mathbb{H}$ and coin tosses σ , $\text{SampGT}(\mu(h); \sigma) = e(\text{SampG}_0(\text{PP}; \sigma), h)$.

Associative: If (g_0, g_1, \dots, g_n) and (h_0, h_1, \dots, h_n) are samples from $\text{SampG}(\text{PP})$ and $\text{SampH}(\text{PP})$ respectively, then for all $i \in [1, n]$, $e(g_0, h_i) = e(g_i, h_0)$.

For *security* we require the following three properties to hold:

Orthogonality: $\tilde{h} \in \text{Kernel}(\mu)$, i.e., $\mu(\tilde{h}) = 1_{\mathbb{G}_T}$.

Non-degeneracy:

1. $\overline{\text{SampH}}_0(\text{PP}, \text{SP}) \cong \tilde{h}^\delta$, where $\delta \leftarrow_R \mathbb{Z}_N$.
2. $\exists \tilde{g} \in \mathbb{G}$ s.t. $\overline{\text{SampG}}_0(\text{PP}, \text{SP}) \cong \tilde{g}^\alpha$, where $\alpha \leftarrow_R \mathbb{Z}_N$.
3. For all $\hat{g}_0 \leftarrow \overline{\text{SampG}}_0(\text{PP}, \text{SP})$, $e(\hat{g}_0, \tilde{h})^\beta$ is uniformly distributed over \mathbb{G}_T , where $\beta \leftarrow_R \mathbb{Z}_N$.

(Here \cong denotes statistical indistinguishability.)

Remark 1. In [CW14a], the non-degeneracy property is defined in a slightly different way. First, they require that for all $\hat{h}_0 \leftarrow \overline{\text{SampH}}_0(\text{PP}, \text{SP})$, \tilde{h} lies in the group generated by \hat{h}_0 , instead of the first point above. And secondly, they do not have any constraint on the output of $\overline{\text{SampG}}_0(\text{PP}, \text{SP})$ like in the second point above. The third property, though, is also present in their definition⁶.

Indistinguishability. For two (positive) polynomials $\text{poly}_1(\cdot)$ and $\text{poly}_2(\cdot)$, define $\mathbf{G}, \mathbf{H}, \hat{\mathbf{G}}, \hat{\mathbf{H}}, \hat{\mathbf{G}}', \hat{\mathbf{H}}'$ as follows:

$$(\text{PP}, \text{SP}) \leftarrow \text{SampP}(1^\lambda, 1^n); \quad \gamma_1, \gamma_2, \dots, \gamma_n \leftarrow_R \mathbb{Z}_N;$$

$$\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{\text{poly}_1(\lambda)} \leftarrow \text{SampG}(\text{PP}); \quad \mathbf{G} := (\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{\text{poly}_1(\lambda)});$$

⁵Note that we have omitted the \mathbb{H} -subgroup property. It is required to construct encryption schemes with key delegation like HIBE. We do not use this property in our constructions.

⁶In the composite-order instantiation of [CW14a], this property holds only in a computational sense.

$$\begin{aligned}
& \mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{\text{poly}_2(\lambda)} \leftarrow \text{SampH}(\text{PP}); \mathbf{H} := (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{\text{poly}_2(\lambda)}); \\
& \forall i \in [1, \text{poly}_1(\lambda)], \quad \hat{\mathbf{g}}_i := (\hat{g}_{i,0}, \dots) \leftarrow \overline{\text{SampG}}(\text{PP}, \text{SP}); \quad \hat{\mathbf{g}}'_i := (1, \hat{g}_{i,0}^{\gamma_1}, \hat{g}_{i,0}^{\gamma_2}, \dots, \hat{g}_{i,0}^{\gamma_n}) \\
& \forall j \in [1, \text{poly}_2(\lambda)], \quad \hat{\mathbf{h}}_j := (\hat{h}_{j,0}, \dots) \leftarrow \overline{\text{SampH}}(\text{PP}, \text{SP}); \quad \hat{\mathbf{h}}'_j := (1, \hat{h}_{j,0}^{\gamma_1}, \hat{h}_{j,0}^{\gamma_2}, \dots, \hat{h}_{j,0}^{\gamma_n}) \\
& \hat{\mathbf{G}} := (\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \dots, \hat{\mathbf{g}}_{\text{poly}_1(\lambda)}); \hat{\mathbf{H}} := (\hat{\mathbf{h}}_1, \hat{\mathbf{h}}_2, \dots, \hat{\mathbf{h}}_{\text{poly}_2(\lambda)}); \\
& \hat{\mathbf{G}}' := (\hat{\mathbf{g}}'_1, \hat{\mathbf{g}}'_2, \dots, \hat{\mathbf{g}}'_{\text{poly}_1(\lambda)}); \hat{\mathbf{H}}' := (\hat{\mathbf{h}}'_1, \hat{\mathbf{h}}'_2, \dots, \hat{\mathbf{h}}'_{\text{poly}_2(\lambda)}).
\end{aligned}$$

We call a dual system group *Left Subgroup Indistinguishable* (LSI), *Right Subgroup Indistinguishable* (RSI) and *Parameter hiding* (PH) if for all polynomials $\text{poly}_1(\cdot)$ and $\text{poly}_2(\cdot)$,

$$\{\text{PP}, \mathbf{G}\} \approx \{\text{PP}, \mathbf{G} \cdot \hat{\mathbf{G}}\}, \quad (5)$$

$$\{\text{PP}, \tilde{h}, \mathbf{G} \cdot \hat{\mathbf{G}}, \mathbf{H}\} \approx \{\text{PP}, \tilde{h}, \mathbf{G} \cdot \hat{\mathbf{G}}, \mathbf{H} \cdot \hat{\mathbf{H}}\}, \text{ and} \quad (6)$$

$$\{\text{PP}, \tilde{h}, \hat{\mathbf{G}}, \hat{\mathbf{H}}\} \equiv \{\text{PP}, \tilde{h}, \hat{\mathbf{G}}' \cdot \hat{\mathbf{G}}, \hat{\mathbf{H}} \cdot \hat{\mathbf{H}}'\} \quad (7)$$

hold respectively. Observe that the two distributions in (5) and (6) are computationally indistinguishable, while the two distributions in (7) are identical.

Instantiations of DSG. The three indistinguishability properties defined above are generalizations of the corresponding ones in Chen and Wee[CW14a]. In Appendix B, we show that the two instantiations of DSG – in composite-order groups under the subgroup decision assumption and in prime-order groups under the decisional linear assumption (d -LIN) – given by [CW14a] satisfy our generalized indistinguishability properties as well as our new definition of non-degeneracy.

Remark 2. In the prime-order instantiation of dual system groups under the d -LIN assumption given by [CW14a], an element from groups \mathbb{G} or \mathbb{H} is represented by $d + 1$ elements from a source prime-order group (an element from \mathbb{G}_T is mapped to just one element of a target prime-order group). Now, suppose we have an encryption scheme in dual system groups where the ciphertext/key consists of elements from \mathbb{G} or \mathbb{H} (and possibly an element from \mathbb{G}_T). Then, a concrete instantiation in prime-order groups would only double the size of ciphertext/key, if we make the SXDH assumption (special case of d -LIN with $d = 1$), and only triple it if we make the DLIN assumption (special case of d -LIN with $d = 2$).

5 Predicate encryption from pair encodings

In this section, we show how to construct a predicate encryption scheme $\Pi_P = (\text{Setup}, \text{Encrypt}, \text{KeyGen}, \text{Decrypt})$ for any predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$ for which we have a pair encoding scheme $\Gamma_P = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$, using dual system groups. The message space for Π_P would be the target group in DSG. Recall that κ specifies a number $N \in \mathbb{N}$ and some additional parameters par .

- $\text{Setup}(1^\lambda, \text{par})$: First run $\text{Param}(\text{par})$ to obtain n , then run $\text{SampP}(1^\lambda, 1^n)$ to obtain PP and SP . Recall that given PP , we know the exponent of group \mathbb{H} and can sample uniformly from it. Output

$$\text{MSK} \leftarrow_R \mathbb{H} \quad \text{MPK} := (\text{PP}, \mu(\text{MSK})).$$

Set $N = \exp(\mathbb{H})$ and $\kappa = (N, \text{par})$.

- $\text{Encrypt}(\text{MPK}, x, m)$: On input an $x \in \mathcal{X}_\kappa$ and an $m \in \mathbb{G}_T$, run $\text{EncC}(x, N)$ to obtain a sequence of w_1 polynomials $(c_1, c_2, \dots, c_{w_1})$ and a $w_2 \in \mathbb{N}$. Draw $w_2 + 1$ samples from SampG :

$$(g_{0,0}, \dots, g_{0,n}) \leftarrow \text{SampG}(\text{PP}; \sigma)$$

$$(g_{1,0}, \dots, g_{1,n}) \leftarrow \text{SampG}(\text{PP}), \dots, (g_{w_2,0}, \dots, g_{w_2,n}) \leftarrow \text{SampG}(\text{PP}),$$

where σ denotes the coin tosses used in drawing the first sample from SampG .

Recall that the polynomial c_ℓ is given by

$$\zeta_\ell s + \sum_{i \in [1, w_2]} \eta_{\ell, i} s_i + \sum_{j \in [1, n]} \theta_{\ell, j} s b_j + \sum_{i \in [1, w_2], j \in [1, n]} \vartheta_{\ell, i, j} s_i b_j,$$

where $\zeta_\ell, \eta_{\ell, i}, \theta_{\ell, j}, \vartheta_{\ell, i, j} \in \mathbb{Z}_N$ are constants. Output $\text{CT} := (\text{CT}_1, \dots, \text{CT}_{w_1}, \text{CT}_{w_1+1})$ as the encryption of m under x where

$$\text{CT}_\ell := g_{0,0}^{\zeta_\ell} \cdot \prod_{i \in [1, w_2]} g_{i,0}^{\eta_{\ell, i}} \cdot \prod_{j \in [1, n]} g_{0,j}^{\theta_{\ell, j}} \cdot \prod_{i \in [1, w_2], j \in [1, n]} g_{i,j}^{\vartheta_{\ell, i, j}}$$

for $\ell \in [1, w_1]$ and $\text{CT}_{w_1+1} := m \cdot \text{SampGT}(\mu(\text{MSK}); \sigma)$. Notice that the monomials $s, s_i, s b_j$, and $s_i b_j$ are mapped to group elements $g_{0,0}, g_{i,0}, g_{0,j}$, and $g_{i,j}$, respectively.

- $\text{KeyGen}(\text{MPK}, \text{MSK}, y)$: On input a $y \in \mathcal{Y}_\kappa$, run $\text{EncK}(y, N)$ to obtain a sequence of m_1 polynomials $(k_1, k_2, \dots, k_{m_1})$ and an $m_2 \in \mathbb{N}$. Draw m_2 samples from SampH :

$$(h_{1,0}, \dots, h_{1,n}) \leftarrow \text{SampH}(\text{PP}), \dots, (h_{m_2,0}, \dots, h_{m_2,n}) \leftarrow \text{SampH}(\text{PP}).$$

Output the key as $\text{SK} := (\text{SK}_1, \text{SK}_2, \dots, \text{SK}_{m_1})$ where for $t \in [1, m_1]$

$$\text{SK}_t := \text{MSK}^{\tau_t} \cdot \prod_{i' \in [1, m_2]} h_{i',0}^{v_{t,i'}} \cdot \prod_{i' \in [1, m_2], j \in [1, n]} h_{i',j}^{\phi_{t,i',j}}.$$

In this case, the variables $\alpha, r_{i'}$, and $r_{i'} b_j$ are mapped to $\text{MSK}, h_{i',0}$, and $h_{i',j}$, respectively.

- $\text{Decrypt}(\text{MPK}, \text{SK}_y, \text{CT}_x)$: On input $\text{SK}_y := (\text{SK}_1, \text{SK}_2, \dots, \text{SK}_{m_1})$ and $\text{CT}_x := (\text{CT}_1, \dots, \text{CT}_{w_1+1})$, run $\text{Pair}(x, y, N)$ to obtain an $m_1 \times w_1$ matrix \mathbf{E} . Output

$$\text{CT}_{w_1+1} \cdot \left[\prod_{t \in [1, m_1], \ell \in [1, w_1]} e(\text{CT}_\ell, \text{SK}_t^{E_{t,\ell}}) \right]^{-1}.$$

Correctness (Sketch). We know that if $P_\kappa(x, y) = 1$, then $\sum_{t \in [1, m_1], \ell \in [1, w_1]} E_{t,\ell} k_t c_\ell = \alpha s$. Consider two polynomials k_t and c_ℓ . When these polynomials are multiplied together, no two monomials – one from k_t and one from c_ℓ – combine to give the same monomial in the product polynomial $k_t c_\ell$, except when

- s is multiplied with $r_{i'} b_j$ and $s b_j$ is multiplied with $r_{i'}$, or
- s_i is multiplied with $r_{i'} b_j$ and $s_i b_j$ is multiplied with $r_{i'}$,

because of the restriction on the form of \mathbf{E} . Now, s is mapped to $g_{0,0}$, $r_{i'}b_j$ is mapped to $h_{i',j}$, sb_j is mapped to $g_{0,j}$ and $r_{i'}$ is mapped to $h_{i',0}$. By the associativity property of dual system groups, we know that $e(g_{0,0}, h_{i',j}) = e(g_{0,j}, h_{i',0})$. Further, we mapped s_i to $g_{i,0}$ and $s_i b_j$ to $g_{i,j}$, and associativity guarantees that $e(g_{i,0}, h_{i',j}) = e(g_{i,j}, h_{i',0})$. Therefore, from the observations above, it follows that

$$\prod_{t \in [1, m_1], \ell \in [1, w_1]} e(\text{CT}_\ell, \text{SK}_t^{E_{t,\ell}}) = e(g_{0,0}, \text{MSK}).$$

Finally, by projective property we know that $e(g_{0,0}, \text{MSK}) = \text{SampGT}(\mu(\text{MSK}); \sigma)$.

Remark 3 (Preserving size). Observe that the output of `Encrypt` consists of $w_1 + 1$ elements, w_1 from \mathbb{G} and 1 from \mathbb{G}_T , where w_1 is the number of polynomials output by `EncC`. Further, any key has the same number of elements from \mathbb{H} as the number of polynomials output by `EncK`. Hence, in particular, if w_1 (resp. m_1) is a constant then ciphertexts (resp. keys) are also of constant size, in terms of dual system group elements. Further, if we instantiate dual system groups in prime-order groups under SXDH or DLIN assumption, then the ciphertexts (resp. keys) would still be of constant size (see Remark 2.)

6 Proof of security

In this section, we show that the encryption scheme Π_P constructed for a predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$ in the previous section is secure using the properties of dual system groups and relaxed perfect security of pair encoding schemes. More formally, we prove the following theorem.

Theorem 1. *For any predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$, if $\Gamma_P = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$ is a **relaxed perfectly secure** pair encoding scheme, then the encryption scheme $\Pi_P = (\text{Setup}, \text{Encrypt}, \text{KeyGen}, \text{Decrypt})$ constructed in Section 5 (using Γ_P) is **semi-adaptively secure**. Furthermore, if the algorithm `Samp` does not depend on input x , then Π_P is **fully secure** (see Definition 1).*

Using Lemma 1, a corollary of the above theorem is that:

Corollary 2. *For any predicate family $P = \{P_\kappa\}_{\kappa \in \mathbb{N}^c}$, if $\Gamma_P = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair}, \text{Samp})$ is a **perfectly secure** pair encoding scheme, then the encryption scheme $\Pi_P = (\text{Setup}, \text{Encrypt}, \text{KeyGen}, \text{Decrypt})$ constructed in Section 5 (using Γ_P) is **fully secure**.*

Recall that dual system groups can be instantiated in prime-order groups under the d -LIN assumption. Together with the above corollary, this gives a useful and interesting result:

Corollary 3. *Every perfectly secure pair encoding scheme proposed by Attrapadung [Att14] has a fully secure predicate encryption scheme in prime order groups under the d -LIN assumption.*

The rest of this section is devoted to the proof of Theorem 1. We first define auxiliary algorithms for encryption and key generation.

- $\overline{\text{Encrypt}}(\text{PP}, x, m; (\mathbf{g}'_0, \mathbf{g}'_1, \dots, \mathbf{g}'_{w_2}), \text{MSK})$: This algorithm is the same as `Encrypt` except that it uses the input $\mathbf{g}'_i \in \mathbb{G}^{n+1}$ instead of choosing samples \mathbf{g}_i from `SampG` for $i \in [0, w_2]$, and sets $\text{CT}_{w_1+1} := m \cdot e(\mathbf{g}'_{0,0}, \text{MSK})$, where $\mathbf{g}'_{0,0}$ is the first element of the vector \mathbf{g}'_0 .
- $\overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}'_1, \dots, \mathbf{h}'_{m_2}))$: This algorithm is the same as `KeyGen` except that it uses \mathbf{h}'_i instead of the samples \mathbf{h}_i from `SampH` for $i \in [1, m_2]$.

Using these algorithms, we define alternate forms for the ciphertext and master secret key:

Table 1: Various types of keys

Type of key	Inputs to $\overline{\text{KeyGen}}$ (besides PP and y)
Normal	MSK; $(\mathbf{h}_1, \dots, \mathbf{h}_{m_2})$
ρ -Intermediate-1	MSK; $(\mathbf{h}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{z_{\rho-1}}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})$
ρ -Intermediate-2	MSK; $(\mathbf{h}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{z_{\rho-1}}, \boxed{\mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{z_\rho}}, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})$
ρ -Intermediate-3	MSK; $(\mathbf{h}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{z_{\rho-1}}, \boxed{\mathbf{h}_\rho \cdot \tilde{h}^{z_\rho}}, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})$
Pseudo-normal noisy	MSK; $(\mathbf{h}_1 \cdot \hat{\mathbf{h}}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{m_2} \cdot \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{z_{m_2}})$
Pseudo-SF noisy	$\boxed{\overline{\text{MSK}}}$; $(\mathbf{h}_1 \cdot \hat{\mathbf{h}}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{m_2} \cdot \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{z_{m_2}})$
SF noisy	$\overline{\text{MSK}}$; $(\mathbf{h}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{m_2} \cdot \tilde{h}^{z_{m_2}})$
ρ -SF-intermediate-1	$\overline{\text{MSK}}$; $(\mathbf{h}_1, \dots, \mathbf{h}_{\rho-1}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{z_\rho}, \mathbf{h}_{\rho+1} \cdot \tilde{h}^{z_{\rho+1}}, \dots, \mathbf{h}_{m_2} \cdot \tilde{h}^{z_{m_2}})$
ρ -SF-intermediate-2	$\overline{\text{MSK}}$; $(\mathbf{h}_1, \dots, \mathbf{h}_{\rho-1}, \boxed{\mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho}, \mathbf{h}_{\rho+1} \cdot \tilde{h}^{z_{\rho+1}}, \dots, \mathbf{h}_{m_2} \cdot \tilde{h}^{z_{m_2}})$
ρ -SF-intermediate-3	$\overline{\text{MSK}}$; $(\mathbf{h}_1, \dots, \mathbf{h}_{\rho-1}, \boxed{\mathbf{h}_\rho}, \mathbf{h}_{\rho+1} \cdot \tilde{h}^{z_{\rho+1}}, \dots, \mathbf{h}_{m_2} \cdot \tilde{h}^{z_{m_2}})$
SF	$\overline{\text{MSK}}$; $(\mathbf{h}_1, \dots, \mathbf{h}_{m_2})$

- *Semi-functional master secret key* is defined to be $\overline{\text{MSK}} := \text{MSK} \cdot \tilde{h}^\beta$ where $\beta \leftarrow_R \mathbb{Z}_N$.
- *Semi-functional ciphertext* is given by $\overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK})$ where $\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{w_2} \leftarrow \text{SampG}(\text{PP})$, $\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \dots, \hat{\mathbf{g}}_{w_2} \leftarrow \overline{\text{SampG}}(\text{PP}, \text{SP})$, $\mathbf{G} := (\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{w_2})$, and $\hat{\mathbf{G}} := (\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \dots, \hat{\mathbf{g}}_{w_2})$. Observe that $\overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G}, \text{MSK})$ is identically distributed to $\text{Encrypt}(\text{MPK}, x, m)$ – the normal ciphertext – by the projective property of dual system groups.

Table 1 defines various forms of keys for $\rho \in [1, m_2]$ and the inputs that need to be passed to $\overline{\text{KeyGen}}$ (besides PP and y) in order to generate them. Intermediate-3 and SF-intermediate-3 keys are also defined for $\rho = 0$ (SF stands for semi-functional). In the table, $\mathbf{h}_1, \dots, \mathbf{h}_{m_2} \leftarrow \text{SampH}(\text{PP})$, $\hat{\mathbf{h}}_1, \dots, \hat{\mathbf{h}}_{m_2} \leftarrow \overline{\text{SampH}}(\text{PP}, \text{SP})$, and $\mathbf{z}_d := (1, z_{d,1}, \dots, z_{d,n})$, where $(z_{d,1}, \dots, z_{d,n}) \leftarrow \text{Samp}(d, x, y, N)$ for all $d \in [1, m_2]$. For convenience in the following, we define a slightly modified form of Samp, called $\overline{\text{Samp}}$, which just prepends 1 to the output of Samp. Note that 0-Intermediate-3 is distributed identically to a normal key and 0-SF-intermediate-3 is distributed identically to a SF noisy key. Since we have many forms of keys, (where appropriate) we use a box to highlight the part of a key which is different from the previous key.

Proof structure: The novelty in our proof is that instead of working at the level of a key, we work at the level of samples that form the key. Let ξ denote the number of queries made by the adversary, and let y_φ denote the φ th query for $\varphi \in [1, \xi]$. Further, let $m_{2,\varphi}$ be the second output of $\text{EncK}(y_\varphi, N)$. We define the following hybrids for $\varphi \in [1, \xi]$ and $\rho \in [1, m_{2,\varphi}]$ (fix any $b \in \{0, 1\}$).

- Hyb_0 : This is the real security game $\text{Expt}_{\mathcal{A}, \Pi_P}^{(b)}(\lambda, \text{par})$ described in Section 2.1.
- Hyb_1 : This game is same as the above except that the ciphertext is semi-functional.
- $\text{Hyb}_{2,\varphi,i,\rho}$ for $i \in \{1, 2, 3\}$: This game is same as the above except that the first $\varphi - 1$ keys are semi-functional, φ th key is of the form ρ -intermediate- i , and rest of the keys are normal.

<i>Indistinguishability</i>	<i>Properties needed</i>	<i>Proof</i>
$\text{Hyb}_0 \approx \text{Hyb}_1$	left subgroup indistinguishability	Lemma 12
$\text{Hyb}_{2,\varphi,3,\rho-1} \approx \text{Hyb}_{2,\varphi,1,\rho}$	right subgroup indistinguishability	Lemma 13
$\text{Hyb}_{2,\varphi,1,\rho} \cong \text{Hyb}_{2,\varphi,2,\rho}$	non-degeneracy, parameter-hiding, RPS (2)	Lemma 2
$\text{Hyb}_{2,\varphi,2,\rho} \approx \text{Hyb}_{2,\varphi,3,\rho}$	right subgroup indistinguishability	similar to Lemma 13
$\text{Hyb}_{2,\varphi,3,m_{2,\varphi}} \approx \text{Hyb}_{2,\varphi,4}$	right subgroup indistinguishability	similar to Lemma 13
$\text{Hyb}_{2,\varphi,4} \cong \text{Hyb}_{2,\varphi,5}$	non-degeneracy, parameter-hiding, RPS (3)	Lemma 3
$\text{Hyb}_{2,\varphi,5} \approx \text{Hyb}_{2,\varphi,6}$	right subgroup indistinguishability	similar to Lemma 13
$\text{Hyb}_{2,\varphi,9,\rho-1} \approx \text{Hyb}_{2,\varphi,7,\rho}$	right subgroup indistinguishability	similar to Lemma 13
$\text{Hyb}_{2,\varphi,7,\rho} \cong \text{Hyb}_{2,\varphi,8,\rho}$	non-degeneracy, parameter-hiding, RPS (2)	similar to Lemma 2
$\text{Hyb}_{2,\varphi,8,\rho} \approx \text{Hyb}_{2,\varphi,9,\rho}$	right subgroup indistinguishability	similar to Lemma 13
$\text{Hyb}_{2,\xi,9,m_{2,\xi}} \cong \text{Hyb}_3$	projective, orthogonality, non-degeneracy	Lemma 14

Table 2: An outline of the proof structure. After proving a lemma for a certain pair of hybrids, we discuss how the proof can be modified to show indistinguishability of other related pairs. In the above, RPS is a shorthand for relaxed perfect security.

- $\text{Hyb}_{2,\varphi,4}$: This game is same as the above except that the φ th key is Pseudo-normal noisy.
- $\text{Hyb}_{2,\varphi,5}$: This game is same as the above except that the φ th key is Pseudo-SF noisy.
- $\text{Hyb}_{2,\varphi,6}$: This game is same as the above except that the φ th key is SF noisy.
- $\text{Hyb}_{2,\varphi,i,\rho}$ for $i \in \{7, 8, 9\}$: This game is same as the above except that the φ th key is of the form ρ -SF-intermediate- $(i - 6)$.
- Hyb_3 : This game is same as $\text{Hyb}_{2,\xi,9,m_{2,\xi}}$ except that the ciphertext is a semi-functional encryption of a random message in \mathbb{G}_T .

Our goal is to show that Hyb_0 and Hyb_3 are computationally indistinguishable from each other, for both values of the bit b used by ChI in the security game $\text{Expt}_{\mathcal{A}, \Pi_P}^{(b)}(\lambda, \text{par})$. Since ChI encrypts a random message in Hyb_3 , there would be no way for a PPT adversary to tell whether m_0 or m_1 was encrypted. This would imply that Π_P is a secure encryption scheme.

Our proof proceeds as follows. We first show that Hyb_0 and Hyb_1 are computationally indistinguishable due to the left subgroup indistinguishability (LSI) property of dual system groups; this takes the ciphertext from normal to semi-functional space (the form of the ciphertext doesn't change after this step). After that, we take the keys one by one from normal to semi-functional space by going through a series of hybrids. We show that $\text{Hyb}_{2,1,3,0}$ (or, equivalently, Hyb_1) is computationally indistinguishable from $\text{Hyb}_{2,1,9,m_{2,1}}$ by following the steps shown in Table 2 for $\varphi = 1$; this makes the first key semi-functional while keeping the rest of the keys unchanged. Then, we show that $\text{Hyb}_{2,2,3,0}$ (or, equivalently, $\text{Hyb}_{2,1,9,m_{2,1}}$) is computationally indistinguishable from $\text{Hyb}_{2,2,9,m_{2,2}}$ by once again following the steps shown in Table 2, but now for $\varphi = 2$; as a result, the second key also moves into the semi-functional space. We continue in the same fashion till all the keys are in the semi-functional space, i.e., we are in the hybrid $\text{Hyb}_{2,\xi,9,m_{2,\xi}}$. The last step of the proof is to show that $\text{Hyb}_{2,\xi,9,m_{2,\xi}}$ and Hyb_3 are statistically close to each other.

We formally prove the indistinguishability of hybrids that require relaxed perfect security, our new information-theoretic notion of security, in Lemma 2 and 3 below, but defer the proof of Lemma 12, 13, and 14 to Appendix C because they follow directly from the properties of dual system groups in a manner similar to Chen and Wee’s security proof for HIBE [CW14a]. Though Lemma 13 shows that $\text{Hyb}_{2,\varphi,3,\rho-1}$ and $\text{Hyb}_{2,\varphi,1,\rho}$ are computationally indistinguishable, we discuss immediately afterwards how the proof can be modified to show indistinguishability of other related pairs of hybrids. We do the same with other lemmas too.

Remark 4 (Full vs. semi-adaptive security.). In transitioning from $\text{Hyb}_{2,\varphi,1,\rho}$ to $\text{Hyb}_{2,\varphi,2,\rho}$ in Lemma 2, we add randomness using the algorithm $\overline{\text{Samp}}$ to the ρ -th sample of the φ -th key. Observe that if $\overline{\text{Samp}}$ depends on input x , then this transition can only take place if x is known *before* any key queries are issued. Therefore, in this case, we can prove semi-adaptive security. On the other hand, if $\overline{\text{Samp}}$ does not depend on x , then we get full security (and as shown in Lemma 1, this is the case for all of the perfectly secure pair encoding schemes of [Att14]).

Remark 5 (Perfectly secure encodings). Recall from the proof of Lemma 1 that for any perfectly secure pair encoding scheme, we can define a dummy sampling algorithm that always outputs a vector of 0s. When this is the case, the security proof can be considerably simplified: we could directly go from Hyb_1 to $\text{Hyb}_{2,\varphi,4}$ and also from $\text{Hyb}_{2,\varphi,5}$ to $\text{Hyb}_{2,\varphi,9,m_2,\varphi}$ using right subgroup indistinguishability.

Remark 6 (Cost of our reduction). There are many complex predicates for which we do not know any perfectly secure pair encoding schemes. But if one can design a scheme that is relaxed perfectly secure, then we show that an encryption scheme can be derived from it, which is secure under standard assumptions. The reduction cost of our security proof, however, is higher than usual: if an adversary makes ξ queries and m_2 is the maximum number of samples used in any key, then the cost is $O(\xi \cdot m_2)$. For instance, this cost only depends on the number of pre-challenge queries in the case of Attrapadung’s computationally secure encodings (Theorem 1 in [Att14]). Note, however, that computational security of the encoding itself is proved under q -type assumptions.

Lemma 2. For every $\varphi \in [1, \xi]$ and $\rho \in [1, m_{2,\varphi}]$, $\text{Hyb}_{2,\varphi,1,\rho} \cong \text{Hyb}_{2,\varphi,2,\rho}$.

Proof. Given PP, MSK and \tilde{h} , one can generate MPK and every key except the φ th (because in order to generate this key and the ciphertext, we need to be able to sample from $\overline{\text{SampH}}$ and $\overline{\text{SampG}}$, for which secret parameters SP are required). Hence, it suffices to show that the following two distributions are statistically close (for clarity, we omit φ in the following):

$$\begin{aligned} & \{ \text{PP, MSK, } \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}), \\ & \quad \overline{\text{KeyGen}}(\text{PP, MSK}, y; (\mathbf{h}_1 \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{\mathbf{z}^{\rho-1}}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})) \}, \\ & \{ \text{PP, MSK, } \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}), \\ & \quad \overline{\text{KeyGen}}(\text{PP, MSK}, y; (\mathbf{h}_1 \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{\mathbf{z}^{\rho-1}}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{\mathbf{z}^\rho}, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})) \}. \end{aligned}$$

But observe that:

$$\begin{aligned} \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}) &= \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G}, \text{MSK}) \cdot \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}}, \text{MSK}), \\ \overline{\text{KeyGen}}(\text{PP, MSK}, y; (\mathbf{h}_1 \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{\mathbf{z}^{\rho-1}}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})) \\ &= \overline{\text{KeyGen}}(\text{PP, MSK}, y; (\mathbf{h}_1 \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{\mathbf{z}^{\rho-1}}, \mathbf{h}_\rho, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})) \\ & \quad \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho, 1, \dots, 1)), \end{aligned}$$

$$\begin{aligned}
& \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{z_{\rho-1}}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{z_\rho}, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})) \\
&= \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}_1 \cdot \tilde{h}^{z_1}, \dots, \mathbf{h}_{\rho-1} \cdot \tilde{h}^{z_{\rho-1}}, \mathbf{h}_\rho, \mathbf{h}_{\rho+1}, \dots, \mathbf{h}_{m_2})) \cdot \\
& \quad \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{z_\rho}, 1, \dots, 1)),
\end{aligned}$$

because of the way Encrypt and KeyGen are defined and bilinearity of e (see the construction in Section 5). The first component on the right hand side of each of the above equations can be generated given PP, MSK and \tilde{h} . Hence, we only need to focus on the second components, i.e., it is enough to show that the following two distributions are statistically close:

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho, 1, \dots, 1))\}, \quad (8)$$

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{z_\rho}, 1, \dots, 1))\}. \quad (9)$$

Let us focus on the first distribution between the two above. By the parameter-hiding property of dual system groups we know that $\{\text{PP}, \tilde{h}, \hat{\mathbf{G}}, \hat{\mathbf{h}}_\rho\}$ and $\{\text{PP}, \tilde{h}, \hat{\mathbf{G}} \cdot \hat{\mathbf{G}}', \hat{\mathbf{h}}_\rho \cdot \hat{\mathbf{h}}'_\rho\}$ are identically distributed. Hence (8) is identically distributed to

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}} \cdot \hat{\mathbf{G}}', \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho \cdot \hat{\mathbf{h}}'_\rho, 1, \dots, 1))\}. \quad (10)$$

Let $\hat{\text{CT}} := (\hat{\text{CT}}_1, \dots, \hat{\text{CT}}_{w_1+1})$ and $\hat{\text{SK}} := (\hat{\text{SK}}_1, \dots, \hat{\text{SK}}_{m_1})$ denote the output of $\overline{\text{Encrypt}}$ and $\overline{\text{KeyGen}}$ respectively. We know that for $\ell \in [1, w_1]$,

$$\hat{\text{CT}}_\ell = \hat{g}_{0,0}^{\zeta_\ell} \cdot \prod_{i \in [1, w_2]} \hat{g}_{i,0}^{\eta_{\ell,i}} \cdot \prod_{j \in [1, n]} (\hat{g}_{0,j} \cdot \hat{g}_{0,0}^{\gamma_j})^{\theta_{\ell,j}} \cdot \prod_{i \in [1, w_2], j \in [1, n]} (\hat{g}_{i,j} \cdot \hat{g}_{i,0}^{\gamma_j})^{\vartheta_{\ell,i,j}},$$

where $(\hat{g}_{i,0}, \dots, \hat{g}_{i,n}) \leftarrow \overline{\text{SampG}}(\text{PP}, \text{SP})$ for $i \in [0, w_2]$ and $\gamma_1, \dots, \gamma_n \leftarrow_R \mathbb{Z}_N$. Also, $\hat{\text{CT}}_{w_1+1} = e(\hat{g}_{0,0}, \text{MSK})$. Using the non-degeneracy property of dual system groups, we can write $\hat{g}_{0,0}$ and $\hat{g}_{i,0}$ as \tilde{g}^δ and \tilde{g}^{δ_i} respectively, for $i \in [1, w_2]$, where $\delta, \delta_1, \dots, \delta_{w_2} \leftarrow_R \mathbb{Z}_N$. Then we consider $\hat{g}_{0,j}$ (and $\hat{g}_{i,j}$) for $j = 1, \dots, n$ to be values sampled from $\overline{\text{SampG}}$ conditioned on the value of $\hat{g}_{0,0}$ (resp. $\hat{g}_{i,0}$). (These values may not be efficiently sampleable.) Therefore, we have

$$\begin{aligned}
\hat{\text{CT}}_\ell &= \tilde{g}^{\zeta_\ell \delta + \sum_{i \in [1, w_2]} \eta_{\ell,i} \delta_i + \sum_{j \in [1, n]} \theta_{\ell,j} \delta \gamma_j + \sum_{i \in [1, w_2], j \in [1, n]} \vartheta_{\ell,i,j} \delta_i \gamma_j} \cdot \\
& \quad \prod_{j \in [1, n]} \hat{g}_{0,j}^{\theta_{\ell,j}} \cdot \prod_{i \in [1, w_2], j \in [1, n]} \hat{g}_{i,j}^{\vartheta_{\ell,i,j}} \quad (11)
\end{aligned}$$

Shifting our focus to the key, we know that its t th component is given by

$$\hat{\text{SK}}_t = \hat{h}_{\rho,0}^{v_{t,\rho}} \cdot \prod_{j \in [1, n]} (\hat{h}_{\rho,j} \cdot \hat{h}_{\rho,0}^{\gamma_j})^{\phi_{t,\rho,j}},$$

for $t \in [1, m_1]$, where $(\hat{h}_{\rho,0}, \dots, \hat{h}_{\rho,n}) \leftarrow \overline{\text{SampH}}(\text{PP}, \text{SP})$. Using non-degeneracy once again, we can write $\hat{h}_{\rho,0}$ as \tilde{h}^ω for an $\omega \leftarrow_R \mathbb{Z}_N$, and consider $\hat{h}_{\rho,j}$ for $j = 1, \dots, n$ to be sampled from $\overline{\text{SampH}}$ conditioned on the value of $\hat{h}_{\rho,0}$. Hence,

$$\hat{\text{SK}}_t = \tilde{h}^{v_{t,\rho} \omega + \sum_{j \in [1, n]} \phi_{t,\rho,j} \omega \gamma_j} \cdot \prod_{j \in [1, n]} \hat{h}_{\rho,j}^{\phi_{t,\rho,j}}. \quad (12)$$

Now, observe the superscripts of \tilde{g} and \tilde{h} in (11) and (12) respectively (over $\ell \in [1, w_1]$ and $t \in [1, m_1]$). We know that $\delta, \delta_1, \dots, \delta_{w_2}, \gamma_1, \dots, \gamma_n$ and ω are randomly chosen from \mathbb{Z}_N . Hence,

we can use the first property (2) of relaxed perfect security to add noise to the ρ -th sample used in the key. But the problem is that in any sample drawn from $\overline{\text{SampG}}$ and $\overline{\text{SampH}}$, elements of the sample may depend on each other. In particular $\hat{g}_{0,j}$ may reveal some information about δ , and similarly for $\hat{g}_{i,j}$ and for $\hat{h}_{\rho,j}$, so we must ensure that (2) applies even given this information. Recall the discussion on *structural restrictions* after the definition of pair encoding schemes. We know that if $\vartheta_{\ell,i,j} \neq 0$ for any $\ell \in [1, w_1]$ and $j \in [1, n]$ (otherwise, we don't need to worry about $\hat{g}_{i,j}$), then δ_i is an *explicit* part of the encoding output by EncC. Similarly, if $\phi_{t,\rho,j} \neq 0$ for any $t \in [1, m_1]$ and $j \in [1, n]$, then ω is an explicit part of the encoding output by EncK. Further, δ is always explicit. Therefore, given a sample from either of the distributions in (2), one can compute the first element of the samples from $\overline{\text{SampG}}$ and $\overline{\text{SampH}}$, and then draw rest of the elements conditioned on the first ones.

In a nutshell, we can apply (2) to conclude that the distribution

$$\{\text{PP}, \text{MSK}, \tilde{h}, (\hat{C}T_1, \dots, \hat{C}T_{w_1+1}), (\hat{S}K_1, \dots, \hat{S}K_{m_1})\}$$

is statistically close to

$$\{\text{PP}, \text{MSK}, \tilde{h}, (\hat{C}T_1, \dots, \hat{C}T_{w_1+1}), (\tilde{S}K_1, \dots, \tilde{S}K_{m_1})\},$$

where

$$\begin{aligned} \tilde{S}K_t &:= \tilde{h}^{v_{t,\rho}\omega + \sum_{j \in [1,n]} \phi_{t,\rho,j}\omega(\gamma_j + z_j)} \cdot \prod_{j \in [1,n]} \hat{h}_{\rho,j}^{\phi_{t,\rho,j}} \\ &= \tilde{h}^{v_{t,\rho}\omega + \sum_{j \in [1,n]} \phi_{t,\rho,j}\omega\gamma_j} \cdot \prod_{j \in [1,n]} (\hat{h}_{\rho,j} \cdot \tilde{h}^{\omega z_j})^{\phi_{t,\rho,j}}, \end{aligned}$$

for $t \in [1, m_1]$, and $\mathbf{z}_\rho = (z_1, \dots, z_n) \leftarrow \text{Samp}(\rho, x, y, N)$. We use the fact that δ is always explicit to generate the $w_1 + 1$ th component of the ciphertext.

Observe that the only difference between $\hat{S}K_t$ and $\tilde{S}K_t$ is that an extra $\tilde{h}^{\omega z_j}$ is multiplied with $\hat{h}_{\rho,j}$ in the latter case. Hence, the key $(\tilde{S}K_1, \dots, \tilde{S}K_{m_1})$ can be generated by giving $\hat{\mathbf{h}}_\rho \cdot \hat{\mathbf{h}}'_\rho \cdot \tilde{h}^{\mathbf{z}_\rho}$ as the ρ -th sample to $\overline{\text{KeyGen}}$ (\mathbf{z}_ρ has the same distribution as $\omega \cdot \mathbf{z}_\rho$ since $\omega \in \mathbb{Z}_N^*$ with high probability). Therefore, (10) is statistically close to

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}} \cdot \hat{\mathbf{G}}', \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho \cdot \hat{\mathbf{h}}'_\rho \cdot \tilde{h}^{\mathbf{z}_\rho}, 1, \dots, 1)).\}$$

Using parameter-hiding once again, we can show that the above distribution is identical to

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, 1, y; (1, \dots, 1, \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{\mathbf{z}_\rho}, 1, \dots, 1)),\}$$

which completes the proof. \square

The above proof can be easily adapted to show that $\text{Hyb}_{2,\varphi,7,\rho} \cong \text{Hyb}_{2,\varphi,8,\rho}$. In this case, we want that the two distributions

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y; (\mathbf{h}_1, \dots, \mathbf{h}_{\rho-1}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho \cdot \tilde{h}^{\mathbf{z}_\rho}, \mathbf{h}_{\rho+1} \cdot \tilde{h}^{\mathbf{z}_{\rho+1}}, \dots, \mathbf{h}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}}))\},$$

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y; (\mathbf{h}_1, \dots, \mathbf{h}_{\rho-1}, \mathbf{h}_\rho \cdot \hat{\mathbf{h}}_\rho, \mathbf{h}_{\rho+1} \cdot \tilde{h}^{\mathbf{z}_{\rho+1}}, \dots, \mathbf{h}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}}))\}.$$

are indistinguishable from each other. Observe that the only difference now is that we have $\overline{\text{MSK}}$ instead of MSK , and noise is present in the samples $\rho + 1, \dots, n$ instead of $1, \dots, \rho - 1$. So, we can split $\overline{\text{Encrypt}}$ and $\overline{\text{KeyGen}}$ in a way similar to the above proof, and once again it suffices to show that exactly the distributions in (8) and (9) are indistinguishable.

Lemma 3. For every $\varphi \in [1, \xi]$, $\text{Hyb}_{2,\varphi,4} \cong \text{Hyb}_{2,\varphi,5}$.

Proof. This proof proceeds in a manner similar to the proof of Lemma 2. To begin with, we observe as before that given PP , MSK and \tilde{h} , one can generate MPK and every key except the φ th (for clarity, we omit φ below). Hence, it suffices to show that the distribution

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}_1 \cdot \hat{\mathbf{h}}_1 \cdot \tilde{h}^{\mathbf{z}_1}, \dots, \mathbf{h}_{m_2} \cdot \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}}))\},$$

is statistically close to a distribution where MSK is replaced by $\overline{\text{MSK}}$, the semi-functional master secret key. Further,

$$\begin{aligned} \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK}) &= \overline{\text{Encrypt}}(\text{PP}, x, m; \mathbf{G}, \text{MSK}) \cdot \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}}, \text{MSK}), \\ \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}_1 \cdot \hat{\mathbf{h}}_1 \cdot \tilde{h}^{\mathbf{z}_1}, \dots, \mathbf{h}_{m_2} \cdot \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}})) &= \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}_1, \dots, \mathbf{h}_{m_2})) \cdot \overline{\text{KeyGen}}(\text{PP}, 1, y; (\hat{\mathbf{h}}_1 \cdot \tilde{h}^{\mathbf{z}_1}, \dots, \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}})), \\ \overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y; (\mathbf{h}_1 \cdot \hat{\mathbf{h}}_1 \cdot \tilde{h}^{\mathbf{z}_1}, \dots, \mathbf{h}_{m_2} \cdot \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}})) &= \overline{\text{KeyGen}}(\text{PP}, \tilde{\text{MSK}}, y; (\mathbf{h}_1(\hat{\mathbf{h}}_1 \cdot \tilde{h}^{\mathbf{z}_1}, \cdot), \dots, \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{\mathbf{z}_{m_2}})) \end{aligned}$$

Shifting our focus to the key, we know that its t th component is given by

$$\hat{\mathbf{SK}}_t = \prod_{i' \in [1, m_2]} \hat{h}_{i', 0}^{v_{t, i'}} \cdot \prod_{i' \in [1, m_2], j \in [1, n]} (\hat{h}_{i', j}^{\gamma_j} \cdot \hat{h}_{i', 0}^{\gamma_j} \cdot \tilde{h}^{z_{i', j}})^{\phi_{t, i', j}},$$

for $t \in [1, m_1]$, where $(\hat{h}_{i', 0}, \dots, \hat{h}_{i', n}) \leftarrow \overline{\text{SampH}}(\text{PP}, \text{SP})$ and $(z_{i', 1}, \dots, z_{i', n}) \leftarrow \text{Samp}(i', x, y, N)$ for $i' \in [1, m_2]$. Using non-degeneracy once again, we can write $\hat{h}_{i', 0}$ as $\tilde{h}^{\omega_{i'}}$ for an $\omega_{i'} \leftarrow_R \mathbb{Z}_N$. Hence,

$$\begin{aligned} \hat{\mathbf{SK}}_t &= \tilde{h}^{\sum_{i' \in [1, m_2]} [v_{t, i'} \omega_{i'} + \sum_{j \in [1, n]} (\phi_{t, i', j} \omega_{i'} \gamma_j + \phi_{t, i', j} z_{i', j})]} \cdot \prod_{i' \in [1, m_2], j \in [1, n]} \hat{h}_{i', j}^{\phi_{t, i', j}} \\ &= \tilde{h}^{\sum_{i' \in [1, m_2]} [v_{t, i'} \omega_{i'} + \sum_{j \in [1, n]} (\phi_{t, i', j} \omega_{i'} (\gamma_j + z_{i', j}))]} \cdot \prod_{i' \in [1, m_2], j \in [1, n]} \hat{h}_{i', j}^{\phi_{t, i', j}}, \end{aligned} \quad (17)$$

since the distribution of $(z_{i', 1}, \dots, z_{i', n})$ is statistically close to $(\omega_{i'} z_{i', 1}, \dots, \omega_{i'} z_{i', n})$ (with high probability $\omega_{i'} \in \mathbb{Z}_N^*$) for all $i' \in [1, m_2]$.

Now, observe the superscripts of \tilde{g} and \tilde{h} in (16) and (17) respectively (over $\ell \in [1, w_1]$ and $t \in [1, m_1]$). We know that $\delta, \delta_1, \dots, \delta_{w_2}, \gamma_1, \dots, \gamma_n$ and $\omega_1, \dots, \omega_{m_2}$ are randomly chosen from \mathbb{Z}_N . Hence, we can use the second property (3) of relaxed perfect security to add noise to the master secret key. (The dependencies between the elements of the samples drawn from $\overline{\text{SampG}}$ and $\overline{\text{SampH}}$ can be handled as in the previous proof.) Therefore, we have that the distribution

$$\{\text{PP}, \text{MSK}, \tilde{h}, (\hat{\mathbf{CT}}_1, \dots, \hat{\mathbf{CT}}_{w_1+1}), (\hat{\mathbf{SK}}_1, \dots, \hat{\mathbf{SK}}_{m_1})\}$$

is statistically close to

$$\{\text{PP}, \text{MSK}, \tilde{h}, (\hat{\mathbf{CT}}_1, \dots, \hat{\mathbf{CT}}_{w_1+1}), (\tilde{\mathbf{SK}}_1, \dots, \tilde{\mathbf{SK}}_{m_1})\},$$

where

$$\tilde{\mathbf{SK}}_t := \tilde{h}^{\tau_t \beta + \sum_{i' \in [1, m_2]} [v_{t, i'} \omega_{i'} + \sum_{j \in [1, n]} (\phi_{t, i', j} \omega_{i'} (\gamma_j + z_{i', j}))]} \cdot \prod_{i' \in [1, m_2], j \in [1, n]} \hat{h}_{i', j}^{\phi_{t, i', j}},$$

for $t \in [1, m_1]$, and $\beta \leftarrow_R \mathbb{Z}_N$. Observe that the only difference between $\hat{\mathbf{SK}}_t$ and $\tilde{\mathbf{SK}}_t$ is that an extra $\tau_t \beta$ is begin added to the exponent of \tilde{h} in the latter case. Hence, the key $(\tilde{\mathbf{SK}}_1, \dots, \tilde{\mathbf{SK}}_{m_1})$ can be generated by providing \tilde{h}^β as master secret key to $\overline{\text{KeyGen}}$. Therefore, (15) is statistically close to

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}} \cdot \hat{\mathbf{G}}', \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, \tilde{h}^\beta, y; (\hat{\mathbf{h}}_1 \cdot \hat{\mathbf{h}}'_1 \cdot \tilde{h}^{z_1}, \dots, \hat{\mathbf{h}}_{m_2} \cdot \hat{\mathbf{h}}'_{m_2} \cdot \tilde{h}^{z_{m_2}}))\}.$$

Using parameter-hiding once again, we can show that the above distribution is identical to

$$\{\text{PP}, \text{MSK}, \tilde{h}, \overline{\text{Encrypt}}(\text{PP}, x, 1; \hat{\mathbf{G}}, \text{MSK}), \overline{\text{KeyGen}}(\text{PP}, \tilde{h}^\beta, y; (\hat{\mathbf{h}}_1 \cdot \tilde{h}^{z_1}, \dots, \hat{\mathbf{h}}_{m_2} \cdot \tilde{h}^{z_{m_2}}))\},$$

which completes the proof. \square

7 Ciphertext-Policy ABE

In this section, we design a relaxed perfectly secure pair encoding scheme for Ciphertext-Policy Attribute Based Encryption (CP-ABE). The access policy is represented by a linear secret sharing (LSS) scheme (\mathbf{A}, π) , where \mathbf{A} is a matrix of size $n_1 \times n_2$ with entries in \mathbb{Z}_N and π is a mapping from $[1, n_1]$ to a universe of attributes \mathcal{U} . Let \mathbf{a}_i denote the i th row of \mathbf{A} for $i \in [1, n_1]$. Let $S \subseteq \mathcal{U}$ be a set of attributes and $\Upsilon = \{i \mid i \in [1, n_1], \pi(i) \in S\}$ be the indices of rows in \mathbf{A} associated with S .

We say that the LSS scheme (\mathbf{A}, π) accepts S if $\mathbf{e} = (1, 0, \dots, 0)$ lies in the span of rows associated with S (otherwise the scheme rejects S). In other words, if S is acceptable, there exists constants $\{\varepsilon_i\}_{i \in \Upsilon}$ such that $\sum_{i \in \Upsilon} \varepsilon_i \mathbf{a}_i = \mathbf{e}$. (This set of constants can be easily computed given S .) An interesting property of LSS schemes that will be useful to us later in the proofs is that if (A, π) rejects S , then there must exist a vector $\mathbf{w} = (w_1, \dots, w_{n_2})$ such that $\langle \mathbf{w}, \mathbf{a}_i \rangle = 0$ for all $i \in \Upsilon$ but $\langle \mathbf{w}, \mathbf{e} \rangle = 1$. This, in particular, implies that $w_1 = 1$. (See [Bei11], Claim 2, for a proof of this and other properties below about secret sharing schemes.)

In order to share a secret $s \in \mathbb{Z}_N$, one picks $v_2, v_3, \dots, v_{n_1} \leftarrow_R \mathbb{Z}_N$, and outputs $\langle \mathbf{a}_i, \mathbf{v} \rangle$ as the i th share for $i \in [1, n_1]$, where $\mathbf{v} = (s, v_2, v_3, \dots, v_{n_1})$. This way of sharing a secret leads to two useful properties:

- **Correctness:** For every S accepted by (\mathbf{A}, π) , every secret $s \in \mathbb{Z}_N$ and any $v_2, v_3, \dots, v_{n_1} \in \mathbb{Z}_N$, $\sum_{i \in \Upsilon} \varepsilon_i \langle \mathbf{a}_i, \mathbf{v} \rangle = \langle \mathbf{v}, \sum_{i \in \Upsilon} \varepsilon_i \mathbf{a}_i \rangle = s$.
- **Privacy:** For every S rejected by (\mathbf{A}, π) , the distribution of $\{\langle \mathbf{a}_i, \mathbf{v} \rangle\}_{i \in \Upsilon}$ is independent of the secret s being shared.

The predicate family for CP-ABE is indexed by $\kappa = (N, n_1, n_2, \mathcal{U}, T)$. \mathcal{X}_κ is the set of all LSS schemes where the matrix is of size $n_1 \times n_2$ with entries in \mathbb{Z}_N and the mapping is from $[1, n_1]$ to \mathcal{U} . \mathcal{Y}_κ is given by the set $\{S \mid S \subseteq \mathcal{U}, |S| \leq T\}$. For all $x \in \mathcal{X}_\kappa$ and $y \in \mathcal{Y}_\kappa$, $P_\kappa(x, y) = 1$ if and only if x accepts y . It is clear from our definition of predicate family that there is a bound on the size of matrices and the number of attributes associated with a key. But there are no other restrictions: the size of attribute universe \mathcal{U} could be arbitrary and π need not be injective. Without loss of generality, we assume \mathcal{U} to be \mathbb{Z}_N .

We are now ready to design a relaxed perfectly secure pair encoding scheme $\Phi_{\text{cp-abe}} = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$ for the CP-ABE predicate family.

7.1 Pair Encoding Scheme

- $\text{Param}(\text{par}) \rightarrow n_1(n_2 + T + 1)$. Let $\mathbf{b} = (\{b_{i,j}\}_{i \in [1, n_1], j \in [1, n_2]}, \{b'_{i,t}\}_{i \in [1, n_1], t \in [0, T]})$.
- $\text{EncC}((A, \pi), N) \rightarrow \mathbf{c}(\mathbf{s}, \mathbf{b}) := (c_1, c_2)$ where

$$c_1 = s \quad c_2 = s \left(\sum_{\substack{i \in [1, n_1] \\ j \in [1, n_2]}} a_{i,j} b_{i,j} + \sum_{\substack{i \in [1, n_1] \\ t \in [0, T]}} \pi(i)^t b'_{i,t} \right),$$

and $\mathbf{s} = (s)$, and $a_{i,j}$ denotes the entry in the i th row and j th column of A .

- $\text{EncK}(S, N) \rightarrow \mathbf{k}(\alpha, \mathbf{r}, \mathbf{b}) := (\{k_{1,i}, k_{2,i,j}, k_{3,i,\ell,j}, k_{4,i,y}, k_{5,i,\ell,t}\}_{i, \ell \in [1, n_1], i \neq \ell, j \in [1, n_2], y \in S, t \in [0, T]})$ where

$$k_{1,i} = r_i \quad k_{2,i,j} = r_i b_{i,j} - v_j \quad k_{3,i,\ell,j} = r_i b_{\ell,j}$$

$$k_{4,i,y} = r_i \sum_{t \in [0,T]} y^t b'_{i,t} \quad k_{5,i,\ell,t} = r_i b'_{\ell,t}$$

and $\mathbf{r} = (r_1, r_2, \dots, r_{n_1}, v_2, \dots, v_{n_2})$ and $v_1 = \alpha$.

We informally discuss how to recover αs by combining the polynomials generated by EncC and EncK, with an intent to provide some intuition about the scheme, and defer a formal proof to Appendix D. We can think of v_2, v_3, \dots, v_{n_1} as the randomness picked in order to share $v_1 = \alpha$ according to the scheme (A, π) . Hence, if we find $\langle \mathbf{a}_i, \mathbf{v} \rangle$ for all $i \in \Upsilon$, we can recover α (ignore s for now). One could start out by multiplying $a_{i,j}$ by $k_{2,i,j}$ and summing over j , for an $i \in \Upsilon$. This does give $\sum_j a_{i,j} v_j$ but also produces an extra term $r_i \sum_j a_{i,j} b_{i,j}$ (ignore r_i for now). We could try to get rid of this term by using c_2 but the product $a_{i,j} b_{i,j}$ there is also summed over i (since we want EncC to produce a constant number of polynomials, we are forced to pack as much into one polynomial as possible). Fortunately, we have the polynomials $k_{3,i,\ell,j}$ for $\ell \neq i$. We can multiply these by $a_{\ell,j}$ and remove the unwanted $a_{i,j} b_{i,j}$ terms. But we are not done yet: we must also remove the term $\sum_{i,t} \pi(i)^t b'_{i,t}$ left in the mix because we used c_2 . If $\pi(i) \in S$, then this is easy: use $k_{4,i,\pi(i)}$ to remove $\sum_t \pi(i)^t b'_{i,t}$, and $k_{5,i,\ell,t} \cdot \pi(\ell)^t$ to remove the rest. However, if $\pi(i) \notin S$, there is no way to do this.

7.2 Relaxed Perfect Security

We now prove that the pair encoding scheme $\Phi_{\text{cp-abe}}$ designed above is relaxed perfectly secure (Definition 3). Towards this, we first define a sampling algorithm Samp as follows. On input an $i \in [1, n_1]$, $(A, \pi) \in \mathcal{X}_\kappa$, $S \in \mathcal{Y}_\kappa$ and N , Samp checks whether $\pi(i) \notin S$. If yes, it picks elements $\hat{b}_{i,1}, \hat{b}_{i,2}, \dots, \hat{b}_{i,n_2}$ independently and uniformly from \mathbb{Z}_N ; otherwise it picks them uniformly but with the constraint that $\sum_{j \in [1, n_2]} a_{i,j} \hat{b}_{i,j} = 0$. Samp outputs

$$\hat{\mathbf{b}}_i := \underbrace{(0, \dots, \dots, 0)}_{(i-1)n_2}, \hat{b}_{i,1}, \hat{b}_{i,2}, \dots, \hat{b}_{i,n_2}, \underbrace{(0, \dots, \dots, \dots, 0)}_{(n_1-i)n_2 + n_1(T+1)}. \quad (18)$$

Observe that the output of Samp depends on (A, π) , the input to EncC. Hence, this sampling algorithm would lead to a semi-adaptively secure scheme.

We consider only those $N \in \mathbb{N}$ which are a product of distinct primes of $\Theta(\lambda)$ bits. This is sufficient for our purposes because the Setup algorithm of the generic construction in Section 5 defines N of exactly this form. We first show that for all $i \in [1, n_1]$ and $N \in \mathbb{N}$,

$$(\mathbf{c}(\mathbf{s}, \mathbf{b}), \mathbf{k}_i(0, r_i, \mathbf{b})) \equiv (\mathbf{c}(\mathbf{s}, \mathbf{b}), \mathbf{k}_i(0, r_i, \mathbf{b} + \hat{\mathbf{b}}_i)), \quad (19)$$

where $\mathbf{s} \leftarrow_R \mathbb{Z}_N^1$, $\mathbf{b} \leftarrow_R \mathbb{Z}_N^{n_2}$, $r_i \leftarrow_R \mathbb{Z}_N$, $\hat{\mathbf{b}}_i \leftarrow \text{Samp}(i, (A, \pi), S, N)$. Recall that \mathbf{k}_i denotes the polynomials in \mathbf{k} obtained by setting all the variables in $\mathbf{r} = (r_1, r_2, \dots, r_{n_1}, v_2, \dots, v_{n_2})$ except the i th to 0. For $i \in [n_1 + 1, n_1 + n_2 - 1]$, the only polynomial in \mathbf{k}_i is $-v_{i-n_1+1}$, or, more importantly, there is no monomial with any b . Hence, the equation above trivially holds for i in this range irrespective of what Samp outputs. (That is why we don't care about defining Samp's behavior on such inputs.)

Let us refer to the left and right distributions in Equation (19) as Δ_L and Δ_R respectively. Fix an arbitrary $i^* \in [1, n_1]$. By the definition of \mathbf{k}_{i^*} , we know that in these two distributions only those components of the key survive which have subscript i^* . Further, in the components $k_{2,i^*,1}, \dots, k_{2,i^*,n_2}$, the variables v_1, \dots, v_{n_2} are all set to 0. Now, focus on the distribution Δ_R . It is clear from Equation (18) that the added randomness $\hat{\mathbf{b}}_{i^*}$ affects only $k_{2,i^*,1}, \dots, k_{2,i^*,n_2}$ components. For $i \in [1, n_1]$ and

$j \in [1, n_2]$, let $\delta_{i,j} := b_{i,j}$ if $i \neq i^*$ and $\delta_{i^*,j} := b_{i^*,j} + \hat{b}_{i^*,j}$. Since $b_{i,j}$ are uniformly and independently distributed, so are $\delta_{i,j}$. The second component of ciphertext encoding, c_2 , can now be rewritten as

$$s \left(\sum_{\substack{i \in [1, n_1] \\ j \in [1, n_2]}} a_{i,j} \delta_{i,j} - \sum_{j \in [1, n_2]} a_{i^*,j} \hat{b}_{i^*,j} + \sum_{t \in [0, T]} \pi(i^*)^t b'_{i^*,t} + \sum_{\substack{i \in [1, n_1], i \neq i^* \\ t \in [0, T]}} \pi(i)^t b'_{i,t} \right).$$

Observe that the only difference between Δ_L and Δ_R is that in the latter case there is an additional term $\text{rand} := \sum_{j \in [1, n_2]} a_{i^*,j} \hat{b}_{i^*,j}$ in c_2 . If $\pi(i^*) \in S$, then this term is 0 by our choice of Samp. On the other hand when $\pi(i^*) \notin S$, we show that $\sum_{t \in [0, T]} \pi(i^*)^t b'_{i^*,t}$ is an independent uniform random variable over \mathbb{Z}_N , and therefore, the additional term rand does not matter. Towards this, consider the polynomial $f(x) = b'_{i^*,T} \cdot x^T + b'_{i^*,T-1} \cdot x^{T-1} + \dots + b'_{i^*,0}$. Since $b'_{i^*,T}, \dots, b'_{i^*,0}$ are chosen at random, any $T + 1$ distinct points on $f(x)$ are uniformly distributed over \mathbb{Z}_N^{T+1} . The only components of the key which depend on $b'_{i^*,T}, \dots, b'_{i^*,0}$ are $\{k_{4,i^*,y}\}_{y \in S}$, which could also be rewritten as $\{r_{i^*} f(y)\}_{y \in S}$. There could be at most T such components because $|S| \leq T$. Therefore, $\sum_{t \in [0, T]} \pi(i^*)^t b'_{i^*,t} = f(\pi(i^*))$ is independently and uniformly distributed.

The second and last step in proving relaxed perfect security is to show that when (A, π) does not accept S , Equation (3) holds, i.e., for large enough values of N , the statistical distance between the distributions,

$$\left(\mathbf{c}(\mathbf{s}, \mathbf{b}), \sum_{i \in [1, n_1 + n_2 - 1]} \mathbf{k}_i(0, r_i, \mathbf{b} + \hat{\mathbf{b}}_i) \right) \text{ and } \left(\mathbf{c}(\mathbf{s}, \mathbf{b}), \sum_{i \in [1, n_1 + n_2 - 1]} \mathbf{k}_i(\alpha, r_i, \mathbf{b} + \hat{\mathbf{b}}_i) \right), \quad (20)$$

is negligible, where $\mathbf{s} \leftarrow_R \mathbb{Z}_N^1$, $\mathbf{b} \leftarrow_R \mathbb{Z}_N^n$, $\mathbf{r} \leftarrow_R \mathbb{Z}_N^{n_1 + n_2 - 1}$, $\alpha \leftarrow_R \mathbb{Z}_N$, and $\hat{\mathbf{b}}_i \leftarrow \text{Samp}(i, (A, \pi), S, N)$ for $i \in [1, n_1 + n_2 - 1]$. Let us denote the left and right distributions in Equation (20) above by Γ_L and Γ_R respectively. The second component of the key in these two distributions is given by

$$k_{2,i,j} = r_i b_{i,j} + r_i \hat{b}_{i,j} - v_j$$

for $i \in [1, n_1]$ and $j \in [1, n_2]$. The only difference between the distributions is in the components $k_{2,1,1}, \dots, k_{2,n_1,1}$. In the case of Γ_L , $v_1 = (n_1 + n_2 - 1)\alpha = 0$, while in the case of Γ_R , it is chosen independently and uniformly from \mathbb{Z}_N .

Let us focus on the distribution Γ_L . Recall that there exists a vector $\mathbf{w} = (w_1, \dots, w_{n_2})$ orthogonal to all the rows associated with S such that $w_1 = 1$. We claim that if we replace the variables $\hat{b}_{i,j}$ by $\hat{b}_{i,j} - r_i^{-1} w_j \alpha$, where $\alpha \leftarrow_R \mathbb{Z}_N$, then Γ_L is not affected. (With high probability $r_i \in \mathbb{Z}_N^*$, so r_i^{-1} exists.) If $\pi(i) \notin S$, we know that $\hat{b}_{i,1}, \hat{b}_{i,2}, \dots, \hat{b}_{i,n_2}$ are independently and uniformly distributed. Hence adding $-r_i^{-1} w_j \alpha$ has no effect on their joint distribution. On the other hand when $\pi(i) \in S$, $\hat{b}_{i,1}, \hat{b}_{i,2}, \dots, \hat{b}_{i,n_2}$ are uniformly chosen with the constraint that $\sum_{j \in [1, n_2]} a_{i,j} \hat{b}_{i,j} = 0$. Now, when $-r_i^{-1} w_j \alpha$ is added,

$$\sum_{j \in [1, n_2]} a_{i,j} (\hat{b}_{i,j} - r_i^{-1} w_j \alpha) = \sum_{j \in [1, n_2]} a_{i,j} \hat{b}_{i,j} - r_i^{-1} \alpha \sum_{j \in [1, n_2]} a_{i,j} w_j = 0$$

because \mathbf{w} is orthogonal to every \mathbf{a}_i such that $\pi(i) \in S$. Hence, the variables $\hat{b}_{i,1}, \hat{b}_{i,2}, \dots, \hat{b}_{i,n_2}$ still satisfy the constraint they did before.

After replacing $\hat{b}_{i,j}$ by $\hat{b}_{i,j} - r_i^{-1} w_j \alpha$, we have that $k_{2,i,j} = r_i b_{i,j} + r_i \hat{b}_{i,j} - w_j \alpha - v_j$ (where $v_1 = 0$). The final step in the proof is to replace the variables $w_1 \alpha, w_2 \alpha + v_2, \dots, w_{n_2} \alpha + v_{n_2}$ by $\alpha, v_2, \dots, v_{n_2}$. This does not affect Γ_L because v_2, \dots, v_{n_2} are picked independently and uniformly from \mathbb{Z}_N (and $w_1 = 1$). But now Γ_L is exactly the distribution Γ_R .

7.3 Instantiation: Constant-size ciphertext

We briefly comment about instantiating the pair encoding scheme $\Phi_{\text{cp-abe}} = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$. Using the generic method in Section 5, one can construct a predicate encryption scheme $\Pi_{\text{cp-abe}} = (\text{Setup}, \text{Encrypt}, \text{KeyGen}, \text{Decrypt})$ for CP-ABE using $\Phi_{\text{cp-abe}}$. According to Theorem 1, $\Pi_{\text{cp-abe}}$ is semi-adaptively secure because the Samp algorithm we defined in the previous subsection depends on the access structure. However, since EncC outputs only two polynomials, Encrypt outputs only two elements from \mathbb{G} (and one element from \mathbb{G}_T). Now, from Remark 2, it follows that one can design a concrete scheme for CP-ABE in prime-order groups where the ciphertext contains *only 4 group elements under the SXDH assumption*, and only 6 elements under the DLIN assumption (plus an additional element from the target group). Furthermore, only a constant number of pairing operations would be required to decrypt a ciphertext.

References

- [AHY15] Nuttapong Attrapadung, Goichiro Hanaoka, and Shota Yamada. Conversions among several classes of predicate encryption and applications to ABE with various compactness tradeoffs. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 575–601. Springer, Heidelberg, November / December 2015.
- [ALdP11] Nuttapong Attrapadung, Benoît Libert, and Elie de Panafieu. Expressive key-policy attribute-based encryption with constant-size ciphertexts. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 90–108. Springer, Heidelberg, March 2011.
- [APG⁺11] Joseph A. Akinyele, Matthew W. Pagano, Matthew D. Green, Christoph U. Lehmann, Zachary N. J. Peterson, and Aviel D. Rubin. Securing electronic medical records using attribute-based encryption on mobile devices. In *SPSM'11, ACM Workshop Security and Privacy in Smartphones and Mobile Devices 2011*, pages 75–86, 2011.
- [Att14] Nuttapong Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 557–577. Springer, Heidelberg, May 2014.
- [Att15] Nuttapong Attrapadung. Dual system encryption framework in prime-order groups. Cryptology ePrint Archive, Report 2015/390, 2015. <http://eprint.iacr.org/2015/390>.
- [BBS⁺09] Randolph Baden, Adam Bender, Neil Spring, Bobby Bhattacharjee, and Daniel Starin. Persona: an online social network with user-defined privacy. In *ACM SIGCOMM 2009 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 135–146, 2009.
- [Bei11] Amos Beimel. Secret-sharing schemes: A survey. In YeowMeng Chee, Zhenbo Guo, San Ling, Fengjing Shao, Yuansheng Tang, Huaxiong Wang, and Chaoping Xing, editors, *Coding and Cryptology*, volume 6639 of *Lecture Notes in Computer Science*, pages 11–46. Springer Berlin Heidelberg, 2011.

- [BRS13] Dan Boneh, Ananth Raghunathan, and Gil Segev. Function-private identity-based encryption: Hiding the function in functional encryption. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 461–478. Springer, Heidelberg, August 2013.
- [BSW07] John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy*, pages 321–334. IEEE Computer Society Press, May 2007.
- [BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Heidelberg, March 2011.
- [BW07] Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 535–554. Springer, Heidelberg, February 2007.
- [CGW15] Jie Chen, Romain Gay, and Hoeteck Wee. Improved dual system ABE in prime-order groups via predicate encodings. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 595–624. Springer, Heidelberg, April 2015.
- [CW13] Jie Chen and Hoeteck Wee. Fully, (almost) tightly secure IBE and dual system groups. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 435–460. Springer, Heidelberg, August 2013.
- [CW14a] Jie Chen and Hoeteck Wee. Dual system groups and its applications — compact HIBE and more. Cryptology ePrint Archive, Report 2014/265, 2014. <http://eprint.iacr.org/2014/265>.
- [CW14b] Jie Chen and Hoeteck Wee. Semi-adaptive attribute-based encryption and improved delegation for boolean formula. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 277–297. Springer, Heidelberg, September 2014.
- [EHK⁺13] Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, August 2013.
- [Fre10] David Mandell Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 44–61. Springer, Heidelberg, May 2010.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.
- [Gui13] Aurore Guillevic. Comparing the pairing efficiency over composite-order and prime-order elliptic curves. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *ACNS 13*, volume 7954 of *LNCS*, pages 357–372. Springer, Heidelberg, June 2013.

- [HHH⁺14] Gottfried Herold, Julia Hesse, Dennis Hofheinz, Carla Ràfols, and Andy Rupp. Polynomial spaces: A new framework for composite-to-prime-order transformations. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 261–279. Springer, Heidelberg, August 2014.
- [KSW08] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 146–162. Springer, Heidelberg, April 2008.
- [Lew12] Allison B. Lewko. Tools for simulating features of composite order bilinear groups in the prime order setting. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 318–335. Springer, Heidelberg, April 2012.
- [LOS⁺10] Allison B. Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 62–91. Springer, Heidelberg, May 2010.
- [LW11] Allison B. Lewko and Brent Waters. Unbounded HIBE and attribute-based encryption. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 547–567. Springer, Heidelberg, May 2011.
- [LW12] Allison B. Lewko and Brent Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 180–198. Springer, Heidelberg, August 2012.
- [OT10] Tatsuaki Okamoto and Katsuyuki Takashima. Fully secure functional encryption with general relations from the decisional linear assumption. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 191–208. Springer, Heidelberg, August 2010.
- [OT12] Tatsuaki Okamoto and Katsuyuki Takashima. Fully secure unbounded inner-product and attribute-based encryption. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 349–366. Springer, Heidelberg, December 2012.
- [PTMW06] Matthew Pirretti, Patrick Traynor, Patrick McDaniel, and Brent Waters. Secure attribute-based systems. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 99–112. ACM Press, October / November 2006.
- [RW13] Yannis Rouselakis and Brent Waters. Practical constructions and new proof methods for large universe attribute-based encryption. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 463–474. ACM Press, November 2013.
- [SRGS12] Nuno Santos, Rodrigo Rodrigues, Krishna P. Gummadi, and Stefan Saroiu. Policy-sealed data: A new abstraction for building trusted cloud services. In *USENIX Security Symposium 2012*, pages 175–188, 2012.
- [SSW09] Emily Shen, Elaine Shi, and Brent Waters. Predicate privacy in encryption systems. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 457–473. Springer, Heidelberg, March 2009.

- [SW05] Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
- [Tak14] Katsuyuki Takashima. Expressive attribute-based encryption with constant-size ciphertexts from the decisional linear assumption. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 298–317. Springer, Heidelberg, September 2014.
- [TBEM08] Patrick Traynor, Kevin R. B. Butler, William Enck, and Patrick McDaniel. Realizing massive-scale conditional access systems through attribute-based cryptosystems. In *NDSS 2008*. The Internet Society, February 2008.
- [Wat09] Brent Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 619–636. Springer, Heidelberg, August 2009.
- [Wee14] Hoeteck Wee. Dual system encryption via predicate encodings. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 616–637. Springer, Heidelberg, February 2014.
- [YAHK14] Shota Yamada, Nuttapong Attrapadung, Goichiro Hanaoka, and Noboru Kunihiro. A framework and compact constructions for non-monotonic attribute-based encryption. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 275–292. Springer, Heidelberg, March 2014.

A Pair encoding schemes

Recall that we consider polynomials of the form:

$$k_t := \tau_t \alpha + \sum_{i' \in [1, m_2]} v_{t, i'} r_{i'} + \sum_{i' \in [1, m_2], j \in [1, n]} \phi_{t, i', j} r_{i'} b_j,$$

$$c_\ell := \zeta_\ell s + \sum_{i \in [1, w_2]} \eta_{\ell, i} s_i + \sum_{j \in [1, n]} \theta_{\ell, j} s b_j + \sum_{i \in [1, w_2], j \in [1, n]} \vartheta_{\ell, i, j} s_i b_j.$$

Here we consider two properties. The first property says that in decryption, an $s_i b_j$ term is never paired with an $r_{i'} b_{j'}$ term. This property holds for all known pair encodings. The second property says that if $s_i b_j$ appears as a term in one of the polynomials in \mathbf{c} , then the polynomial s_i also appears in \mathbf{c} , and similarly for $r_{i'}$ (also, s is always a part of \mathbf{c}). We will show that w.l.o.g. we can assume that if the first property holds, then the second holds as well, in that we can always construct an equivalent scheme for which it does hold. This structure allows us to simplify our relaxed perfect security property for pair encodings.

Definition 4. *We say that a pair encoding is dual-system-group-compatible (DSG-compatible) if for all x, y, N , for all outputs of $\text{EncC}(x, N)$, $\text{EncK}(y, N)$, $\text{Pair}(x, y, N)$, and for all $t \in [1, m_1]$, $\ell \in [1, w_1]$, if there exist $i' \in [1, m_2]$, $j' \in [1, n]$ such that $\phi_{t, i', j'} \neq 0$ and there exist $i \in [1, w_2]$, $j \in [1, n]$ such that one of $\theta_{\ell, j}$ or $\vartheta_{\ell, i, j}$ is not zero, then $E_{t, \ell}$ must be 0.*

Definition 5. We say that a pair encoding has specified variables if for all x, y, N and for all outputs of $\text{EncC}(x, N)$, $\text{EncK}(y, N)$ it holds that

- $s \in \mathbf{c}$;
- for all $i \in [1, w_2]$, if $\vartheta_{\ell, i, j} \neq 0$ for any $\ell \in [1, w_1]$ and $j \in [1, n]$, then $s_i \in \mathbf{c}$; and,
- for all $i' \in [1, m_2]$, if $\phi_{t, i', j} \neq 0$ for any $t \in [1, m_1]$ and $j \in [1, n]$, then $r_{i'} \in \mathbf{k}$.

Lemma 4. We show that for any pair encoding which is DSG-compatible, we can construct an equivalent encoding which also has specified variables. Here by equivalent we mean that the number of polynomials produced by EncC and EncK together is increased by at most 1, and that perfect security and computational security defined by Attrapadung [Att14] are preserved.

Proof. First of all, we can assume without loss of generality that for all $x \in \mathcal{X}$, there is at least some $y \in \mathcal{Y}$ such that $P(x, y) = 1$. Let c_1, \dots, c_{w_1} be the polynomials output by $\text{EncC}(x, N)$, \mathbf{E} be the matrix output by $\text{Pair}(x, y, N)$, and τ_1, \dots, τ_t be the coefficients of α in the t polynomials output by $\text{EncK}(y, N)$. Then, if c_1, \dots, c_{w_1} are linearly combined using $\sum_t E_{t,1} \tau_t, \dots, \sum_t E_{t,w_1} \tau_t$, the result should just be s . (If there is any other term σ in the linear combination, then on combining polynomials produced by EncC and EncK according to \mathbf{E} , we will get an unwanted term $\alpha\sigma$.) Hence, we can make s an explicit part of the vector of polynomials \mathbf{c} without affecting any of the security properties. Note that if there exists *some* linear combination of c_1, \dots, c_{w_1} that gives s , then it can be found efficiently.

Next, we describe a process that iteratively removes any s_i which is not a part of \mathbf{c} , but for which there is a polynomial with a non-zero $s_i b_j$ term. Call a polynomial *simple* if it doesn't have any of the $s b_j$, $s_i b_j$, and $r_{i'} b_j$ terms (and non-simple otherwise). We consider two separate cases.

1. First, suppose there exists a simple polynomial $c_\ell = \zeta_\ell s + \sum_{i'} \eta_{\ell, i'} s_{i'}$ where the coefficient $\eta_{\ell, i}$ of s_i is not zero. In this case we will replace s_i with $1/\eta_{\ell, i}(s_i - (\zeta_\ell s + \sum_{i' \neq i} \eta_{\ell, i'} s_{i'}))$ every where in the encoding. Thus, in particular, c_ℓ becomes just s_i in the new encoding. It is easy to see that this does not affect DSG compatibility or correctness. Security is preserved because the distribution of *new* s_i is still uniform and independent of other variables.

Note that to carry out the aforementioned replacement, it is crucial that c_ℓ is simple, for otherwise, the substitute for s_i will have an $s b_{j'}$ or $s_{i'} b_{j'}$ term, and substitution in $s_i b_j$ will result in $s b_j b_{j'}$ or $s_{i'} b_j b_{j'}$.

We can repeat the above process as long as it applies. Each time we create a polynomial consisting of a single s_i (which makes s_i specified), and this polynomial is not affected by later steps. Further, since no new variables are introduced, there are at most w_2 iterations.

2. If the first condition is not satisfied for an s_i , then it must be the case that all polynomials that contain it are non-simple. Due to DSG-compatibility, we know that such polynomials cannot be paired with a polynomial containing $r_{i'} b_j$ for any i', j . Thus, s_i is never paired with $r_{i'} b_j$ for any i', j . Further, the coefficients of $s_i b_j \cdot r_{i'}$ and $s_i \cdot r_{i'} b_j$ must sum to 0 by correctness of encoding. Thus, we conclude that the coefficients of $s_i b_j \cdot r_{i'}$ sum to 0 for all i', j . So we can replace $s_i b_j$ with a new variable s'_j (for all j for which the coefficient of $s_i b_j$ is not zero in at least one polynomial) without affecting correctness. (Note that this transformation does increase the total number of variables, however it does not change the number of polynomials in \mathbf{c}, \mathbf{k} .) To see that perfect security is preserved, note that replacing $s_i b_j$ with a new variable has essentially the same result on the distributions in Definition 2 as adding a uniform random variable to both sides: if the original distributions were indistinguishable, the resulting ones

will be as well. Similarly, for Attrapadung’s selective and co-selective security, it can only decrease the adversary’s advantage.

Observe that this step is quite different from the first one. It removes all occurrences of $s_i b_j$ so that we don’t have to worry about the second condition in Definition 5 anymore. In the process, some new variables are introduced, but they only appear by themselves (i.e., never with a b_j). So, we need to repeat this process at most w_2 times.

We can apply the same two steps discussed above for the $r_{i'}$ variables, but a problem occurs when an $r_{i'}$ falls into the first category, i.e. there exists a simple k_t containing $r_{i'}$. In this case, when $r_{i'}$ is replaced with $1/v_{t,i'}(r_{i'} - (\tau_t \alpha + \sum_{i'' \neq i'} v_{t,i''} r_{i''}))$ in the term $r_{i'} b_j$, we get terms of the form αb_j (if $\tau_t \neq 0$), which is not allowed. Hence, for simplicity, we first remove α from all but one polynomial.

Suppose there exists a simple polynomial, say k_1 , containing α with coefficient $\tau_1 \neq 0$ (if not, then the problem described before does not arise). Replace any other polynomial k_t with $\tau_1 k_t - \tau_t k_1$, where τ_t is the coefficient of α in k_t . At the same time, change the pairing coefficient $E_{1,\ell}$ with $E_{1,\ell} + \tau_t/\tau_1 E_{t,\ell}$ and $E_{t,\ell}$ with $E_{t,\ell}/\tau_1$ for all ℓ and $t \neq 1$. It is easy to see that correctness still holds. Further, the modified encoding remains DSG-compatible. No extra variables or polynomials are generated during this process.

We now have only one polynomial $k_1 = \tau_1 \alpha + p$ containing α , where p does not have terms of the form $r_{i'} b_j$. When this polynomial is paired with different polynomials output by EncC, the result should just be $(1/\tau_1)s$, i.e. $\sum_{\ell} E_{1,\ell} c_{\ell} = (1/\tau_1)s$. Hence, we can as well pair k_1 with the new polynomial that has just s . (If this polynomial is denoted by c_0 , then $E_{1,0}$ will be set to $1/\tau_1$, while $E_{2,0}, \dots, E_{m_1,0}$ will remain 0.)

Suppose we now apply the two steps described before for handling s_i variables that do not satisfy the second condition to $r_{i'}$ variables that do not satisfy the third condition in Definition 5. If there is a simple polynomial besides k_1 that contains $r_{i'}$, then the first step will work, and if all polynomials containing $r_{i'}$ are non-simple, then the second step works. The only case left to analyze is when k_1 is the only simple polynomial $r_{i'}$ is a part of—all others are non-simple. We know that these non-simple polynomials cannot be paired with any polynomial containing $s b_j$ or $s_i b_j$, and that k_1 too is only being paired with s . Hence, $r_{i'}$ never gets paired with $s b_j$ or $s_i b_j$ for any i, j . Therefore, by making an argument similar to the second step, we can replace $r_{i'} b_j$ with a new variable $r'_{i'}$ without affecting correctness or any of the security properties. \square

B Dual System Groups

Chen and Wee instantiate dual system groups under the subgroup decision assumption in composite-order groups as well as the decisional linear assumption (d -LIN) in prime-order groups. We show that both these instantiations satisfy the generalized indistinguishability properties and the new non-degeneracy property we defined (first two conditions) below. (For the rest of the properties, the proofs given in [CW14a] carry over.)

Remark 7 (Sampling algorithms). In the two concrete constructions of dual system groups discussed below, the running time of the four sampling algorithms (SampG, SampH, SampG, SampH) depends linearly on the number of elements we require from a sample. This could significantly improve the efficiency of encryption schemes built on top of dual system groups. For example, if we need only the first and third elements from a sample of SampG (which consists of $n + 1$ elements), then we could just pass 1 and 3 to SampG (after modifying its definition suitably) and get the required elements, saving a considerable amount of time.

B.1 Composite-order construction

A composite-order bilinear group generator \mathcal{G} takes the security parameter λ as input and outputs $(N, G_N, G_T, g_1, g_2, g_3, e)$. G_N and G_T are two multiplicative cyclic groups of order $N = p_1 p_2 p_3$, where p_1, p_2 and p_3 are three distinct primes of $\Theta(\lambda)$ bits each. e is an efficiently computable non-degenerate bilinear map which maps two elements of G_N to an element of G_T . g_1, g_2 and g_3 denote the generators of G_{p_1}, G_{p_2} and G_{p_3} respectively, where for every divisor n of N , we use G_n to denote the subgroup of G_N of order n . We require that the following two subgroup decision assumptions hold with respect to \mathcal{G} .

Definition 6 (Assumption 1). *Consider the following distribution:*

$$\begin{aligned} (N, G_N, G_T, g_1, g_2, g_3, e) &\leftarrow \mathcal{G}(1^\lambda); \\ h_{123} &\leftarrow_R G_N; \\ D &:= ((N, G_N, G_T, e); g_1, g_3, h_{123}); \\ T_0 &\leftarrow_R G_{p_1}, T_1 \leftarrow_R G_{p_1 p_2}. \end{aligned}$$

We assume that for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{SD1}}(\lambda) := |\Pr[\mathcal{A}(D, T_0) = 1] - \Pr[\mathcal{A}(D, T_1) = 1]|$$

is negligible in λ .

Definition 7 (Assumption 2). *Consider the following distribution:*

$$\begin{aligned} (N, G_N, G_T, g_1, g_2, g_3, e) &\leftarrow \mathcal{G}(1^\lambda); \\ h_{123} &\leftarrow_R G_N, h_{23} \leftarrow_R G_{p_2 p_3}, g_{12} \leftarrow_R G_{p_1 p_2}; \\ D &:= ((N, G_N, G_T, e); g_1, g_3, h_{123}, h_{23}, g_{12}); \\ T_0 &\leftarrow_R G_{p_1 p_3}, T_1 \leftarrow_R G_N. \end{aligned}$$

We assume that for any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{SD2}}(\lambda) := |\Pr[\mathcal{A}(D, T_0) = 1] - \Pr[\mathcal{A}(D, T_1) = 1]|$$

is negligible in λ .

We show that the construction given in Section 5.2 of [CW14a] satisfies non-degeneracy, LSI, RSI and parameter-hiding properties.

Non-degeneracy: The statistical distance between \tilde{h}^δ and $(g^2 \cdot g^3)^{\hat{r}}$, where $\delta \leftarrow_R \mathbb{Z}_N$ and $\hat{r} \leftarrow_R \mathbb{Z}_N^*$, is at most $1/p_2 + 1/p_3$, which is negligible in λ . Similarly, if we set $g := g_2$, then the statistical distance between g^μ and $g_2^{\hat{s}}$ for $\mu \in \mathbb{Z}_N$ and $\hat{s} \in \mathbb{Z}_N^*$, is at most $1/p_2$.

Lemma 5 (SD1 to LSI). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{LSI}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{SD1}}(\lambda) + 2/p_1 + 2/p_2 + 1/p_3.$$

Proof. The adversary \mathcal{B} gets as input $((N, G_N, G_T, e); g_1, g_3, h_{123}, T)$, where T is chosen uniformly at random from G_{p_1} or $G_{p_1 p_2}$. Using this input, \mathcal{B} simulates the public parameters as follows. It picks $\mathbf{w} \leftarrow_R \mathbb{Z}_N^n$ and gives

$$\text{PP} := ((N, G_N, G_T, e); g_1, g_1^{\mathbf{w}}, g_3, h_{123})$$

to \mathcal{A} . Note that PP is properly distributed if h_{123} is a generator of G_N , which happens with probability at least $1 - 1/p_1 - 1/p_2 - 1/p_3$.

Consider any (positive) polynomial $\text{poly}(x)$, and let $\ell := \text{poly}(\lambda)$. \mathcal{B} picks ℓ numbers $u_1, u_2, \dots, u_\ell \leftarrow_R \mathbb{Z}_N$ such that $u_i \bmod p_2 \neq 0$ for $i \in [1, \ell]$, and gives

$$\mathbf{G}' = ((T^{u_1}, T^{u_1 \mathbf{w}}), \dots, (T^{u_\ell}, T^{u_\ell \mathbf{w}}))$$

as the challenge to \mathcal{A} . If $T \leftarrow G_{p_1}$, then \mathbf{G}' is identically distributed to \mathbf{G} when T is a generator of G_{p_1} , which happens with probability $1 - 1/p_1$. On the other hand when $T \leftarrow G_{p_1 p_2}$, then \mathbf{G}' is identically distributed to $\mathbf{G} \cdot \hat{\mathbf{G}}$ when T is a generator of $G_{p_1 p_2}$, which happens with probability at least $1 - 1/p_1 - 1/p_2$. \square

Lemma 6 (SD2 to RS1). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{RS1}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{SD2}}(\lambda) + 3/p_1 + 4/p_2 + 3/p_3.$$

Proof. The adversary \mathcal{B} gets as input $((N, G_N, G_T, e); g_1, g_3, h_{123}, h_{23}, g_{12}, T)$, where T is chosen uniformly at random from $G_{p_1 p_3}$ or from the whole group. Let poly_1 and poly_2 be two (positive) polynomials. Let ℓ and m denote $\text{poly}_1(\lambda)$ and $\text{poly}_2(\lambda)$ respectively. \mathcal{B} picks $\mathbf{w} \leftarrow_R \mathbb{Z}_N^n$ and ℓ numbers $u_1, u_2, \dots, u_\ell \leftarrow_R \mathbb{Z}_N$ such that $u_i \bmod p_2 \neq 0$ for $i \in [1, \ell]$, and gives

$$\text{PP} := ((N, G_N, G_T, e); g_1, g_1^{\mathbf{w}}, g_3, h_{123})$$

$$\tilde{h} := h_{23} \quad \text{and} \quad \mathbf{G} \cdot \hat{\mathbf{G}} := ((g_{12}^{u_1}, g_{12}^{u_1 \mathbf{w}}), \dots, (g_{12}^{u_\ell}, g_{12}^{u_\ell \mathbf{w}}))$$

to \mathcal{A} . In order for $(\text{PP}, \tilde{h}, \mathbf{G} \cdot \hat{\mathbf{G}})$ to be properly distributed, we need h_{123}, h_{23} and g_{12} to be generators of $G_N, G_{p_2 p_3}$ and $G_{p_1 p_2}$ respectively, which happens with probability at least $1 - 2/p_1 - 3/p_2 - 2/p_3$.

Now to simulate the challenge, \mathcal{B} picks m numbers $v_1, v_2, \dots, v_m \leftarrow_R \mathbb{Z}_N$ such that $v_j \bmod p_2 \neq 0$ for $j \in [1, m]$ and vectors $\mathbf{X}'_1, \mathbf{X}'_2, \dots, \mathbf{X}'_m \leftarrow_R G_{p_3}^n$ (using g_3), and outputs

$$\mathbf{H}' = ((T^{v_1}, T^{v_1 \mathbf{w}} \cdot \mathbf{X}'_1), \dots, (T^{v_m}, T^{v_m \mathbf{w}} \cdot \mathbf{X}'_m)).$$

If $T \leftarrow G_{p_1 p_3}$, then \mathbf{H}' is identically distributed to \mathbf{H} , except when T is not a generator of $G_{p_1 p_3}$, which happens with probability at most $1/p_1 + 1/p_3$. On the other hand when $T \leftarrow G_N$, then \mathbf{H}' is identically distributed to $\mathbf{H} \cdot \hat{\mathbf{H}}$, except when T is not a generator G_N , which happens with probability at most $1/p_1 + 1/p_2 + 1/p_3$. \square

Lemma 7 (Parameter-hiding). *For any polynomials $\text{poly}_1(x)$ and $\text{poly}_2(x)$, the following distributions are identical:*

$$\{\text{PP}, \tilde{h}, ((g_2^{\hat{s}_1}, g_2^{\hat{s}_1 \mathbf{w}}), \dots, (g_2^{\hat{s}_\ell}, g_2^{\hat{s}_\ell \mathbf{w}})), ((g_2^{\hat{r}_1} \cdot g_3^{\hat{r}_1}, g_2^{\hat{r}_1 \mathbf{w}} \cdot \mathbf{X}_1), \dots, (g_2^{\hat{r}_m} \cdot g_3^{\hat{r}_m}, g_2^{\hat{r}_m \mathbf{w}} \cdot \mathbf{X}_m))\}$$

and

$$\{\text{PP}, \tilde{h}, ((g_2^{\hat{s}_1}, g_2^{\hat{s}_1(\mathbf{w}+\mathbf{w}')}), \dots, (g_2^{\hat{s}_\ell}, g_2^{\hat{s}_\ell(\mathbf{w}+\mathbf{w}')}), ((g_2^{\hat{r}_1} \cdot g_3^{\hat{r}_1}, g_2^{\hat{r}_1(\mathbf{w}+\mathbf{w}')}) \cdot \mathbf{X}_1), \dots, (g_2^{\hat{r}_m} \cdot g_3^{\hat{r}_m}, g_2^{\hat{r}_m(\mathbf{w}+\mathbf{w}')}) \cdot \mathbf{X}_m))\},$$

where

$$\begin{aligned}
\ell &:= \text{poly}_1(\lambda) \quad \text{and} \quad m := \text{poly}_2(\lambda); \\
(\text{PP}, \text{SP}) &\leftarrow \text{SampP}(1^\lambda, 1^n); \\
\mathbf{w}, \mathbf{w}' &\leftarrow_R \mathbb{Z}_N^n; \\
\hat{s}_1, \dots, \hat{s}_\ell &\leftarrow_R \mathbb{Z}_N^*; \\
\hat{r}_1, \dots, \hat{r}_m &\leftarrow_R \mathbb{Z}_N^*; \\
\mathbf{X}_1, \dots, \mathbf{X}_m &\leftarrow_R G_{p_3}^m.
\end{aligned}$$

Proof. Note that \mathbf{w} appears in the public parameters PP in the form $g_1^{\mathbf{w}}$. Hence, $\mathbf{w} \bmod p_2$ is a uniformly random number in \mathbb{Z}_{p_2} given PP (by Chinese remainder theorem), and the lemma follows. \square

B.2 Prime-order construction

A prime-order bilinear group generator \mathcal{G} takes the security parameter λ as input and outputs $(p, G_1, G_2, G_T, g_1, g_2, e)$. G_1, G_2 and G_T are three multiplicative groups of order p , where p is a prime of $\Theta(\lambda)$ bits. e is an efficiently computable non-degenerate bilinear map which maps an element of G_1 and an element of G_2 to an element of G_T . g_1 and g_2 are generators of G_1 and G_2 respectively.

We first define the following generalization of d -LIN assumption with respect to \mathcal{G} , whose security follows tightly from d -LIN itself.

Definition 8 (*gen- d -LIN Assumption*). *Let $\text{poly}(x)$ be a (positive) polynomial in x . Given a group generator \mathcal{G} , we define the following distribution:*

$$\begin{aligned}
(p, G_1, G_2, G_T, g_1, g_2, e) &\leftarrow \mathcal{G}(1^\lambda); \\
m &:= \text{poly}(\lambda); \\
s_{1,1}, \dots, s_{1,d}, \dots, s_{m,1}, \dots, s_{m,d} &\leftarrow_R \mathbb{Z}_p; \\
a_1, a_2, \dots, a_{d+1}, s_{1,d+1}, \dots, s_{m,d+1} &\leftarrow_R \mathbb{Z}_p^*; \\
D &:= ((p, G_1, G_2, G_T, e); g_1, g_2, g_1^{a_1}, \dots, g_1^{a_d}, g_1^{a_{d+1}}, g_1^{a_1 s_{1,1}}, \dots, g_1^{a_d s_{1,d}}, \dots, g_1^{a_1 s_{m,1}}, \dots, g_1^{a_d s_{m,d}}); \\
T_0 &:= (g_1^{a_{d+1}(s_{1,1} + \dots + s_{1,d})}, \dots, g_1^{a_{d+1}(s_{m,1} + \dots + s_{m,d})}); \\
T_1 &:= (g_1^{a_{d+1}(s_{1,1} + \dots + s_{1,d}) + s_{1,d+1}}, \dots, g_1^{a_{d+1}(s_{m,1} + \dots + s_{m,d}) + s_{m,d+1}}).
\end{aligned}$$

We assume that for any polynomial p and any PPT algorithm \mathcal{A} ,

$$\text{Adv}_{\mathcal{A}}^{\text{gen-}d\text{-LIN}}(\lambda) := |\Pr[\mathcal{A}(D, T_0) = 1] - \Pr[\mathcal{A}(D, T_1) = 1]|$$

is negligible in λ .

The above assumption is defined with respect to the first group G_1 output by \mathcal{G} . We also assume that this assumption holds with respect to the second group G_2 . We now show how the d -LIN assumption (which is a special case of the above with $m = 1$) can be reduced to *gen- d -LIN*.

Lemma 8 (*d -LIN to gen- d -LIN*). *If the d -LIN assumption holds for a group generator \mathcal{G} , then the gen- d -LIN assumption stated in Definition 8 also holds in \mathcal{G} .*

Proof. Consider any (positive) polynomial $\text{poly}(x)$, and let $m := \text{poly}(\lambda)$. Let \mathcal{A} be a PPT algorithm that gets a non-negligible advantage in the $\text{gen-}d\text{-LIN}$ security game w.r.t. to the polynomial p . We construct a PPT algorithm \mathcal{B} which uses \mathcal{A} to break the $d\text{-LIN}$ assumption as follows. \mathcal{B} obtains

$$((p, G_1, G_2, G_T, e); g_1, g_2, g_1^{a_1}, \dots, g_1^{a_d}, g_1^{a_{d+1}}, g_1^{a_1 r_1}, \dots, g_1^{a_d r_d}, T := g_1^{a_{d+1}(r_1 + \dots + r_d) + r_{d+1}})$$

as input, where r_{d+1} is either 0 or uniformly chosen from \mathbb{Z}_p^* . It picks

$$s_{1,1}, \dots, s_{1,d}, \dots, s_{m,1}, \dots, s_{m,d} \leftarrow_R \mathbb{Z}_p,$$

$$s_{1,d+1}, \dots, s_{m,d+1} \leftarrow_R \mathbb{Z}_p^*,$$

and computes

$$[g_1^{a_1 r_1} \cdot (g_1^{a_1})^{s_{i,1}}]^{s_{i,d+1}}, \dots, [g_1^{a_d r_d} \cdot (g_1^{a_d})^{s_{i,d}}]^{s_{i,d+1}},$$

$$\begin{aligned} T'_i &:= [T \cdot (g_1^{a_{d+1}})^{s_{i,1}} \cdot (g_1^{a_{d+1}})^{s_{i,2}} \dots (g_1^{a_{d+1}})^{s_{i,d}}]^{s_{i,d+1}} \\ &= g_1^{a_{d+1}\{s_{i,d+1}(r_1 + s_{i,1}) + \dots + s_{i,d+1}(r_d + s_{i,d})\} + r_{d+1}s_{i,d+1}}, \end{aligned}$$

for every $i \in [1, m]$. \mathcal{B} then gives $((p, G_1, G_2, G_T, e); g_1, g_2, g_1^{a_1}, \dots, g_1^{a_d}, g_1^{a_{d+1}})$ along with the group elements computed above as challenge to \mathcal{A} . It is easy to see that the challenge has the right distribution. \square

We are now ready to show how the prime order construction satisfies the security properties we desire.

Non-degeneracy: We know that \tilde{h} is defined to be $g_2^{f^*}$ and $\overline{\text{SampH}}_0$ is distributed as $g_2^{\hat{r}f^*}$ for $\hat{r} \leftarrow_R \mathbb{Z}_p^*$. Hence, the statistical distance between \tilde{h}^δ , when $\delta \leftarrow_R \mathbb{Z}_p$, and $\overline{\text{SampH}}_0$ can be at most $1/p$. Similarly, if we set $g := g_1^f$, we can show that the second condition is satisfied.

Lemma 9 ($\text{gen-}d\text{-LIN}$ to LSI). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{LSI}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{gen-}d\text{-LIN}}(\lambda).$$

Proof. Consider any (positive) polynomial $\text{poly}(x)$, and let $m := \text{poly}(\lambda)$. We first write $(\text{PP}, \mathbf{g}_1, \dots, \mathbf{g}_m, \mathbf{g}_1 \cdot \hat{\mathbf{g}}_1, \dots, \mathbf{g}_m \cdot \hat{\mathbf{g}}_m)$ in terms of the prime order construction.

$$\text{PP} := ((p, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, e); g_1^{\rho_L(\mathbf{B})}, g_1^{\rho_L(\mathbf{BA}_1)}, \dots, g_1^{\rho_L(\mathbf{BA}_n)}, g_2^{\rho_L(\mathbf{B}^* \mathbf{R})}, g_1^{\rho_L(\mathbf{B}^* \mathbf{A}_1^T \mathbf{R})}, \dots, g_1^{\rho_L(\mathbf{B}^* \mathbf{A}_n^T \mathbf{R})}),$$

$$\forall i \in [1, m] \quad \mathbf{g}_i := \left(g_1^{\mathbf{B} \begin{pmatrix} \mathbf{s}_i \\ 0 \end{pmatrix}}, g_1^{\mathbf{BA}_1 \begin{pmatrix} \mathbf{s}_i \\ 0 \end{pmatrix}}, \dots, g_1^{\mathbf{BA}_n \begin{pmatrix} \mathbf{s}_i \\ 0 \end{pmatrix}} \right)$$

$$\forall i \in [1, m] \quad \mathbf{g}_i \cdot \hat{\mathbf{g}}_i := \left(g_1^{\mathbf{B} \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}}, g_1^{\mathbf{BA}_1 \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}}, \dots, g_1^{\mathbf{BA}_n \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}} \right)$$

In the above, $\mathbf{s}_1, \dots, \mathbf{s}_m \leftarrow_R \mathbb{Z}_p^d$ and $\hat{s}_1, \dots, \hat{s}_m \leftarrow_R \mathbb{Z}_p^*$.

\mathcal{B} obtains as input

$$((p, G_1, G_2, G_T, e); g_1, g_2, g_1^{a_1}, \dots, g_1^{a_d}, g_1^{a_{d+1}}, g_1^{a_1 s_{1,1}}, \dots, g_1^{a_d s_{1,d}}, \dots, g_1^{a_1 s_{m,1}}, \dots, g_1^{a_d s_{m,d}}, g_1^{a_{d+1}(s_{1,1} + \dots + s_{1,d}) + s_{1,d+1}}, \dots, g_1^{a_{d+1}(s_{m,1} + \dots + s_{m,d}) + s_{m,d+1}}),$$

where $s_{1,d+1}, \dots, s_{m,d+1}$ are all 0 or uniformly chosen from \mathbb{Z}_p^* .

To begin with, \mathcal{B} implicitly sets

$$\mathbf{s}_i := (s_{i,1}, s_{i,2}, \dots, s_{i,d}),$$

$$\hat{s}_i := s_{i,d+1},$$

for all $i \in [1, m]$. It programs $\mathbf{B}, \mathbf{B}^*, \mathbf{A}_1, \dots, \mathbf{A}_n, \mathbf{R}$ and simulates the public parameters PP in the same manner as the proof of Lemma 10 in [CW14a]. This involves defining

$$\mathbf{W} := \begin{pmatrix} a_1 & & & & & \\ & a_2 & & & & \\ & & \ddots & & & \\ & & & a_d & & \\ a_{d+1} & a_{d+1} & \dots & a_{d+1} & 1 & \end{pmatrix},$$

and sampling $\tilde{\mathbf{B}} \leftarrow_R \text{GL}_{d+1}(\mathbb{Z}_p)$ along with $\tilde{\mathbf{A}}_1, \dots, \tilde{\mathbf{A}}_n \leftarrow_R \mathbb{Z}_p^{(d+1) \times (d+1)}$.

Now, observe that

$$\mathbf{W} \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix} = \begin{pmatrix} a_1 s_{i,1} \\ \vdots \\ a_d s_{i,d} \\ a_{d+1}(s_{i,1} + \dots + s_{i,d}) + s_{i,d+1} \end{pmatrix},$$

and hence \mathcal{B} can compute

$$g_1^{\mathbf{W} \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}}$$

for all $i \in [1, m]$ using its input. Lastly, it outputs the challenge as

$$g_1^{\mathbf{B} \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}} = g_1^{\tilde{\mathbf{B}} \mathbf{W} \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}} \quad \text{and} \quad g_1^{\mathbf{B} \mathbf{A}_j \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}} = g_1^{\tilde{\mathbf{B}} \tilde{\mathbf{A}}_j \mathbf{W} \begin{pmatrix} \mathbf{s}_i \\ \hat{s}_i \end{pmatrix}}$$

for all $j \in [1, n]$ and $i \in [1, m]$. If $s_{1,d+1}, \dots, s_{m,d+1}$ are all 0, implying that $\hat{s}_1, \dots, \hat{s}_m$ are 0 as well, then the view of \mathcal{A} is identically distributed to $(\text{PP}, \mathbf{g}_1, \dots, \mathbf{g}_m)$, otherwise the view is distributed according to $(\text{PP}, \mathbf{g}_1 \cdot \hat{\mathbf{g}}_1, \dots, \mathbf{g}_m \cdot \hat{\mathbf{g}}_m)$. \square

Lemma 10 (*gen-d-LIN to RSI*). *For any PPT adversary \mathcal{A} , there exist a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{\text{RSI}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{gen-d-LIN}}(\lambda) + \text{poly}(\lambda)/p,$$

where $\text{poly}(\lambda)$ is independent of $\text{Adv}_{\mathcal{B}}^{\text{gen-d-LIN}}(\lambda)$.

Proof. Consider two (positive) polynomials $\text{poly}_1(x)$ and $\text{poly}_2(x)$. Define $\ell := \text{poly}_1(\lambda)$ and $m := \text{poly}_2(\lambda)$. We first write $(\text{PP}, \tilde{h}, \mathbf{g}_1 \cdot \hat{\mathbf{g}}_1, \dots, \mathbf{g}_\ell \cdot \hat{\mathbf{g}}_\ell, \mathbf{h}_1, \dots, \mathbf{h}_m, \mathbf{h}_1 \cdot \hat{\mathbf{h}}_1, \dots, \mathbf{h}_m \cdot \hat{\mathbf{h}}_m)$ in terms of the prime order construction.

$$\text{PP} := ((p, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, e); g_1^{\rho_L(\mathbf{B})}, g_1^{\rho_L(\mathbf{BA}_1)}, \dots, g_1^{\rho_L(\mathbf{BA}_n)}, g_2^{\rho_L(\mathbf{B}^*\mathbf{R})}, g_1^{\rho_L(\mathbf{B}^*\mathbf{A}_1^T\mathbf{R})}, \dots, g_1^{\rho_L(\mathbf{B}^*\mathbf{A}_n^T\mathbf{R})}),$$

$$\tilde{h} := g_2^{\rho_R(\mathbf{B}^*\mathbf{R})},$$

$$\forall j \in [1, \ell] \quad \mathbf{g}_j \cdot \hat{\mathbf{g}}_j := \left(g_1 \begin{pmatrix} \mathbf{B} \begin{pmatrix} \mathbf{s}_j \\ \hat{\mathbf{s}}_j \end{pmatrix} \\ \hat{\mathbf{s}}_j \end{pmatrix}, g_1 \begin{pmatrix} \mathbf{BA}_1 \begin{pmatrix} \mathbf{s}_j \\ \hat{\mathbf{s}}_j \end{pmatrix} \\ \hat{\mathbf{s}}_j \end{pmatrix}, \dots, g_1 \begin{pmatrix} \mathbf{BA}_n \begin{pmatrix} \mathbf{s}_j \\ \hat{\mathbf{s}}_j \end{pmatrix} \\ \hat{\mathbf{s}}_j \end{pmatrix} \right),$$

$$\forall i \in [1, m] \quad \mathbf{h}_i := \left(g_2 \begin{pmatrix} \mathbf{B}^*\mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ 0 \end{pmatrix} \\ 0 \end{pmatrix}, g_2 \begin{pmatrix} \mathbf{B}^*\mathbf{A}_1^T\mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ 0 \end{pmatrix} \\ 0 \end{pmatrix}, \dots, g_2 \begin{pmatrix} \mathbf{B}^*\mathbf{A}_n^T\mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ 0 \end{pmatrix} \\ 0 \end{pmatrix} \right),$$

$$\forall i \in [1, m] \quad \mathbf{h}_i \cdot \hat{\mathbf{h}}_i := \left(g_2 \begin{pmatrix} \mathbf{B}^*\mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} \\ \hat{r}_i \end{pmatrix}, g_2 \begin{pmatrix} \mathbf{B}^*\mathbf{A}_1^T\mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} \\ \hat{r}_i \end{pmatrix}, \dots, g_2 \begin{pmatrix} \mathbf{B}^*\mathbf{A}_n^T\mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} \\ \hat{r}_i \end{pmatrix} \right).$$

In the above, $\mathbf{s}_1, \dots, \mathbf{s}_\ell, \mathbf{r}_1, \dots, \mathbf{r}_m \leftarrow_R \mathbb{Z}_p^d$ and $\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_\ell, \hat{r}_1, \dots, \hat{r}_m \leftarrow_R \mathbb{Z}_p^*$.
 \mathcal{B} obtains as input

$$((p, G_1, G_2, G_T, e); g_1, g_2, g_1^{a_1}, \dots, g_1^{a_d}, g_1^{a_{d+1}}, g_1^{a_1 r_{1,1}}, \dots, g_1^{a_d r_{1,d}}, \dots, g_1^{a_1 r_{m,1}}, \dots, g_1^{a_d r_{m,d}}, g_1^{a_{d+1}(r_{1,1} + \dots + r_{1,d}) + r_{1,d+1}}, \dots, g_1^{a_{d+1}(r_{m,1} + \dots + r_{m,d+1}) + r_{m,d+1}}),$$

where $r_{1,d+1}, \dots, r_{m,d+1}$ are all 0 or uniformly chosen from \mathbb{Z}_p^* .

To begin with, \mathcal{B} picks $\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_d \leftarrow_R \mathbb{Z}_p^*$ and implicitly sets

$$\mathbf{R} := \begin{pmatrix} a_1 \tilde{r}_1 & & & & \\ & a_2 \tilde{r}_2 & & & \\ & & \ddots & & \\ & & & a_d \tilde{r}_d & \\ & & & & 1 \end{pmatrix},$$

$$\mathbf{r}_i := (\tilde{r}_1^{-1} r_{i,1}, \tilde{r}_2^{-1} r_{i,2}, \dots, \tilde{r}_d^{-1} r_{i,d}),$$

$$\hat{r}_i := r_{i,d+1},$$

for all $i \in [1, m]$. It programs $\mathbf{B}, \mathbf{B}^*, \mathbf{A}_1, \dots, \mathbf{A}_n$ and simulates PP, \tilde{h} along with $\mathbf{g}_1 \cdot \hat{\mathbf{g}}_1, \dots, \mathbf{g}_\ell \cdot \hat{\mathbf{g}}_\ell$ (there is an error of ℓ/p in simulating this) in the same manner as the proof of Lemma 11 in [CW14a]. This involves defining

$$\mathbf{W}^* := \begin{pmatrix} 1 & & & & \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ a_1^{-1}a_{d+1} & a_2^{-1}a_{d+1} & \dots & a_d^{-1}a_{d+1} & 1 \end{pmatrix},$$

sampling $\tilde{\mathbf{B}} \leftarrow_R \text{GL}_{d+1}(\mathbb{Z}_p)$ along with $\tilde{\mathbf{A}}_1, \dots, \tilde{\mathbf{A}}_n \leftarrow_R \mathbb{Z}_p^{(d+1) \times (d+1)}$, and setting $\tilde{\mathbf{B}}^* := (\tilde{\mathbf{B}}^{-1})^T$. Now, observe that

$$\mathbf{W}^* \mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} = \begin{pmatrix} a_1 r_{i,1} \\ \vdots \\ a_d r_{i,d} \\ a_{d+1}(r_{i,1} + \dots + r_{i,d}) + r_{i,d+1} \end{pmatrix},$$

and hence \mathcal{B} can compute

$$g_2 \begin{pmatrix} \mathbf{W}^* \mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} \end{pmatrix}$$

for all $i \in [1, m]$ using its input. Lastly, it outputs the challenge as

$$g_2 \mathbf{B}^* \mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} = g_2 \tilde{\mathbf{B}}^* \mathbf{W}^* \mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} \quad \text{and} \quad g_2 \mathbf{B}^* \mathbf{A}_j^T \mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix} = g_2 \tilde{\mathbf{B}}^* \tilde{\mathbf{A}}_j^T \mathbf{W}^* \mathbf{R} \begin{pmatrix} \mathbf{r}_i \\ \hat{r}_i \end{pmatrix}$$

for all $j \in [1, n]$ and $i \in [1, m]$. If $r_{1,d+1}, \dots, r_{m,d+1}$ are all 0, implying that $\hat{r}_1, \dots, \hat{r}_m$ are 0 as well, then the view of \mathcal{A} is identically distributed to $(\text{PP}, \tilde{h}, \mathbf{g}_1 \cdot \hat{\mathbf{g}}_1, \dots, \mathbf{g}_\ell \cdot \hat{\mathbf{g}}_\ell, \mathbf{h}_1, \dots, \mathbf{h}_m)$, otherwise the view is distributed according to $(\text{PP}, \tilde{h}, \mathbf{g}_1 \cdot \hat{\mathbf{g}}_1, \dots, \mathbf{g}_\ell \cdot \hat{\mathbf{g}}_\ell, \mathbf{h}_1 \cdot \hat{\mathbf{h}}_1, \dots, \mathbf{h}_m \cdot \hat{\mathbf{h}}_m)$. \square

Lemma 11 (Parameter-hiding). *For any (positive) polynomials $\text{poly}_1(x)$ and $\text{poly}_2(x)$, the following distributions are identical:*

$$\left\{ \text{PP}, g_2^{\mathbf{f}^*}, (g_1^{\hat{s}_1 \mathbf{f}}, g_1^{\hat{s}_1 \mathbf{f}_1}, \dots, g_1^{\hat{s}_1 \mathbf{f}_n}), \dots, (g_1^{\hat{s}_\ell \mathbf{f}}, g_1^{\hat{s}_\ell \mathbf{f}_1}, \dots, g_1^{\hat{s}_\ell \mathbf{f}_n}), \right. \\ \left. (g_2^{\hat{r}_1 \mathbf{f}^*}, g_2^{\hat{r}_1 \mathbf{f}_1^*}, \dots, g_2^{\hat{r}_1 \mathbf{f}_n^*}), \dots, (g_2^{\hat{r}_m \mathbf{f}^*}, g_2^{\hat{r}_m \mathbf{f}_1^*}, \dots, g_2^{\hat{r}_m \mathbf{f}_n^*}) \right\}$$

and

$$\left\{ \text{PP}, g_2^{\mathbf{f}^*}, (g_1^{\hat{s}_1 \mathbf{f}}, g_1^{\hat{s}_1(\mathbf{f}_1 + \gamma_1 \mathbf{f})}, \dots, g_1^{\hat{s}_1(\mathbf{f}_n + \gamma_n \mathbf{f})}), \dots, (g_1^{\hat{s}_\ell \mathbf{f}}, g_1^{\hat{s}_\ell(\mathbf{f}_1 + \gamma_1 \mathbf{f})}, \dots, g_1^{\hat{s}_\ell(\mathbf{f}_n + \gamma_n \mathbf{f})}), \right. \\ \left. (g_2^{\hat{r}_1 \mathbf{f}^*}, g_2^{\hat{r}_1(\mathbf{f}_1^* + \gamma_1 \mathbf{f}^*)}, \dots, g_2^{\hat{r}_1(\mathbf{f}_n^* + \gamma_n \mathbf{f}^*)}), \dots, (g_2^{\hat{r}_m \mathbf{f}^*}, g_2^{\hat{r}_m(\mathbf{f}_1^* + \gamma_1 \mathbf{f}^*)}, \dots, g_2^{\hat{r}_m(\mathbf{f}_n^* + \gamma_n \mathbf{f}^*)}) \right\}$$

where

$$\begin{aligned} \ell &:= \text{poly}_1(\lambda) \quad \text{and} \quad m := \text{poly}_2(\lambda); \\ (\text{PP}, \text{SP}) &\leftarrow \text{SampP}(1^\lambda, 1^n); \\ \hat{s}_1, \dots, \hat{s}_\ell, \hat{r}_1, \dots, \hat{r}_m &\leftarrow_R \mathbb{Z}_p^*; \\ \gamma_1, \dots, \gamma_n &\leftarrow_R \mathbb{Z}_p. \end{aligned}$$

Proof. Our proof closely follows the one given for Lemma 12 in [CW14a]. In a manner similar to their's, we could define

$$\mathbf{A}'_i := \mathbf{A}_i + \gamma_i \mathbf{V}$$

for $i \in [1, n]$, where \mathbf{V} is a matrix which is 0 everywhere except the bottom right entry which is 1. We run SampP with $(\mathbf{A}'_1, \dots, \mathbf{A}'_n)$ instead of $(\mathbf{A}_1, \dots, \mathbf{A}_n)$ to generate (PP, SP) and

$$\begin{aligned} \hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \dots, \hat{\mathbf{g}}_\ell &\leftarrow \overline{\text{SampG}}(\text{PP}, \text{SP}); \\ \hat{\mathbf{h}}_1, \hat{\mathbf{h}}_2, \dots, \hat{\mathbf{h}}_m &\leftarrow \overline{\text{SampH}}(\text{PP}, \text{SP}). \end{aligned}$$

Since all the samples above share the same $\gamma_1, \dots, \gamma_n$, one can easily verify that rest of the proof in [CW14a] goes through for the present generalized case as well. \square

C Proof of security

We use $\text{Adv}_{\mathcal{A}}^{p-q}(\lambda)$ to denote the advantage of an adversary \mathcal{A} in distinguishing Hyb_p from Hyb_q when the security parameter is λ .

Lemma 12. *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{0-1}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{LSI}}(\lambda).$$

Proof. \mathcal{B} gets as input (PP, \mathbf{G}') where \mathbf{G}' is either \mathbf{G} or $\mathbf{G} \cdot \hat{\mathbf{G}}$. While \mathbf{G} is an ordered set of $w_2 + 1$ samples from SampG, $\hat{\mathbf{G}}$ is an ordered set of the same size with samples from $\overline{\text{SampG}}$ (recall that LSI property holds for every polynomial, and in particular, for $w_2 + 1$). \mathcal{B} first picks $\text{MSK} \leftarrow_R \mathbb{H}$ and outputs $(\text{PP}, \mu(\text{MSK}))$ as the master public key. When \mathcal{A} sends a challenge x^* and two messages m_0, m_1 , \mathcal{B} responds with $\overline{\text{Encrypt}}(\text{PP}, x^*, m_b; \mathbf{G}', \text{MSK})$ as the ciphertext, where b is uniformly chosen bit. Further, when \mathcal{A} issues a key query y (either before or after the challenge), \mathcal{B} responds with $\overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y; (\mathbf{h}_1, \dots, \mathbf{h}_{m_2}))$ by sampling $\mathbf{h}_1, \dots, \mathbf{h}_{m_2}$ from SampH.

When $\mathbf{G}' = \mathbf{G}$, then the view of \mathcal{A} is identically distributed as in Hyb_0 . On the other hand, when $\mathbf{G}' = \mathbf{G} \cdot \hat{\mathbf{G}}$, it is easy to see that view of \mathcal{A} is identical to Hyb_1 . \square

Lemma 13. *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\mathcal{A}}^{(2,\varphi,3,\rho-1)-(2,\varphi,1,\rho)}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{RSI}}(\lambda),$$

for every $\varphi \in [1, \xi]$ and $\rho \in [1, m_{2,\varphi}]$.

Proof. \mathcal{B} gets as input $(\text{PP}, \tilde{h}, \mathbf{G} \cdot \hat{\mathbf{G}}, \mathbf{h}')$ where \mathbf{h}' is either \mathbf{h} or $\mathbf{h} \cdot \hat{\mathbf{h}}$ (special case of RSI with $\text{poly}_2(x) = 1$). \mathcal{B} first picks $\text{MSK} \leftarrow_R \mathbb{H}$ and outputs $(\text{PP}, \mu(\text{MSK}))$ as the master public key. When \mathcal{A} sends a challenge x^* and two messages m_0, m_1 , \mathcal{B} responds with $\overline{\text{Encrypt}}(\text{PP}, x^*, m_b; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK})$ as the ciphertext, where b is uniformly chosen bit.

\mathcal{B} picks a $\beta \leftarrow_R \mathbb{Z}_N$ and sets $\overline{\text{MSK}} := \text{MSK} \cdot (\tilde{h})^\beta$. When \mathcal{A} issues ς th key query y_ς , it responds with

$$\text{SK}_{y_\varsigma} := \begin{cases} \overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y_\varsigma; (\mathbf{h}_1^{(\varsigma)}, \dots, \mathbf{h}_{m_2, \varsigma}^{(\varsigma)})) & \text{if } \varsigma < \varphi \\ \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y_\varsigma; (\mathbf{h}_1^{(\varsigma)} \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}_{\rho-1}^{(\varsigma)} \cdot \tilde{h}^{\mathbf{z}^{\rho-1}}, \mathbf{h}', \mathbf{h}_{\rho+1}^{(\varsigma)}, \dots, \mathbf{h}_{m_2, \varsigma}^{(\varsigma)})) & \text{if } \varsigma = \varphi \\ \overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y_\varsigma; (\mathbf{h}_1^{(\varsigma)}, \dots, \mathbf{h}_{m_2, \varsigma}^{(\varsigma)})) & \text{otherwise,} \end{cases}$$

where for every $\varsigma \in [1, \xi]$ and $i \in [1, m_{2, \varsigma}]$ (except when $\varsigma = \varphi$ and $i = \rho$), $\mathbf{h}_i^{(\varsigma)} \leftarrow \text{SampH}(\text{PP})$, and for every $j \in [1, \rho - 1]$, $\mathbf{z}_j \leftarrow \overline{\text{Samp}}(j, x, y, N)$. It is easy to see that when $\mathbf{h}' = \mathbf{h}$, then the view of \mathcal{A} is identically distributed to $\text{Hyb}_{2, \varphi, 3, \rho-1}$, and when $\mathbf{h}' = \mathbf{h} \cdot \hat{\mathbf{h}}$, then it is identically distributed to $\text{Hyb}_{2, \varphi, 1, \rho}$. \square

We now see how the above proof can be adapted to show indistinguishability between other pairs of hybrids. Below, we only describe the changes that need to be made; other details can be easily worked out.

- $\text{Hyb}_{2, \varphi, 2, \rho} \approx \text{Hyb}_{2, \varphi, 3, \rho}$: In order to generate the φ th key, use

$$\overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y_\varphi; (\mathbf{h}_1^{(\varphi)} \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}_{\rho-1}^{(\varphi)} \cdot \tilde{h}^{\mathbf{z}^{\rho-1}}, \mathbf{h}' \cdot \tilde{h}^{\mathbf{z}^\rho}, \mathbf{h}_{\rho+1}^{(\varphi)}, \dots, \mathbf{h}_{m_2, \varphi}^{(\varphi)})),$$

where $\mathbf{z}_\rho \leftarrow \overline{\text{Samp}}(\rho, x, y, N)$.

- $\text{Hyb}_{2, \varphi, 3, m_2, \varphi} \approx \text{Hyb}_{2, \varphi, 4}$: Assume that \mathcal{B} gets $(\text{PP}, \tilde{h}, \mathbf{G} \cdot \hat{\mathbf{G}}, \mathbf{H}')$ as input where $\mathbf{H}' := (\mathbf{h}'_1, \dots, \mathbf{h}'_{m_2, \varphi})$ is a vector of $m_2 + 1$ samples (instead of just 1). In order to generate the φ th key, it uses

$$\overline{\text{KeyGen}}(\text{PP}, \text{MSK}, y_\varphi; (\mathbf{h}'_1 \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}'_{m_2, \varphi} \cdot \tilde{h}^{\mathbf{z}^{m_2, \varphi}})),$$

where $\mathbf{z}_j \leftarrow \overline{\text{Samp}}(j, x, y, N)$ for all $j \in [1, m_{2, \varphi}]$.

- $\text{Hyb}_{2, \varphi, 5} \approx \text{Hyb}_{2, \varphi, 6}$: Once again assume that \mathcal{B} gets $(\text{PP}, \tilde{h}, \mathbf{G} \cdot \hat{\mathbf{G}}, \mathbf{H}')$ as input where $\mathbf{H}' := (\mathbf{h}'_1, \dots, \mathbf{h}'_{m_2, \varphi})$ is a vector of $m_2 + 1$ samples. In order to generate the φ th key, it uses

$$\overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y_\varphi; (\mathbf{h}'_1 \cdot \tilde{h}^{\mathbf{z}^1}, \dots, \mathbf{h}'_{m_2, \varphi} \cdot \tilde{h}^{\mathbf{z}^{m_2, \varphi}})),$$

where $\mathbf{z}_j \leftarrow \overline{\text{Samp}}(j, x, y, N)$ for $j \in [1, m_{2, \varphi}]$.

- $\text{Hyb}_{2, \varphi, 9, \rho-1} \approx \text{Hyb}_{2, \varphi, 7, \rho}$: In order to generate the φ th key, use

$$\overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y_\varphi; (\mathbf{h}_1^{(\varphi)}, \dots, \mathbf{h}_{\rho-1}^{(\varphi)}, \mathbf{h}' \cdot \tilde{h}^{\mathbf{z}^\rho}, \mathbf{h}_{\rho+1}^{(\varphi)} \cdot \tilde{h}^{\mathbf{z}^{\rho+1}}, \dots, \mathbf{h}_{m_2, \varphi}^{(\varphi)} \cdot \tilde{h}^{\mathbf{z}^{m_2, \varphi}})),$$

where $\mathbf{z}_j \leftarrow \overline{\text{Samp}}(j, x, y, N)$ for $j \in [\rho, m_{2, \varphi}]$.

- $\text{Hyb}_{2, \varphi, 8, \rho} \approx \text{Hyb}_{2, \varphi, 9, \rho}$: In order to generate the φ th key, use

$$\overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y_\varphi; (\mathbf{h}_1^{(\varphi)}, \dots, \mathbf{h}_{\rho-1}^{(\varphi)}, \mathbf{h}', \mathbf{h}_{\rho+1}^{(\varphi)} \cdot \tilde{h}^{\mathbf{z}^{\rho+1}}, \dots, \mathbf{h}_{m_2, \varphi}^{(\varphi)} \cdot \tilde{h}^{\mathbf{z}^{m_2, \varphi}})),$$

where $\mathbf{z}_j \leftarrow \overline{\text{Samp}}(j, x, y, N)$ for $j \in [\rho + 1, m_{2, \varphi}]$.

Lemma 14. $\text{Hyb}_{2, \xi, 9, m_2, \xi} \cong \text{Hyb}_3$.

Proof. The only difference between the hybrids $\text{Hyb}_{2,\xi,9,m_2,\xi}$ and Hyb_3 is that in the former case the message m_b is encrypted, while in the latter case a random message is encrypted; all the keys as well as the ciphertext in both the cases are in the semi-functional space. The following line of argument is very similar to the one in [CW14a] for the corresponding lemma.

We can assume that MSK and $\overline{\text{MSK}}$ are sampled as follows: first pick $\overline{\text{MSK}} \leftarrow_R \mathbb{H}$ and then set $\text{MSK} := \overline{\text{MSK}} \cdot \tilde{h}^\beta$, where $\beta \leftarrow_R \mathbb{Z}_N$. Observe that

$$\mu(\text{MSK}) = \mu(\overline{\text{MSK}} \cdot \tilde{h}^\beta) = \mu(\overline{\text{MSK}}) \cdot \mu(\tilde{h})^\beta = \mu(\overline{\text{MSK}}) \quad (21)$$

due to the linearity of μ and the orthogonality property ($\mu(\tilde{h}) = 1$). Further, for any public parameters PP and coin tosses σ ,

$$e(\text{SampG}_0(\text{PP}; \sigma), \tilde{h}) = \text{SampGT}(\mu(\tilde{h}); \sigma) = \text{SampGT}(1; \sigma) = e(\text{SampG}_0(\text{PP}; \sigma), 1) = 1 \quad (22)$$

due to the projective and orthogonality properties.

We now show that the view of any adversary \mathcal{A} in both the hybrids can be simulated given PP and $\mathbf{G} \cdot \hat{\mathbf{G}}$ only. First pick $\overline{\text{MSK}}$ and MSK as described above. Output $(\text{PP}, \mu(\overline{\text{MSK}}))$ as the master public key MPK ; using (21), this is identically distributed to $(\text{PP}, \mu(\text{MSK}))$. When \mathcal{A} issues a key query y , respond with $\overline{\text{KeyGen}}(\text{PP}, \overline{\text{MSK}}, y; (\mathbf{h}_1, \dots, \mathbf{h}_{m_2}))$, where $\mathbf{h}_1, \dots, \mathbf{h}_{m_2} \leftarrow \text{SampH}(\text{PP})$. When \mathcal{A} sends a pair of messages (m_0, m_1) and an x , where $m_0, m_1 \in \mathbb{G}_T$, output $\text{CT} := \overline{\text{Encrypt}}(\text{PP}, x, m_b; \mathbf{G} \cdot \hat{\mathbf{G}}, \text{MSK})$. It is clear that the view of \mathcal{A} in this experiment is identically distributed to its view in $\text{Hyb}_{2,\xi,9,m_2,\xi}$. In order to prove that this view is also identically distributed to the view in Hyb_3 , we only need to show that CT is the encryption of a random message.

We know that CT has $w_1 + 1$ components. The first w_1 components depend on PP , x and $\mathbf{G} \cdot \hat{\mathbf{G}}$, while the last one, $\text{CT}_{w_1+1} := m_b \cdot e(g_{0,0} \cdot \hat{g}_{0,0}, \text{MSK})$, depends on m_b and MSK (see Section 5). Now,

$$\begin{aligned} e(g_{0,0} \cdot \hat{g}_{0,0}, \text{MSK}) &= e(g_{0,0} \cdot \hat{g}_{0,0}, \overline{\text{MSK}} \cdot \tilde{h}^\beta) \\ &= e(g_{0,0} \cdot \hat{g}_{0,0}, \overline{\text{MSK}}) \cdot e(\hat{g}_{0,0}, \tilde{h}^\beta) \cdot e(g_{0,0}, \tilde{h}^\beta) \\ &= e(g_{0,0} \cdot \hat{g}_{0,0}, \overline{\text{MSK}}) \cdot e(\hat{g}_{0,0}, \tilde{h}^\beta) && \text{(due to (22))} \\ &= e(g_{0,0} \cdot \hat{g}_{0,0}, \overline{\text{MSK}}) \cdot e(\hat{g}_{0,0}, \tilde{h})^\beta. \end{aligned}$$

Observe that MPK , the keys and other parts of the ciphertext do not depend on β , which is chosen uniformly from \mathbb{Z}_N . Therefore, $e(\hat{g}_{0,0}, \tilde{h})^\beta$ is uniformly distributed over \mathbb{G}_T from the non-degeneracy property. This implies that CT is identically distributed to the encryption of a random message. \square

D Ciphertext-Policy ABE

Correctness: If (A, π) accepts S , then we know that there exists constants $\{\varepsilon_i\}_{i \in \Upsilon}$ such that $\sum_{i \in \Upsilon} \varepsilon_i \sum_{j \in [1, n_2]} a_{i,j} v_j = v_1 = \alpha$, where $\Upsilon = \{i \mid i \in [1, n_1], \pi(i) \in S\}$. Below we show how to combine the polynomials generated by EncC and EncK using $\{\varepsilon_i\}_{i \in \Upsilon}$ in order to recover αs (this

implicitly defines the output of $\text{Pair}((A, \pi), S, N)$:

$$\begin{aligned}
& c_2 \left(\sum_{i \in \Upsilon} \varepsilon_i k_{1,i} \right) - c_1 \left[\sum_{i \in \Upsilon} \varepsilon_i \left(\sum_{j \in [1, n_2]} a_{i,j} k_{2,i,j} + \sum_{\substack{\ell \in [1, n_1], \ell \neq i \\ j \in [1, n_2]}} a_{\ell,j} k_{3,i,\ell,j} \right. \right. \\
& \qquad \qquad \qquad \left. \left. + k_{4,i,\pi(i)} + \sum_{\substack{\ell \in [1, n_1], \ell \neq i \\ t \in [0, T]}} \pi(\ell)^t k_{5,i,\ell,t} \right) \right] \\
&= \sum_{i \in \Upsilon} \varepsilon_i \left[c_2 k_{1,i} - c_1 \left(\sum_{j \in [1, n_2]} a_{i,j} k_{2,i,j} + \sum_{\substack{\ell \in [1, n_1], \ell \neq i \\ j \in [1, n_2]}} a_{\ell,j} k_{3,i,\ell,j} \right. \right. \\
& \qquad \qquad \qquad \left. \left. + k_{4,i,\pi(i)} + \sum_{\substack{\ell \in [1, n_1], \ell \neq i \\ t \in [0, T]}} \pi(\ell)^t k_{5,i,\ell,t} \right) \right] \\
&= \sum_{i \in \Upsilon} \varepsilon_i \left(sr_i \sum_{\substack{i' \in [1, n_1] \\ j \in [1, n_2]}} a_{i,j} b_{i',j} + sr_i \sum_{\substack{i' \in [1, n_1] \\ t \in [0, T]}} \pi(i)^t b'_{i',t} \right. \\
& \qquad \qquad \qquad - sr_i \sum_{j \in [1, n_2]} a_{i,j} b_{i,j} + s \sum_{j \in [1, n_2]} a_{i,j} v_j - sr_i \sum_{\substack{\ell \in [1, n_1], \ell \neq i \\ j \in [1, n_2]}} a_{\ell,j} b_{\ell,j} \\
& \qquad \qquad \qquad \left. - sr_i \sum_{t \in [0, T]} \pi(i)^t b'_{i,t} - sr_i \sum_{\substack{\ell \in [1, n_1], \ell \neq i \\ t \in [0, T]}} \pi(\ell)^t b'_{\ell,t} \right) \\
&= s \sum_{i \in \Upsilon} \varepsilon_i \sum_{j \in [1, n_2]} a_{i,j} v_j = \alpha s.
\end{aligned}$$

E Key Policy ABE

The following pair encoding scheme $\Phi_{\text{kp-abe}} = (\text{Param}, \text{EncC}, \text{EncK}, \text{Pair})$ for the KP-ABE predicate family (which is the dual of the predicate family for CP-ABE described in the Section 7) is taken (verbatim) from Attrapadung [Att14]. While this pair encoding scheme is not secure under the stricter information-theoretic property of Attrapadung, we show that it satisfies our relaxed definition. Thus we get a semi-adaptively secure KP-ABE scheme with short ciphertexts under a standard assumption, while Attrapadung proved (full) security under q -type assumptions.

- Param(par) $\rightarrow T + 6$. Let $\mathbf{b} = (\mathbf{b}_0, \mathbf{b}_1, \dots, \mathbf{b}_{T+1}, \phi_1, \phi_2, \phi_3, \eta)$.
- EncC(S, N) $\rightarrow \mathbf{c}(\mathbf{s}, \mathbf{b}) := (c_1, c_2, c_3, c_4, c_5, c_6)$ where

$$c_1 = s, \quad c_2 = s\eta, \quad c_3 = s\phi_1 + w\phi_2,$$

$$c_4 = w, \quad c_5 = w\phi_3 + \tilde{s}(b_0 + b_1z_0 + \dots + b_{T+1}z_T), \quad c_6 = \tilde{s},$$

$\mathbf{s} = (s, w, \tilde{s})$, and z_i is the coefficient of x^i in $p(x) := \prod_{y \in S}(x - y)$.

- EncK($(\mathbf{A}, \pi), N$) $\rightarrow \mathbf{k}(\alpha, \mathbf{r}, \mathbf{b}) := (k_1, k_2, k_3, \{k_{4,i}, k_{5,i}, \mathbf{k}_{6,i}\}_{i \in [1, n_1]})$ where

$$k_1 = \alpha + r\phi_1 + u\eta, \quad k_2 = u, \quad k_3 = r,$$

$$k_{4,i} = \mathbf{A}_i \mathbf{v}^T + r_i\phi_3, \quad k_{5,i} = r_i,$$

$$\mathbf{k}_{6,i} = (r_i b_0, r_i(b_2 - b_1\pi(i)), \dots, r_i(b_{T+1} - b_1\pi(i)^T)),$$

$$\mathbf{v} = (v_1, \dots, v_{n_2}), v_1 = r\phi_2 \text{ and } \mathbf{r} = (r, u, r_1, r_2, \dots, r_{n_1}, v_2, \dots, v_{n_2}).$$

(See [Att14] for the Pair algorithm.)

Let an algorithm Samp be defined as follows. On input an $i \in [1, n_1]$, S , (A, π) , and N , Samp checks whether $\pi(i) \notin S$. If yes, it picks $\hat{\phi}_3 \leftarrow_R \mathbb{Z}_N$ and outputs $\hat{\mathbf{b}}_i := (0, \dots, 0, \hat{\phi}_3, 0)$; otherwise, a zero vector is output.

Lemma 15. $\Phi_{\text{kp-abe}}$ is relaxed perfect secure w.r.t. the algorithm Samp defined above.

Proof (sketch). It is instructive to see why this scheme is not perfectly secure. Consider (\mathbf{A}, π) and S such that the policy is not satisfied but there exists a j such that $\pi(j) \in S$. If we take a linear combination of the terms in $\mathbf{k}_{6,j}$ with $(1, z_1, \dots, z_T)$, we get $r_j(b_0 + (b_2 - b_1\pi(j))z_1 + \dots + (b_{T+1} - b_1\pi(j)^T)z_T)$, which can be simplified to $r_j(b_0 + b_1z_0 + \dots + b_{T+1}z_T)$ since $\pi(j)$ is a root of $p(x)$. From this and $k_{5,j}$, we know $b_0 + b_1z_0 + \dots + b_{T+1}z_T$. Further, using c_4 and c_6 , we can recover ϕ_3 from c_5 . Now, using $k_{4,i}$ and $k_{5,i}$, we get $\mathbf{A}_i \mathbf{v}^T$ for every $i \in [1, n_1]$, revealing all the shares of the secret $v_1 = r\phi_2$. One can go further and show that this exposes α ; hence perfect security is violated.

So, the problem is that even if S has only one *good* attribute (i.e., a j s.t. $\pi(j) \in S$), *all* the shares of the secret are revealed. If we can somehow hide the shares corresponding to the rows for which S does not have a valid attribute, it will be enough. The relaxed property enables us to do this by focusing on just one r_i , or one row of the access matrix \mathbf{A} , at a time—rest of the elements in \mathbf{r} are set to zero. If $\pi(i) \notin S$, then linearly combining $\mathbf{k}_{6,i}$ with $(1, z_1, \dots, z_T)$ gives an additional term $r_i b_1 \pi(i)$, so the above line of argument breaks down; in particular, one cannot recover ϕ_3 that way anymore. Therefore, we can add a random element to ϕ_3 without changing the distribution. We can, similarly, add *fresh randomness* to ϕ_3 for every i such that $\pi(i) \notin S$. Once noise has been added this way, and we look at the distribution of the whole key, shares for which S doesn't have an attribute will remain hidden. Hence, it wouldn't be possible to learn α .

Converting the above argument into a formal proof of relaxed perfect security using Samp is not hard. We leave it to the reader. \square