Tools for Simulating Features of Composite Order Bilinear Groups in the Prime Order Setting

Allison Lewko*

The University of Texas at Austin
alewko@cs.utexas.edu

Abstract. In this paper, we explore a general methodology for converting composite order pairing-based cryptosystems into the prime order setting. We employ the dual pairing vector space approach initiated by Okamoto and Takashima and formulate versatile tools in this framework that can be used to translate composite order schemes for which the prior techniques of Freeman were insufficient. Our techniques are typically applicable for composite order schemes relying on the canceling property and proven secure from variants of the subgroup decision assumption, and will result in prime order schemes that are proven secure from the decisional linear assumption. As an instructive example, we obtain a translation of the Lewko-Waters composite order IBE scheme. This provides a close analog of the Boneh-Boyen IBE scheme that is proven fully secure from the decisional linear assumption. In the full version of this paper, we also provide a translation of the Lewko-Waters unbounded HIBE scheme.

1 Introduction

Recently, several cryptosystems have been constructed in composite order bilinear groups and proven secure from instances (and close variants) of the general subgroup decision assumption defined in [3]. For example, the systems presented in [27,25,29,28,26] provide diverse and advanced functionalities like identity-based encryption (IBE), hierarchical identity-based encryption (HIBE), and attribute-based encryption with strong security guarantees (e.g. full security, leakage-resilience) proven from static assumptions. These works leverage convenient features of composite order bilinear groups that are not shared by prime order bilinear groups, most notably the presence of orthogonal subgroups of coprime orders. Up to isomorphism, a composite order bilinear group has the structure of a direct product of prime order subgroups, so every group element can be decomposed as the product of components in the separate subgroups. However, when the group order is hard to factor, such a decomposition is hard to compute. The orthogonality of these subgroups means that they can function as independent spaces, allowing a system designer to use them in different ways

* Supported by a Microsoft Research Ph.D. Fellowship.

D. Pointcheval and T. Johansson (Eds.): EUROCRYPT 2012, LNCS 7237, pp. 318–335, 2012.

without any cross interactions between them destroying correctness. Security relies on the assumption that these subgroups are essentially inseparable: given a random group element, it should be hard to decide which subgroups contribute non-trivial components to it.

Though composite order bilinear groups have appealing features, it is desirable to obtain the same functionalities and strong guarantees achieved in composite order groups from other assumptions, particularly from the decisional linear assumption (DLIN) in prime order bilinear groups. The ability to work with prime order bilinear groups instead of composite order ones offers several advantages. First, we can obtain security under the more standard decisional linear assumption. Second, we can achieve much more efficient systems for the same security levels. This is because in composite order groups, security typically relies on the hardness of factoring the group order. This requires the use of large group orders, which results in considerably slower pairing operations.

There have been many previous examples of cryptosystems that were first built in composite order groups while later analogs were obtained in prime order groups. These include Groth-Ostrovsky-Sahai proofs [22,21], the Boneh-Sahai-Waters traitor tracing scheme [10,15], and the functional encryption schemes of Lewko-Okamoto-Sahai-Takashima-Waters [25,33]. Waters also notes that the dual system encryption techniques in [38] used to obtain prime order systems were first instantiated in composite order groups. These results already suggest that there are strong parallels between the composite order and prime order settings, but the translation techniques are developed in system-specific ways.

Beyond improving the assumptions and efficiency for particular schemes, our goal in this paper is to expand our general understanding of how tools that are conveniently inherent in the composite order setting can be simulated in the prime order setting. We begin by asking: what are the basic features of composite order bilinear groups that are typically exploited by cryptographic constructions and security proofs? Freeman considers this question in [14] and identifies two such features, called *projecting* and *canceling* (we also refer to canceling as "orthogonality"). Freeman then provides examples of how to construct either of these properties using pairings of vectors of group elements in prime order groups. Notably, Freeman does not provide a way of simultaneously achieving *both* projecting and canceling. There may be good reason for this, since Meiklejohn, Shacham, and Freeman [30] have shown that both properties cannot be simultaneously achieved in prime order groups when one relies on the decisional linear assumption in a "natural way". By instantiating either projecting or canceling in prime order groups, Freeman [14] successfully translates several composite order schemes into prime order schemes: the Boneh-Goh-Nissim encryption scheme [9], the Boneh-Sahai-Waters traitor tracing system [10], and the Katz-Sahai-Waters predicate encryption scheme [24]. These translations use a three step process. The first step is to write the scheme in an abstract framework (replacing subgroups by subspaces of vectors in the exponent), the second step is to translate the assumptions into prime order analogs, and the third step is to transfer the security proof.

There are two aspects of Freeman's approach that can render the results unsatisfying in certain cases. First, the step of translating the assumptions often does not result in standard assumptions like DLIN. A reduction to DLIN is only provided for the most basic variant of the subgroup decision assumption, and does not extend (for example) to the general subgroup decision assumption from [3]. Second, the step of translating the proof fails for many schemes, including all of the recent composite order schemes employing the dual system encryption proof methodology [27,25,29,28,26]. These schemes use only canceling and not projecting, and so this is unrelated to the limitations discussed in [30].

The reason for this failure is instructive to examine. As Freeman points out, "the recent identity-based encryption scheme of Lewko and Waters [27] uses explicitly in its security proof the fact that the group G has two subgroups of relatively prime order". The major obstacle here is not translating the description of the scheme or its assumptions - instead the problem lies in translating a trick in the security proof. The trick works as follows. Suppose we have a group G of order $N = p_1 p_2 \dots p_m$, where $p_1, \dots, p_m$ are distinct primes. Then if we take an element $g_1 \in G$ of order p_1 (i.e. an element of the subgroup of G with order p_1) and a random exponent $a \in \mathbb{Z}_N$, the group element g_1^a *reveals no information* about the value of a modulo the other primes. Only $a \bmod p_1$ is revealed. The fact that $a \bmod p_2$, for instance, is uniformly random even conditioned on $a \bmod p_1$ follows from the Chinese Remainder Theorem. In the security proof of the Lewko-Waters scheme, there are elements of the form g_1^a in the public parameters, and the fact that $a \bmod p_2$ remains information-theoretically hidden is later used to argue that all the keys and ciphertext received by the attacker are properly distributed in the midst of a hybrid argument.

Clearly, in a prime order group, we cannot hope to construct subgroups with coprime orders. There are a few possible paths for resolving this difficulty. We could start by reworking proofs in the composite order setting to avoid using this trick and then hope to apply the techniques of [14] without modification. This approach is likely to result in more complicated (though still static) assumptions in the composite order setting, which will translate into more complicated assumptions in the prime order setting. Since we prefer to rely only on the decisional linear assumption, we follow an alternate strategy: finding a version of this trick in prime order groups that does not rely on coprimeness. This is possible because coprimeness here is used a mechanism for achieving "parameter hiding," meaning that some useful information is information-theoretically hidden from the attacker, even after the public parameters are revealed. We can construct an alternate mechanism in prime order groups that similarly enables a form of parameter hiding.

Our Contribution. We present versatile tools that can be used to translate composite order bilinear systems relying on canceling to prime order bilinear systems, particularly those whose security proofs rely on general subgroup decision assumptions and employ the coprime mechanism discussed above. This includes schemes like [27], which could not be handled by Freeman's methods. Our tools are based in the dual pairing vector space framework initiated by Okamoto and

Takashima [31,32]. We observe that dual pairing vector spaces provide a mechanism for parameter hiding that can be used in place of coprimeness. We then formulate an assumption in prime order groups that can be used to mimic the effect of the general subgroup decision assumption in composite order groups. We prove that this assumption is implied by DLIN. Putting these ingredients together, we obtain a flexible toolkit for turning a class of composite order constructions into prime order constructions that can be proven secure from DLIN.

We demonstrate the use of our toolkit by providing a translation of the composite order Lewko-Waters IBE construction [27]. This yields a prime order IBE construction that is proven fully secure from DLIN and also inherits the intuitive structure of the Boneh-Boyen IBE [5]. Compared to the fully secure prime order IBE construction in [38], our scheme achieves comparable efficiency and security with a simpler structure. As a second application, we provide a translation of the Lewko-Waters unbounded HIBE scheme [29] in the full version. This additionally demonstrates how to handle delegation of secret keys with our tools.

We note that some composite order systems employing dual system encryption, such as the attribute-based encryption scheme in [25], already have analogs in prime order groups proven secure from DLIN using dual pairing vector spaces. In [33], Okamoto and Takashima provide a functional encryption scheme in prime order bilinear groups that is proven fully secure under DLIN. Their construction encompasses both attribute-based and inner product encryption, and their proof relies on dual system encryption techniques, similarly to [25]. While they focus on providing a particular construction and proof, our goal is to formulate a more general strategy for translating composite order schemes into prime order schemes with analogous proofs.

Other Related Work. The concept of identity-based encryption was first proposed by Shamir [36] and later constructed by Boneh and Franklin [8] and Cocks [13]. In an identity-based encryption scheme, users are associated with identities and obtain secret keys from a master authority. Encryption to any identity can be done knowing only the identity and some global public parameters. Both of the initial constructions of IBE were proven secure in the random oracle model. The first standard model constructions, by Canetti, Halevi, and Katz [11] and Boneh and Boyen [5] relied on selective security, which is a more restrictive security model requiring the attacker to announce the identity to be attacker prior to viewing the public parameters. Subsequently, Boneh and Boyen [6], Gentry [16], and Waters [37,38] provided constructions proven fully secure in the standard model from various assumptions. Except for the scheme of [13], which relied on the quadratic residuousity assumption, all of the schemes we have cited above rely on bilinear groups. A lattice-based IBE construction was first provided by Gentry, Peikert, and Vaikuntanathan in [18].

Hierarchical identity-based encryption was proposed by Horwitz and Lynn [23] and then constructed by Gentry and Silverberg [19] in the random oracle model. In a HIBE scheme, users are associated with identity vectors that indicate their places in a hierarchy (a user Alice is a superior of the user Bob if her identity vector is a prefix of his). Any user can obtain a secret key for his identity vector

either from the master authority or from one of his superiors (i.e. a mechanism for key delegation to subordinates is provided). Selectively secure standard model constructions of HIBE were provided by Boneh and Boyen [5] and Boneh, Boyen, and Goh [7] in the bilinear setting and by Cash, Hofheinz, Kiltz, and Peikert [12] and Agrawal, Boneh, and Boyen [1,2] in the lattice-based setting. Fully secure constructions allowing polynomial depth were given by Gentry and Halevi [17], Waters [38], and Lewko and Waters [27]. The first unbounded construction (meaning that the maximal depth is not bounded by the public parameters) was given by Lewko and Waters in [29].

Attribute-based encryption (ABE) is a more flexible functionality than (H)IBE, first introduced by Sahai and Waters in [35]. In an ABE scheme, keys and ciphertexts are associated with attributes and access policies instead of identities. In a ciphertext-policy ABE scheme, keys are associated with attributes and ciphertexts are associated with access policies. In a key-policy ABE scheme, keys are associated with access policies and ciphertexts are associated with attributes. In both cases, a key can decrypt a ciphertext if and only if the attributes satisfy the formula. There are several constructions of both kinds of ABE schemes, e.g. [35,20,34,4,25,33,39].

The dual system encryption methodology was introduced by Waters in [38] as a tool for proving full security of advanced functionalities such as (H)IBE and ABE. It was further developed in several subsequent works [27,25,33,26,29,28]. Most of these works have used composite order groups as a convenient setting for instantiating the dual system methodology, with the exception of [33]. Here, we extend and generalize the techniques of [33] to demonstrate that this use of composite order groups can be viewed as an intermediary step in the development of prime order systems whose security relies on the DLIN assumption.

2 Background

2.1 Composite Order Bilinear Groups

When G is a bilinear group of composite order $N = p_1 p_2 \dots p_m$ (where p_1, p_2, …, p_m are distinct primes), we let $e : G \times G \to G_T$ denote its bilinear map (also referred to as a pairing). We note that both G and G_T are cyclic groups of order N. For each p_i, G has a subgroup of order p_i denoted by G_{p_i}. We let $g_1, \dots, g_m$ denote generators of G_{p_1} through G_{p_m} respectively. Each element $g \in G$ can be expressed as $g = g_1^{a_1} g_2^{a_2} \cdots g_m^{a_m}$ for some $a_1, \dots, a_m \in \mathbb{Z}_N$, where each a_i is unique modulo p_i. We will refer to $g_i^{a_i}$ as the "G_{p_i} component" of g. When a_i is congruent to zero modulo p_i, we say that g has no G_{p_i} component. The subgroups $G_{p_1}, \dots, G_{p_m}$ are "orthogonal" under the bilinear map e, meaning that if $h \in G_{p_i}$ and $u \in G_{p_j}$ for $i \neq j$, then $e(h, u) = 1$, where 1 denotes the identity element in G_T.

General Subgroup Decision Assumption. The general subgroup decision assumption for composite order bilinear groups (formulated in [3]) is a family of static complexity assumptions based on the intuition that it should be hard to determine which components are present in a random group element, except for

what can be trivially determined by testing for orthogonality with other given group elements. More precisely, for each non-empty subset $S \subseteq [m]$, there is an associated subgroup of order $\prod_{i \in S} p_i$ in G, which we will denote by G_S. For two distinct, non-empty subsets S_0 and S_1, we assume it is hard to distinguish a random element of G_{S_0} from a random element of G_{S_1}, when one is only given random elements of $G_{S_2}, \ldots, G_{S_k}$ where for each $2 \leq j \leq k$, either $S_j \cap S_0 = \emptyset = S_j \cap S_1$ or $S_j \cap S_0 \neq \emptyset \neq S_j \cap S_1$.

More formally, we let $\mathcal{G}$ denote a group generation algorithm, which takes in m and a security parameter λ and outputs a bilinear group G of order $N = p_1 \cdots p_m$, where $p_1, \ldots, p_m$ are distinct primes. The General Subgroup Decision Assumption with respect to $\mathcal{G}$ is defined as follows.

Definition 1. *General Subgroup Decision Assumption. Let $S_0, S_1, S_2, \ldots, S_k$ be non-empty subsets of $[m]$ such that for each $2 \leq j \leq k$, either $S_j \cap S_0 = \emptyset = S_j \cap S_1$ or $S_j \cap S_0 \neq \emptyset \neq S_j \cap S_1$. Given a group generator $\mathcal{G}$, we define the following distribution:*

$$\mathbb{G} := (N = p_1 \cdots p_m, G, G_T, e) \xleftarrow{R} \mathcal{G},$$

$$Z_0 \xleftarrow{R} G_{S_0},\ Z_1 \xleftarrow{R} G_{S_1},\ Z_2 \xleftarrow{R} G_{S_2}, \ldots, Z_k \xleftarrow{R} G_{S_k},$$

$$D := (\mathbb{G}, Z_2, \ldots, Z_k).$$

We assume that for any PPT algorithm $\mathcal{A}$ (with output in $\{0,1\}$),

$$Adv_{\mathcal{G},\mathcal{A}} := |\mathbb{P}[\mathcal{A}(D, Z_0) = 1] - \mathbb{P}[\mathcal{A}(D, Z_1) = 1]|$$

is negligible in the security parameter λ.

We note that this assumption holds in the generic group model, assuming it is hard to find a non-trivial factor of the group order N.

Restricting to Challenge Sets Differing by One Element. We observe that it suffices to consider challenge sets S_0 and S_1 of the form $S_1 = S_0 \cup \{i\}$ for some $i \in [m]$, $i \notin S_0$. We refer to this restricted class of subgroup decision assumptions as the 1-General Subgroup Decision Assumption. To see that the 1-general subgroup decision assumption implies the general subgroup decision assumption, we show that any instance of the general subgroup decision assumption is implied by a sequence of the more restricted instances. More precisely, for general S_0, S_1, we let U denote the set $S_0 \cup S_1 - S_0$. For any i in U, the 1-general subgroup decision assumption implies that it hard to distinguish a random element of G_{S_0} from a random element of $G_{S_0 \cup \{i\}}$, even given random elements from $G_{S_2}, \ldots, G_{S_k}$. That is because each of the sets $S_2, \ldots, S_k$ either does not intersect S_1 or S_0 and hence does not intersect S_0 or $S_0 \cup \{i\} \subseteq S_1$, or intersects both S_0 and $S_0 \cup \{i\}$. We can now incrementally add the other elements of U using instances of the 1-general subgroup decision assumption, ultimately showing that it is hard to distinguish a random element of G_{S_0} from a random element of $G_{S_0 \cup S_1}$. We can reverse the process and subtract one element at a time from $S_0 \cup S_1$ until we arrive at S_1. Thus, the seemingly more restrictive 1-general subgroup decision assumption implies the general subgroup decision assumption.

2.2 Prime Order Bilinear Groups

We now let G denote a bilinear group of prime order p, with bilinear map $e : G \times G \to G_T$. More generally, one may have a bilinear map $e : G \times H \to G_T$, where G and H are different groups. For simplicity in this paper, we will always consider groups where $G = H$.

In addition to referring to individual elements of G, we will also consider "vectors" of group elements. For $\boldsymbol{v} = (v_1, \ldots, v_n) \in \mathbb{Z}_p^n$ and $g \in G$, we write $g^{\boldsymbol{v}}$ to denote a n-tuple of elements of G:

$$g^{\boldsymbol{v}} := (g^{v_1}, g^{v_2}, \ldots, g^{v_n}).$$

We can also perform scalar multiplication and vector addition in the exponent. For any $a \in \mathbb{Z}_p$ and $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{Z}_p^n$, we have:

$$g^{a\boldsymbol{v}} := (g^{av_1}, \ldots, g^{av_n}), \quad g^{\boldsymbol{v}+\boldsymbol{w}} = (g^{v_1+w_1}, \ldots, g^{v_n+w_n}).$$

We define e_n to denote the product of the componentwise pairings:

$$e_n(g^{\boldsymbol{v}}, g^{\boldsymbol{w}}) := \prod_{i=1}^{n} e(g^{v_i}, g^{w_i}) = e(g, g)^{\boldsymbol{v} \cdot \boldsymbol{w}}.$$

Here, the dot product is taken modulo p.

Dual Pairing Vector Spaces. We will employ the concept of dual pairing vector spaces from [31,32]. For a fixed (constant) dimension n, we will choose two random bases $\mathbb{B} := (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ and $\mathbb{B}^* := (\boldsymbol{b}_1^*, \ldots, \boldsymbol{b}_n^*)$ of $\mathbb{Z}_p^n$, subject to the constraint that they are "dual orthonormal", meaning that

$$\boldsymbol{b}_i \cdot \boldsymbol{b}_j^* = 0 \pmod{p},$$

whenever $i \neq j$, and

$$\boldsymbol{b}_i \cdot \boldsymbol{b}_i^* = \psi$$

for all i, where ψ is a uniformly random element of $\mathbb{Z}_p$. (This is a slight abuse of the terminology "orthonormal", since ψ is not constrained to be 1.)

For a generator $g \in G$, we note that

$$e_n(g^{\boldsymbol{b}_i}, g^{\boldsymbol{b}_j^*}) = 1$$

whenever $i \neq j$, where 1 here denotes the identity element in G_T.

We note that choosing random dual orthonormal bases $(\mathbb{B}, \mathbb{B}^*)$ can equivalently be thought of as choosing a random basis $\mathbb{B}$, choosing a random vector $\boldsymbol{b}_1^*$ subject to the constraint that it is orthogonal to $\boldsymbol{b}_2, \ldots, \boldsymbol{b}_n$, defining $\psi = \boldsymbol{b}_1 \cdot \boldsymbol{b}_1^*$, and then choosing $\boldsymbol{b}_2^*$ so that it is orthogonal to $\boldsymbol{b}_1, \boldsymbol{b}_3, \ldots, \boldsymbol{b}_n$, and has dot product with $\boldsymbol{b}_2$ equal to ψ, and so on. We will later use the notation $(\mathbb{D}, \mathbb{D}^*)$ and $\boldsymbol{d}_1, \ldots$, etc. to also denote dual orthonormal bases and their vectors (and even $\mathbb{F}, \mathbb{F}^*$ and $\boldsymbol{f}_1$, etc.). This is because we will sometimes be handling more than one pair of dual orthonormal bases at a time, and we use different notation to avoid confusing them.

Decisional Linear Assumption. The complexity assumption we will rely on in prime order bilinear groups is the Decisional Linear Assumption. To define this formally, we let $\mathcal{G}$ denote a group generation algorithm, which takes in a security parameter λ and outputs a bilinear group G of order p.

Definition 2. *Decisional Linear Assumption. Given a group generator $\mathcal{G}$, we define the following distribution:*

$$\mathbb{G} := (p, G, G_T, e) \xleftarrow{R} \mathcal{G},$$

$$g, f, v, w \xleftarrow{R} G,\ c_1, c_2, w \xleftarrow{R} \mathbb{Z}_p,$$

$$D := (g, f, v, f^{c_1}, v^{c_2}).$$

We assume that for any PPT algorithm $\mathcal{A}$ (with output in $\{0,1\}$),

$$Adv_{\mathcal{G},\mathcal{A}} := \left|\mathbb{P}\left[\mathcal{A}(D, g^{c_1+c_2}) = 1\right] - \mathbb{P}\left[\mathcal{A}(D, g^{c_1+c_2+w}) = 1\right]\right|$$

is negligible in the security parameter λ.

3 Our Main Tools

There is an additional feature of composite order groups that is often exploited along with canceling/orthogonality in the security proofs for composite order constructions: we call this *parameter hiding*. In composite order groups, parameter hiding takes the following form. Consider a composite order group G of order $N = p_1p_2$ and an element $g_1 \in G_{p_1}$ (an element of order p_1). Then if we sample a uniformly random exponent $a \in Z_N$ and produce g_1^a, this reveals nothing about the value of a modulo p_2. More precisely, the Chinese Remainder theorem guarantees that the value of a modulo p_2 conditioned on the value of a modulo p_1 is still uniformly random, and g_1^a only depends on the value of a modulo p_1. This allows a party choosing a to publish g_1^a and still *hide* some information about a, namely its value modulo p_2. Note that this party only needs to know N and g_1: it does not need to know the factorization of N.

This is an extremely useful tool in security proofs, enabling a simulator to choose some secret random exponents, publish the public parameters by raising known subgroup elements to these exponents, and still information-theoretically hide the values of these exponents modulo some of the primes. These hidden values can be leveraged later in the security game to argue that something looks well-distributed in the attacker's view, even if this does not hold in the simulator's view. This sort of trick is crucial in proofs employing the dual system encryption methodology.

Replicating this trick in prime order groups seems challenging, since if one is given g and g^a in a prime order group, a is completely revealed modulo p in an information-theoretic sense. To resolve this issue, we use dual pairing vector spaces. We observe that a form of parameter hiding is achieved by using dual orthonormal bases: one can generate a random pair of dual orthonormal bases

$(\mathbb{B}, \mathbb{B}^*)$ for $\mathbb{Z}_p^n$, apply an invertible change of basis matrix A to a subset of these basis vectors, and produce a new pair of dual orthonormal bases which is also randomly distributed, *independently of* A. This allows us to *hide* a random matrix A. We formulate this precisely below.

3.1 Parameter Hiding in Dual Orthonormal Bases

We consider taking dual orthonormal bases and applying a linear change of basis to a subset of their vectors. We do this in such a way that we produce new dual orthonormal bases. In this subsection, we prove that if we start with randomly sampled dual orthonormal bases, then the resulting bases will also be random - in particular, the distribution of the final bases reveals nothing about the change of basis matrix that was employed. This "hidden" matrix can then be leveraged in security proofs as a way of separating the simulator's view from the attacker's.

To describe this formally, we let $m \leq n$ be fixed positive integers and $A \in \mathbb{Z}_p^{m \times m}$ be an invertible matrix. We let $S_m \subseteq [n]$ be a subset of size m ($|S| = m$). For any dual orthonormal bases $\mathbb{B}, \mathbb{B}^*$, we can then define new dual orthonormal bases $\mathbb{B}_A, \mathbb{B}_A^*$ as follows. We let B_m denote the $n \times m$ matrix over $\mathbb{Z}_p$ whose columns are the vectors $\boldsymbol{b}_i \in \mathbb{B}$ such that $i \in S_m$. Then $B_m A$ is also an $n \times m$ matrix. We form $\mathbb{B}_A$ by retaining all of the vectors $\boldsymbol{b}_i \in \mathbb{B}$ for $i \notin S_m$ and exchanging the $\boldsymbol{b}_i$ for $i \in S_m$ with the columns of $B_m A$. To define $\mathbb{B}_A^*$, we similarly let B_m^* denote the $n \times m$ matrix over $\mathbb{Z}_p$ whose columns are the vectors $\boldsymbol{b}_i^* \in \mathbb{B}^*$ such that $i \in S_m$. Then $B_m^*(A^{-1})^t$ is also an $n \times m$ matrix, where $(A^{-1})^t$ denotes the transpose of A^{-1}. We form $\mathbb{B}_A^*$ by retaining all of the vectors $\boldsymbol{b}_i^* \in \mathbb{B}^*$ for $i \notin S_m$ and exchanging the $\boldsymbol{b}_i$ for $i \in S_m$ with the columns of $B_m^*(A^{-1})^t$.

To see that $\mathbb{B}_A$ and $\mathbb{B}_A^*$ are dual orthonormal bases, note that for $i \in S_m$, the corresponding basis vector in $\mathbb{B}_A$ can be expressed as a linear combination of the basis vectors $\boldsymbol{b}_j \in \mathbb{B}$ with $j \in S_m$, and the coefficients of this linear combination correspond to a column of A, say the ℓ^{th} column (equivalently, say i is the ℓ^{th} element of S_m). When $\ell \neq \ell'$, the ℓ^{th} column of A is orthogonal to the $(\ell')^{th}$ column of $(A^{-1})^t$. This means that the i^{th} vector of $\mathbb{B}_A$ will be orthogonal to the $(i')^{th}$ vector of $\mathbb{B}_A^*$ whenever $i \neq i'$. Moreover, the ℓ^{th} column of A and the ℓ^{th} column of $(A^{-1})^t$ have dot product equal to 1, so the dot product of the i^{th} vector of $\mathbb{B}_A$ and the i^{th} vector of $\mathbb{B}_A^*$ will be equal to the same value ψ as in the original bases $\mathbb{B}$ and $\mathbb{B}^*$.

For a fixed dimension n and prime p, we let $(\mathbb{B}, \mathbb{B}^*) \xleftarrow{R} \mathit{Dual}(\mathbb{Z}_p^d)$ denote choosing random dual orthonormal bases $\mathbb{B}$ and $\mathbb{B}^*$ of $\mathbb{Z}_p^n$. Here, $\mathit{Dual}(\mathbb{Z}_p^n)$ denotes the set of dual orthonormal bases.

Lemma 1. *For any fixed positive integers $m \leq n$, any fixed invertible $A \in \mathbb{Z}_p^{m \times m}$ and set $S_m \subseteq [n]$ of size m, if $(\mathbb{B}, \mathbb{B}^*) \xleftarrow{R} \mathit{Dual}(\mathbb{Z}_p^d)$, then $(\mathbb{B}_A, \mathbb{B}_A^*)$ is also distributed as a random sample from $\mathit{Dual}(\mathbb{Z}_p^d)$. In particular, the distribution of $(\mathbb{B}_A, \mathbb{B}_A^*)$ is independent of A.*

Proof. There is a one-to-one correspondence between $(\mathbb{B}, \mathbb{B}^*)$ and $(\mathbb{B}_A, \mathbb{B}_A^*)$: given $(\mathbb{B}_A, \mathbb{B}_A^*)$, one can recover $(\mathbb{B}, \mathbb{B}^*)$ by applying A^{-1} to the vectors in $\mathbb{B}_A$ whose

indices are in S_m, and applying A^t to the corresponding vectors in $\mathbb{B}_A^*$. This shows that every pair of dual orthonormal bases is equally likely to occur as $\mathbb{B}_A, \mathbb{B}_A^*$.

3.2 The Subspace Assumption

We now state a complexity assumption in prime order groups that we will use to simulate the effects of subgroup decision assumptions in composite order groups. We call this the Subspace Assumption. In the full version, we show that the subspace assumption is implied by the decisional linear assumption.

In prime order groups, basis vectors in the exponent take the place of subgroups. Since we are using dual orthonormal bases, our new concept of orthogonality between "subgroups" becomes asymmetric. If we have dual orthonormal bases $\mathbb{B}, \mathbb{B}^*$ and we think of "subgroup 1" in $\mathbb{B}$ as corresponding to the span of $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_4$, then this is not orthogonal to the other vectors in $\mathbb{B}$, but it is orthogonal to vectors $\boldsymbol{b}_5^*, \ldots, \boldsymbol{b}_n^*$ in $\mathbb{B}^*$. Essentially, the notion of a single subgroup has now been split into a pair of "subgroups", one for each side of the pairing, and orthogonality between different subgroups now only holds for elements on opposite sides.

This sort of asymmetry can be quite useful. For example, consider an instance of the general subgroup decision assumption in composite order groups, where the task is to distinguish a random element of G_{p_1} from $G_{p_1 p_2}$. In this case, we cannot give out an element of G_{p_2}, since it can trivially be used to break the assumption by pairing it with the challenge term and seeing if the result is the identity. If we instead use dual orthonormal bases in a prime order group, the situation is a bit different. Suppose that given $g^{\boldsymbol{v}}$, the task is to distinguish whether the exponent vector $\boldsymbol{v}$ is in the span of $\boldsymbol{b}_1^*, \boldsymbol{b}_2^*$ or in the larger span of $\boldsymbol{b}_1^*, \boldsymbol{b}_2^*, \boldsymbol{b}_3^*$. We cannot give out $g^{\boldsymbol{b}_3}$, since one could then break the assumption by testing if $e_n(g^{\boldsymbol{v}}, g^{\boldsymbol{b}_3}) = e(g,g)^{\boldsymbol{v} \cdot \boldsymbol{b}_3}$ is the identity, but *we can give out* $g^{\boldsymbol{b}_3^*}$.

Our definition of the subspace assumption is motivated by this and our observation in Section 2.1 that the general subgroup decision assumption in composite order groups can be restricted to distinguishing between sets that differ by one element. What this means is that to simulate the uses of the general subgroup decision in composite order groups, one can focus merely on creating an analog for expansion into one new "subgroup" at a time. At its core, our subspace assumption says that if one is given $g^{\boldsymbol{v}}$, then it is hard to tell if $\boldsymbol{v}$ is randomly chosen from the span of $\boldsymbol{b}_1^*, \boldsymbol{b}_2^*$ or from the larger span of $\boldsymbol{b}_1^*, \boldsymbol{b}_2^*, \boldsymbol{b}_3^*$, even if one is given scalar multiples of all bases vectors in $\mathbb{B}$ and $\mathbb{B}^*$ in the exponent, *except for* $\boldsymbol{b}_3$. We augment this by also given out a random linear combination of $\boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3$ in the exponent. We then generalize this by replicating the same structure for k 3-tuples of vectors, with the random linear combinations having the *same* coefficients. (The fact that these coefficients are the same prevents this from following immediately from the assumption for a single 3-tuple applied in hybrid fashion.)

We now give the formal description of the subspace assumption. For a fixed dimension $n \geq 3$ and prime p, we recall that $(\mathbb{B}, \mathbb{B}^*) \xleftarrow{R} \mathit{Dual}(\mathbb{Z}_p^n)$ denotes choosing random dual orthonormal bases $\mathbb{B}$ and $\mathbb{B}^*$ of $\mathbb{Z}_p^n$, and $\mathit{Dual}(\mathbb{Z}_p^n)$ denotes the set of dual orthonormal bases. Our assumption is additionally parameterized by a positive integer $k \leq \frac{n}{3}$.

Definition 3. *(Subspace Assumption) Given a group generator $\mathcal{G}$, we define the following distribution:*

$$\mathbb{G} := (p, G, G_T, e) \xleftarrow{R} \mathcal{G},\ (\mathbb{B}, \mathbb{B}^*) \xleftarrow{R} Dual(\mathbb{Z}_p^n),$$

$$g \xleftarrow{R} G,\ \eta, \beta, \tau_1, \tau_2, \tau_3, \mu_1, \mu_2, \mu_3 \xleftarrow{R} \mathbb{Z}_p,$$

$$U_1 := g^{\mu_1 \boldsymbol{b}_1 + \mu_2 \boldsymbol{b}_{k+1} + \mu_3 \boldsymbol{b}_{2k+1}},\ U_2 := g^{\mu_1 \boldsymbol{b}_2 + \mu_2 \boldsymbol{b}_{k+2} + \mu_3 \boldsymbol{b}_{2k+2}},\ \ldots,$$

$$U_k := g^{\mu_1 \boldsymbol{b}_k + \mu_2 \boldsymbol{b}_{2k} + \mu_3 \boldsymbol{b}_{3k}}, V_1 := g^{\tau_1 \eta \boldsymbol{b}_1^* + \tau_2 \beta \boldsymbol{b}_{k+1}^*},\ V_2 := g^{\tau_1 \eta \boldsymbol{b}_2^* + \tau_2 \beta \boldsymbol{b}_{k+2}^*},\ \ldots,$$

$$V_k := g^{\tau_1 \eta \boldsymbol{b}_k^* + \tau_2 \beta \boldsymbol{b}_{2k}^*},\ W_1 := g^{\tau_1 \eta \boldsymbol{b}_1^* + \tau_2 \beta \boldsymbol{b}_{k+1}^* + \tau_3 \boldsymbol{b}_{2k+1}^*},$$

$$W_2 := g^{\tau_1 \eta \boldsymbol{b}_2^* + \tau_2 \beta \boldsymbol{b}_{k+2}^* + \tau_3 \boldsymbol{b}_{2k+2}^*},\ \ldots, W_k := g^{\tau_1 \eta \boldsymbol{b}_k^* + \tau_2 \beta \boldsymbol{b}_{2k}^* + \tau_3 \boldsymbol{b}_{3k}^*}$$

$$D := \big(g^{\boldsymbol{b}_1}, g^{\boldsymbol{b}_2}, \ldots, g^{\boldsymbol{b}_{2k}}, g^{\boldsymbol{b}_{3k+1}}, \ldots, g^{\boldsymbol{b}_n}, g^{\eta \boldsymbol{b}_1^*}, \ldots, g^{\eta \boldsymbol{b}_k^*},$$

$$g^{\beta \boldsymbol{b}_{k+1}^*}, \ldots, g^{\beta \boldsymbol{b}_{2k}^*}, g^{\boldsymbol{b}_{2k+1}^*}, \ldots, g^{\boldsymbol{b}_n^*}, U_1, U_2, \ldots, U_k, \mu_3\big).$$

We assume that for any PPT algorithm $\mathcal{A}$ (with output in $\{0,1\}$),

$$Adv_{\mathcal{G},\mathcal{A}} := |\mathbb{P}\left[\mathcal{A}(D, V_1, \ldots, V_k) = 1\right] - \mathbb{P}\left[\mathcal{A}(D, W_1, \ldots, W_k) = 1\right]|$$

is negligible in the security parameter λ.

We have included in D more terms than will be necessary for many applications of this assumption. We will work exclusively with the $k = 1$ and $k = 2$ cases. We present the assumption in the form above to make it more versatile for use in future applications. We additionally note that the form stated above can be further generalized to involve multiple, independently generated dual orthonormal bases $(\mathbb{B}_1, \mathbb{B}_1^*), (\mathbb{B}_2, \mathbb{B}_2^*), \ldots, (\mathbb{B}_j, \mathbb{B}_j^*)$, for any fixed j. The terms in the assumption would be duplicated for each pair of bases, with the same values of $\eta, \beta, \tau_1, \tau_2, \tau_3, \mu_1, \mu_2, \mu_3$. We will not need this generalization for the applications we present. To help the reader see the main structure of this assumption through the burdensome notation, we include a heuristic illustration of the $k = 2$ case.

In the diagram, the top rows illustrate the U terms, while the bottom rows illustrate the V, W terms. The solid ovals and rectangles indicate the presence of basis vectors. The crossed rectangles indicate basis elements of $\mathbb{B}$ which are present in U_1, U_2 but are not given out in isolation. The dotted ovals adorned by question marks indicate the basis vectors whose presence depends on whether we consider the V's or the W's.

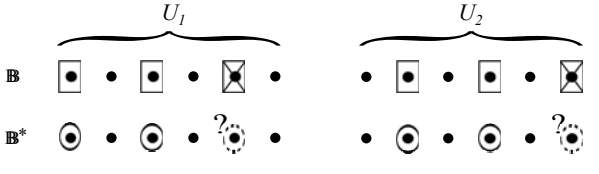

Fig. 1. Subspace Assumption with $k = 2$

4 Analog of the Boneh-Boyen IBE Scheme

In this section, we employ our subspace assumption and our parameter hiding technique for dual orthonormal bases to prove full security for a close analog of the Boneh-Boyen IBE scheme from the decisional linear assumption. This is the same security guarantee achieved for the IBE scheme in [38] and our efficiency is also similar. The advantage of our scheme is that it is a much closer analog to the original Boneh-Boyen IBE, and resultingly has a simpler, more intuitive structure.

Our security proof essentially mirrors the structure of the security proof given in [27], which provides a fully secure variant of the Boneh-Boyen IBE scheme in composite order groups. This serves as an illustrative example of how our techniques can be used to simulate dual system encryption proofs in the prime order setting that were originally presented in composite order groups.

4.1 Our Construction

We will use dual orthonormal bases $(\mathbb{D}, \mathbb{D}^*)$ of $\mathbb{Z}_p^6$, where p is the prime order of our bilinear group G. Public parameters and ciphertexts will have exponents described in terms of the basis vectors in $\mathbb{D}$, while secret keys will have exponents described in terms of $\mathbb{D}^*$. The first four basis vectors of each will constitute the "normal space" (like G_{p_1} in the LW scheme), and the last two basis vectors of each will constitute the "semi-functional space" (like G_{p_2} in the LW scheme).

By using dual pairing vector spaces, we avoid the need to simulate G_{p_3}. In the LW scheme, the purpose of G_{p_3} is to allow the creation of other semi-functional keys while a challenge key is changing from normal to semi-functional. More precisely, it allows the subgroup decision assumption to give out an element of $G_{p_2 p_3}$ that can be used to generate semi-functional keys when the task is to distinguish a random element of $G_{p_1 p_3}$ from a random element of G. We note that if we did not use G_{p_3} here and instead tried to create all of the semi-functional keys from a term in $G_{p_1 p_2}$, then these keys would not be properly randomized in the G_{p_2} subgroup because the structure of the scheme is enforced in the G_{p_1} subgroup. Pairwise independence cannot save us here because there are many keys. However, the asymmetry of dual pairing vector spaces avoids this issue: while we are expanding the challenge key into the "semi-functional space" in $\mathbb{D}^*$, we can still know a basis for the semi-functional space of $\mathbb{D}^*$ in the exponent - it is only the corresponding terms in the semi-functional space of $\mathbb{D}$ that we do not have access to in isolation. This allows us to make the other semi-functional keys without needing to create an analog of the G_{p_3} subgroup.

The core of the Boneh-Boyen scheme is a cancelation between terms in two pairings, one with the identity appearing on the ciphertext side and the other with the identity appearing on the key side. This is combined with a mechanism for preventing multiplication manipulation of the identity. In our scheme, this core cancelation is duplicated: instead of having one cancelation, we have two, each with its own random coefficients. The first cancelation will occur for the $\boldsymbol{d}_1, \boldsymbol{d}_2$ and $\boldsymbol{d}_1^*, \boldsymbol{d}_2^*$ components, and the second will occur for the $\boldsymbol{d}_3, \boldsymbol{d}_4$ and $\boldsymbol{d}_3^*, \boldsymbol{d}_4^*$ components.

This expansion gives us room to use the subspace assumption with parameter $k = 2$ to transition from 4-dimensional exponents for normal keys and ciphertexts to 6-dimensional exponents for semi-functional keys and ciphertexts. Having a 2-dimensional semi-functional space allows us to implement nominal semi-functionality. To prevent multiplicative manipulations of the identities in our scheme is rather easy, since the orthogonality of the dual bases allows us to "tie" all the components of the keys and ciphertexts together without causing cross interactions that interfere with decryption.

We assume that messages M are elements of G_T (the target group of the bilinear map) and that identities ID are elements of $\mathbb{Z}_p$.

Setup$(\lambda) \rightarrow \mathrm{MSK}, \mathrm{PP}$. The setup algorithm takes in the security parameter λ and chooses a bilinear group G of sufficiently large prime order p. We let $e : G \times G \rightarrow G_T$ denote the bilinear map. We set $n = 6$. The algorithm samples random dual orthonormal bases, $(\mathbb{D}, \mathbb{D}^*) \xleftarrow{R} Dual(\mathbb{Z}_p^n)$. We let $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_6$ denote the elements of $\mathbb{D}$ and $\boldsymbol{d}_1^*, \ldots, \boldsymbol{d}_6^*$ denote the elements of $\mathbb{D}^*$. It also chooses random values $\alpha, \theta, \sigma \in \mathbb{Z}_p$. The public parameters are computed as:

$$\mathrm{PP} := \left\{ G, p, e(g,g)^{\alpha\theta \boldsymbol{d}_1 \cdot \boldsymbol{d}_1^*}, g^{\boldsymbol{d}_1}, \ldots, g^{\boldsymbol{d}_4} \right\}.$$

(We note that $\boldsymbol{d}_1 \cdot \boldsymbol{d}_1^* = \psi$ by definition of $\mathbb{D}, \mathbb{D}^*$, but we write out the dot product when we feel it is more instructive.) The master secret key is:

$$\mathrm{MSK} := \left\{ g^{\theta \boldsymbol{d}_1^*}, g^{\alpha\theta \boldsymbol{d}_1^*}, g^{\theta \boldsymbol{d}_2^*}, g^{\sigma \boldsymbol{d}_3^*}, g^{\sigma \boldsymbol{d}_4^*} \right\}.$$

KeyGen$(\mathrm{MSK}, ID) \rightarrow \mathrm{SK}_{ID}$. The key generation algorithm chooses random values $r_1, r_2 \in \mathbb{Z}_p$ and forms the secret key as:

$$\mathrm{SK}_{ID} := g^{(\alpha + r_1 ID)\theta \boldsymbol{d}_1^* - r_1 \theta \boldsymbol{d}_2^* + r_2 ID \sigma \boldsymbol{d}_3^* - r_2 \sigma \boldsymbol{d}_4^*}.$$

Encrypt$(M, ID, \mathrm{PP}) \rightarrow \mathrm{CT}$. The encryption algorithm chooses random values $s_1, s_2 \in \mathbb{Z}_p$ and forms the ciphertext as:

$$\mathrm{CT} := \left\{ C_1 := M \left(e(g,g)^{\alpha\theta \boldsymbol{d}_1 \cdot \boldsymbol{d}_1^*} \right)^{s_1}, \;\; C_2 := g^{s_1 \boldsymbol{d}_1 + s_1 ID \boldsymbol{d}_2 + s_2 \boldsymbol{d}_3 + s_2 ID \boldsymbol{d}_4} \right\}.$$

Decrypt$(\mathrm{CT}, \mathrm{SK}_{ID}) \rightarrow M$. The decryption algorithm computes the message as:

$$M := C_1 / e_n(\mathrm{SK}_{ID}, C_2).$$

Recall that $n = 6$, so this requires six pairings.

4.2 Semi-functional Algorithms

We choose to define our semi-functional objects by providing algorithms that generate them. We note that these algorithms are only provided for definitional purposes, and are not part of the IBE system. In particular, they do not need to be efficiently computable from the public parameters and master secret key alone.

KeyGenSF. The semi-functional key generation algorithm chooses random values $r_1, r_2, t_5, t_6 \in \mathbb{Z}_p$ and forms the secret key as

$$\mathrm{SK}_{ID} := g^{(\alpha + r_1 ID)\theta \boldsymbol{d}_1^* - r_1 \theta \boldsymbol{d}_2^* + r_2 ID \sigma \boldsymbol{d}_3^* - r_2 \sigma \boldsymbol{d}_4^* + t_5 \boldsymbol{d}_5^* + t_6 \boldsymbol{d}_6^*}.$$

This is distributed like a normal key with additional random multiples of $\boldsymbol{d}_5^*$ and $\boldsymbol{d}_6^*$ added in the exponent.

EncryptSF. The semi-functional encryption algorithm chooses random values $s_1, s_2, z_5, z_6 \in \mathbb{Z}_p$ and forms the ciphertext as:

$$\mathrm{CT} := \left\{ C_1 := M \left(e(g,g)^{\alpha\theta \boldsymbol{d}_1 \cdot \boldsymbol{d}_1^*} \right)^{s_1}, C_2 := g^{s_1 \boldsymbol{d}_1 + s_1 ID \boldsymbol{d}_2 + s_2 \boldsymbol{d}_3 + s_2 ID \boldsymbol{d}_4 + z_5 \boldsymbol{d}_5 + z_6 \boldsymbol{d}_6} \right\}.$$

This is distributed like a normal ciphertext with additional random multiples of $\boldsymbol{d}_5$ and $\boldsymbol{d}_6$ added in the exponent.

We observe that if one applies the decryption procedure with a semi-functional key and a normal ciphertext, decryption will succeed because $\boldsymbol{d}_5^*, \boldsymbol{d}_6^*$ are orthogonal to all of the vectors in exponent of C_2, and hence have no effect on decryption. Similarly, decryption of a semi-functional ciphertext by a normal key will also succeed because $\boldsymbol{d}_5, \boldsymbol{d}_6$ are orthogonal to all of the vectors in the exponent of the key. When *both* the ciphertext and key are semi-functional, the result of $e_n(\mathrm{SK}_{ID}, C_2)$ will have an additional term, namely $e(g,g)^{t_5 z_5 \boldsymbol{d}_5 \cdot \boldsymbol{d}_5^* + t_6 z_6 \boldsymbol{d}_6 \cdot \boldsymbol{d}_6^*} = e(g,g)^{(t_5 z_5 + t_6 z_6)\psi}$. Decryption will then fail unless $t_5 z_5 + t_6 z_6 \equiv 0 \bmod p$. If this modular equation holds, we say that the key and ciphertext pair is *nominally semi-functional*. We note that this is possible, even when none of t_5, z_5, t_6, z_6 are congruent to zero modulo p (this is why we have designated a semi-functional space of dimension two).

In the full version, we prove the following theorem. Here, we sketch the outline of the proof.

Theorem 1. *Under the decisional linear assumption, the IBE scheme presented in Section 4.1 is fully secure.*

We prove this using a hybrid argument over a sequence of games, following the LW strategy. We start with the real security game, denoted by Game_{real}. We let q denote the number of keys requested by the attacker. We define the following additional games.

Game_i for $i = 0, 1, \ldots, q$. Game_i is like Game_{real}, except the ciphertext given to the attacker is semi-functional (i.e. generated by a call to EncryptSF instead of Encrypt) and the first i keys given to the attacker are semi-functional (generated by KeyGenSF). The remaining keys are normal. We note that in Game_0, all of the keys are normal, and in Game_q, all of the keys are semi-functional.

Game_{final}. Game_{final} is like Game_q, except that the ciphertext is a semi-functional encryption of a *random* message in G_T, instead of one of the messages supplied by the attacker.

We transition from Game_{real} to Game_0, then to Game_1, and so on, until we arrive at Game_q. We prove that with each transition, the attacker's advantage cannot change by a non-negligible amount. As a last step, we transition to Game_{final}, where it is clear that the attacker's advantage is zero. These transitions are accomplished in the following lemmas, all using the subspace assumption. We let $Adv_{\mathcal{A}}^{real}$ denote the advantage of an algorithm $\mathcal{A}$ in the real game, $Adv_{\mathcal{A}}^{i}$ denote its advantage in Game_i, and $Adv_{\mathcal{A}}^{final}$ denote its advantage in Game_{final}.

We begin with the transition from Game_{real} to Game_0. At the analogous step in the LW proof, a subgroup decision assumption is used to expand the ciphertext from G_{p_1} into $G_{p_1p_2}$. Here, we use the subspace assumption with $k = 2$ to expand the ciphertext exponent vector from the span of $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_4$ into the larger span of $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_6$. We use a very basic instance of the parameter hiding technique to argue that the resulting coefficients of $\boldsymbol{d}_5$ and $\boldsymbol{d}_6$ are randomly distributed: this is done by initially embedding a random 2×2 change of basis matrix A into our setting of the basis vectors $\boldsymbol{d}_5, \boldsymbol{d}_6$.

We now handle the transition from Game_{i-1} to Game_i. At this step in the LW proof, a subgroup decision assumption is used to expand the i^{th} secret key from $G_{p_1p_3}$ into $G = G_{p_1p_2p_3}$. Analogously, we will use the subspace assumption to expand the i^{th} secret key exponent vector from the span of $\boldsymbol{d}_1^*, \ldots, \boldsymbol{d}_4^*$ into the larger span of $\boldsymbol{d}_1^*, \ldots, \boldsymbol{d}_6^*$. We will embed a 2×2 change of basis matrix A and set $\mathbb{D} = \mathbb{B}_A$ and $\mathbb{D}^* = \mathbb{B}_A^*$, where A is applied to $\boldsymbol{b}_5, \boldsymbol{b}_6$ to form $\boldsymbol{d}_5, \boldsymbol{d}_6$. As in the LW proof, we cannot be given an object that resides solely in the semi-functional space of the ciphertext (e.g. we cannot be given $g^{\boldsymbol{d}_5}, g^{\boldsymbol{d}_6}$), but we are given objects that have semi-functional components attached to normal components, and we can use these to create the semi-functional ciphertext. In the LW proof, a term in $G_{p_1p_2}$ in used. Here, an exponent vector that is a linear combination of $\boldsymbol{b}_1, \boldsymbol{b}_3, \boldsymbol{b}_5$ and another exponent vector that is a linear combination of $\boldsymbol{b}_2, \boldsymbol{b}_4, \boldsymbol{b}_6$ are used. In our case, making the other normal and semi-functional keys is straightforward, since we are given scalar multiples of all of the vectors of $\mathbb{D}^*$ in the exponent. We use the fact that the matrix A is hidden from the attacker in order to argue that the semi-functional parts of the ciphertext and i^{th} key appear well-distributed.

The final step of the LW proof uses an assumption that it is not technically an instance of the general subgroup decision assumption, but is of a similar flavor. In our case, we use a slightly different strategy: we use the subspace assumption with $k = 1$ twice to randomize each appearance of s_1 in the C_2 term of the ciphertext, thereby severing its link with the blinding factor. The end result is the same - we obtain a semi-functional encryption of a random message. This randomization of s_1 is accomplished by first expanding an exponent vector from the span of $\boldsymbol{d}_5, \boldsymbol{d}_6$ into the larger span of $\boldsymbol{d}_5, \boldsymbol{d}_6, \boldsymbol{d}_2$ and then expanding an exponent vector from the span of $\boldsymbol{d}_5, \boldsymbol{d}_6$ into the larger span of $\boldsymbol{d}_5, \boldsymbol{d}_6, \boldsymbol{d}_1$. We note that the knowledge of the μ_3 value in the subspace assumption is used here to ensure that while we are doing the first expansion, for example, we can make the two occurrences of r_1 in the keys match consistently (this is necessary because $g^{\boldsymbol{d}_2^*}$ by itself will not be known during this step).

5 Further Applications

As a second demonstration of our tools, in the full version of this paper we consider a variant of the Lewko-Waters unbounded HIBE construction [29]. The composite order construction we present is simpler than the one presented in [29], at the cost of using more subgroups. Since we will ultimately simulate these subgroups in a prime order group, such a cost is no longer a significant detriment. In designing our prime order translation and proof, we will proceed along a path that is very similar to the path we took to translate the more basic IBE scheme. However, we now must take care to preserve delegation ability throughout our proof. As a result, we employ a different strategy for the final step of the proof. The details of our composite order construction, its prime order translation, and security proofs in both settings can be found in the full version of this paper.

In applying our tools to the both IBE and unbounded HIBE applications, we see that there is some flexibility in how we choose the construction, organize the hybrid games, and embed the subspace assumption in our reductions. All of these considerations interact, allowing us to make tradeoffs. The amount of flexibility available in applying our tools make them suitably versatile to handle a wider variety of applications as well. In particular, they can be applied in the attribute-based encryption setting. We suspect that applying our techniques to the composite order ABE constructions in [25] would result in a system and proof quite similar to the functional encryption schemes presented by Okamoto and Takashima in [33], who obtain security from the decisional linear assumption through dual pairing vector spaces.

References

1. Agrawal, S., Boneh, D., Boyen, X.: Efficient Lattice (H)IBE in the Standard Model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg (2010)
2. Agrawal, S., Boneh, D., Boyen, X.: Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98–115. Springer, Heidelberg (2010)
3. Bellare, M., Waters, B., Yilek, S.: Identity-Based Encryption Secure against Selective Opening Attack. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 235–252. Springer, Heidelberg (2011)
4. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: Proceedings of the IEEE Symposium on Security and Privacy, pp. 321–334.
5. Boneh, D., Boyen, X.: Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 223–238. Springer, Heidelberg (2004)
6. Boneh, D., Boyen, X.: Secure Identity Based Encryption Without Random Oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (2004)
7. Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical Identity Based Encryption with Constant Size Ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)

8. Boneh, D., Franklin, M.: Identity-Based Encryption from the Weil Pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
9. Boneh, D., Goh, E.-J., Nissim, K.: Evaluating 2-DNF Formulas on Ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Heidelberg (2005)
10. Boneh, D., Sahai, A., Waters, B.: Fully Collusion Resistant Traitor Tracing with Short Ciphertexts and Private Keys. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 573–592. Springer, Heidelberg (2006)
11. Canetti, R., Halevi, S., Katz, J.: A Forward-Secure Public-Key Encryption Scheme. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 255–271. Springer, Heidelberg (2003)
12. Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai Trees, or How to Delegate a Lattice Basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer, Heidelberg (2010)
13. Cocks, C.: An Identity Based Encryption Scheme Based on Quadratic Residues. In: Honary, B. (ed.) Cryptography and Coding 2001. LNCS, vol. 2260, pp. 360–363. Springer, Heidelberg (2001)
14. Freeman, D.M.: Converting Pairing-Based Cryptosystems from Composite-Order Groups to Prime-Order Groups. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 44–61. Springer, Heidelberg (2010)
15. Garg, S., Kumarasubramanian, A., Sahai, A., Waters, B.: Building efficient fully collusion-resilient traitor tracing and revocation schemes. In: ACM Conference on Computer and Communications Security, pp. 121–130 (2010)
16. Gentry, C.: Practical Identity-Based Encryption Without Random Oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 445–464. Springer, Heidelberg (2006)
17. Gentry, C., Halevi, S.: Hierarchical Identity Based Encryption with Polynomially Many Levels. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 437–456. Springer, Heidelberg (2009)
18. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing, pp. 197–206 (2008)
19. Gentry, C., Silverberg, A.: Hierarchical ID-Based Cryptography. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 548–566. Springer, Heidelberg (2002)
20. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute based encryption for fine-grained access control of encrypted data. In: ACM Conference on Computer and Communications Security, pp. 89–98 (2006)
21. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive Zaps and New Techniques for NIZK. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 97–111. Springer, Heidelberg (2006)
22. Groth, J., Ostrovsky, R., Sahai, A.: Perfect Non-interactive Zero Knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg (2006)
23. Horwitz, J., Lynn, B.: Toward Hierarchical Identity-Based Encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 466–481. Springer, Heidelberg (2002)
24. Katz, J., Sahai, A., Waters, B.: Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 146–162. Springer, Heidelberg (2008)

25. Lewko, A., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully Secure Functional Encryption: Attribute-Based Encryption and (Hierarchical) Inner Product Encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 62–91. Springer, Heidelberg (2010)
26. Lewko, A., Rouselakis, Y., Waters, B.: Achieving Leakage Resilience through Dual System Encryption. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 70–88. Springer, Heidelberg (2011)
27. Lewko, A., Waters, B.: New Techniques for Dual System Encryption and Fully Secure HIBE with Short Ciphertexts. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 455–479. Springer, Heidelberg (2010)
28. Lewko, A., Waters, B.: Decentralizing Attribute-Based Encryption. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 568–588. Springer, Heidelberg (2011)
29. Lewko, A., Waters, B.: Unbounded HIBE and Attribute-Based Encryption. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 547–567. Springer, Heidelberg (2011)
30. Meiklejohn, S., Shacham, H., Freeman, D.M.: Limitations on Transformations from Composite-Order to Prime-Order Groups: The Case of Round-Optimal Blind Signatures. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 519–538. Springer, Heidelberg (2010)
31. Okamoto, T., Takashima, K.: Homomorphic Encryption and Signatures from Vector Decomposition. In: Galbraith, S.D., Paterson, K.G. (eds.) Pairing 2008. LNCS, vol. 5209, pp. 57–74. Springer, Heidelberg (2008)
32. Okamoto, T., Takashima, K.: Hierarchical Predicate Encryption for Inner-Products. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 214–231. Springer, Heidelberg (2009)
33. Okamoto, T., Takashima, K.: Fully Secure Functional Encryption with General Relations from the Decisional Linear Assumption. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 191–208. Springer, Heidelberg (2010)
34. Ostrovksy, R., Sahai, A., Waters, B.: Attribute based encryption with non-monotonic access structures. In: ACM Conference on Computer and Communications Security, pp. 195–203 (2007)
35. Sahai, A., Waters, B.: Fuzzy Identity-Based Encryption. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473. Springer, Heidelberg (2005)
36. Shamir, A.: Identity-Based Cryptosystems and Signature Schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
37. Waters, B.: Efficient Identity-Based Encryption Without Random Oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)
38. Waters, B.: Dual System Encryption: Realizing Fully Secure IBE and HIBE under Simple Assumptions. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg (2009)
39. Waters, B.: Ciphertext-Policy Attribute-Based Encryption: An Expressive, Efficient, and Provably Secure Realization. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 53–70. Springer, Heidelberg (2011)