A Profitable Sub-prime Loan: Obtaining the Advantages of Composite Order in Prime-Order Bilinear Groups

Allison Lewko[1(✉)] and Sarah Meiklejohn[2]

[1] Columbia University, New York, USA
alewko@cs.columbia.edu
[2] University College London, London, UK
s.meiklejohn@ucl.ac.uk

Abstract. Composite-order bilinear groups provide many structural features that are useful for both constructing cryptographic primitives and enabling security reductions. Despite these convenient features, however, composite-order bilinear groups are less desirable than prime-order bilinear groups for reasons of both efficiency and security. A recent line of work has therefore focused on translating these structural features from the composite-order to the prime-order setting; much of this work focused on two such features, projecting and canceling, in isolation, but a result due to Seo and Cheon showed that both features can be obtained simultaneously in the prime-order setting.

In this paper, we reinterpret the construction of Seo and Cheon in the context of dual pairing vector spaces (which provide canceling as well as useful parameter hiding features) to obtain a unified framework that simulates all of these composite-order features in the prime-order setting. We demonstrate the strength of this framework by providing two applications: one that adds dual pairing vector spaces to the existing projection in the Boneh-Goh-Nissim encryption scheme to obtain leakage resilience, and another that adds the concept of projecting to the existing dual pairing vector spaces in an IND-CPA-secure IBE scheme to "boost" its security to IND-CCA1. Our leakage-resilient BGN application is of independent interest, and it is not clear how to achieve it from pure composite-order techniques without mixing in additional vector space tools. Both applications rely solely on the Symmetric External Diffie Hellman assumption (SXDH).

1 Introduction

Since their introduction in 2005 by Boneh, Goh, and Nissim [9], composite-order bilinear groups have been used to construct a diverse set of advanced cryptographic primitives, including (hierarchical) identity-based encryption [30, 32], group signatures [12,13], functional encryption [26,29], and attribute-based encryption [31]. The main assumptions used to prove the security of such schemes are variants of the *subgroup decision* assumption, which (in the simplest case)

J. Katz (Ed.): PKC 2015, LNCS 9020, pp. 377–398, 2015.
DOI: 10.1007/978-3-662-46447-2_17

states that, for a bilinear group G of order $N = pq$, without an element of order q it should be hard to distinguish a random element of G from a random element of order p. Such assumptions crucially rely on the hardness of factoring N.

Beyond this basic assumption and its close variants, many of these schemes have exploited additional structural properties that are inherent in composite-order bilinear groups. Two such properties, *projecting* and *canceling*, were formally identified by Freeman [18]; projecting requires (roughly) that there exists a trapdoor projection map from G into its p-order subgroup (and a related map in the target group G_T), and canceling requires that elements in the p-order and q-order subgroups cancel each other out (i.e., yield the identity when paired). Additionally, Lewko [27] identified another property, *parameter hiding*, that requires (again, roughly) that elements in the p-order subgroup reveal nothing about seemingly correlated elements in the q-order subgroup.

While therefore quite attractive and rich from a structural standpoint, the use of composite-order bilinear groups comes with a number of drawbacks, both in terms of efficiency and security. Until a recent construction of Boneh, Rubin, and Silverberg [11], all known composite-order bilinear groups were on supersingular, or Type-1 [19], curves. Even in the prime-order setting, supersingular curves are already less efficient than their ordinary counterparts: speed records for the former [4,42] are approximately six times slower than speed records for the latter [5]. In the composite-order setting, it is furthermore necessary to increase the size of the modulus by at least a factor of 10 (from 160 to at least 1024 bits) in order to make the assumption that N is hard to factor plausible. Operations performed in composite-order bilinear groups are therefore significantly slower; for example, Guillevic [22] recently observed that computing a pairing was 254 times slower. (This slowdown also extends to the non-supersingular construction of Boneh et al., and indeed to any composite-order bilinear group.) Furthermore, from a security standpoint, a number of recent results [1,2,21,23,25] demonstrate that it is possible to efficiently compute discrete logarithms in common types of supersingular curves, so that one must be significantly more careful when working over supersingular curves than when working over their non-supersingular counterparts.

One natural question to ask is: to what extent is it possible to obtain the structural advantages of composite-order bilinear groups without the disadvantages? Although the structural properties described above might seem specific to composite-order groups, both Freeman and Lewko are in fact able to express them rather abstractly and then describe how to construct prime-order bilinear groups in which each of these individual properties are met; they also show how to translate the subgroup decision assumption into a generalized version, that in prime-order groups is implied by either Decision Linear [8] or Symmetric External Diffie Hellman (SXDH) [6]. Lewko's approach is based on the framework of dual pairing vector spaces, as developed by Okamoto and Takashima [38,39]. This framework has been particularly useful for enabling translations of cryptosystems employing the dual system encryption methodology in their security reductions.

In contrast, Meiklejohn, Shacham, and Freeman [36] showed that it was impossible to achieve projecting and canceling simultaneously under a "natural" usage

of Decision Linear; as a motivation, they presented a blind signature scheme that seemingly relied upon both projecting and canceling for its proof of security. Recently, Seo and Cheon [44] showed that it was actually possible to achieve both projecting and canceling simultaneously in prime-order groups, and Seo [43] explored both possibility and impossibility results for projecting. To derive hardness of subgroup decision in their setting, however, Seo and Cheon rely on a non-standard assumption and show that this implies the hardness of subgroup decision only in a very limited case. They also provide a prime-order version of the Meiklejohn et al. blind signature that is somewhat divorced from their setting: rather than prove its security directly using projecting and canceling, they instead alter the blind signature, introduce a new property called *translating*, and then show that the modified blind signature is secure not in the projecting and canceling setting, but rather in a separate projecting and translating setting.

Subsequently, Herold et al. [24] presented a new translation framework called "polynomial spaces" that achieves projecting in a natural and elegant way, and can also be augmented to simultaneously achieve canceling. Like the prior result of Seo and Cheon, they employ a non-standard hardness assumption to obtain subgroup decision hardness when projecting and canceling are both supported. Interestingly, their approach does not seem to provide a way of achieving just canceling with subgroup decision problems relying on standard assumptions like SXDH or DLIN, as is achieved by dual pairing vector spaces. Integrating the benefits of dual pairing vector spaces into something like the polynomial spaces approach remains a worthwhile goal for future work. The framework in [24] also extends to the setting of multilinear groups, as do approaches based on eigenspaces, as demonstrated for example in [20].

Our Contributions. In this paper, we present in Section 3 an abstract presentation of the projecting and canceling pairing due to Seo and Cheon [44]. Our presentation is based on dual pairing vector spaces (DPVS) [38,39], and it can be parameterized to yield projection properties of varying strength. This perspective yields several advantages. First, all the power of DPVS is embedded inside this construction and can thus be exploited as in prior works. Second, we observe that many instances of subgroup decision problems in this framework are implied by the relatively simple SXDH assumption.

The advantages of our perspective are most clear for our BGN application, which we present in Section 4. If one starts with the goal of making the composite-order BGN scheme leakage resilient (i.e., providing provable security even when some bits of the secret key may have been leaked), the first obstacle one faces is the uniqueness of secret keys. Since the secret key is a factorization of the group order, there is only one secret key for each public key, making the common kind of hash proof argument for leakage resilience (as codified by Naor and Segev [37], for example) inapplicable. The DPVS techniques baked into our projecting and canceling prime-order construction remove this barrier quite naturally by allowing secret keys to be vectors that still serve as projection maps but can now be sampled from subspaces containing exponentially many potential

keys. This demonstrates the benefits of adding canceling and parameter hiding to applications that are designed around projection.

As an additional application, in Section 5, we present an IND-CCA1-secure identity-based encryption (IBE) scheme that uses canceling, parameter hiding, and weak projecting properties in its proof of security. Although efficient constructions of IND-CCA2-secure IBE schemes have been previously obtained by combining IND-CPA-secure HIBE schemes with signatures [15], we nevertheless view our IBE construction as a demonstration of the applicability of our unified framework. Furthermore, our new construction does not aim to amplify security by adding new primitives; instead, it explores the existing security of the IND-CPA-secure IBE due to Boneh and Boyen [7] (which cannot be IND-CCA2 secure, as it has re-randomizable ciphertexts), and observes that, by modifying the scheme in a rather organic way and exploiting the (weak) projecting and canceling properties of the setting, we can prove IND-CCA1 security directly. Hence, we view this as an exploration of the security properties that can be proved solely from the minimalistic spirit of the Boneh-Boyen scheme.

Our two applications serve as a proof of concept for the usefulness of obtaining projecting and canceling simultaneously in the prime-order setting, and a demonstration of how to leverage such properties while relying only on relatively simple assumptions like SXDH. We believe that the usefulness of our framework extends beyond these specific examples, and we intend our work to facilitate future applications of these combined properties.

Our Techniques. To obtain a more user-friendly interpretation of the projecting and canceling pairing construction over prime-order groups, we begin by observing that it is essentially a concatenation of DPVS. Dual pairing vector spaces were first used in prime-order bilinear groups by Okamoto and Takashima [38,39] and have since been employed in many works, in particular to instantiate dual system technique [45] in the prime-order setting [27,29,40]. These previous uses of DPVS typically relied on the canceling property, variants of subgroup decision problems, and certain parameter hiding properties that are present by design in DPVS. One particularly nice feature of DPVS constructions is that a large family of useful subgroup decision variants can be proven to follow from standard assumptions like SXDH for asymmetric groups and DLIN for symmetric groups; viewing the construction of a projecting and canceling pairing as a natural extension of DPVS therefore has the twin benefits that it provides a clear guide on how to derive certain subgroup decision variants from standard assumptions, and that it comes with all the built-in tools that DPVS offers.

In particular, DPVS includes a suite of vector-space-based tools for proving leakage resilience, similar to ones used in previous works [14,16,17,34,35,37]. This enables us to combine the projecting-supported limited homomorphic functionality of the BGN encryption scheme with provable leakage resilience. DPVS also supports a toolkit developed for dual system proofs (e.g., [29,40,41]), which is what enables us to boost our IBE to full IND-CCA1 security with just the addition of projection.

2 Definitions and Notation

In this section, we define bilinear groups and the three functional properties we would like them to satisfy: projecting, canceling, and parameter hiding. For the first two, we use the definitions of Freeman [18] (albeit in a somewhat modified form); for parameter hiding, on the other hand, we come up with a new formal framework. In addition to these functional properties, we consider the notion of subgroup decision in bilinear groups, in which a random element of a subgroup should be indistinguishable from a random element of the full group. The variant we define, called generalized correlated subgroup decision, is very general: in addition to seeing random elements of subgroups, we allow an attacker to see elements *correlated* across subgroups (e.g., elements of different subgroups with correlated randomness), and require that it is still difficult for him to distinguish between correlated elements of different subgroups. We then see in Section 3that many specific instances of this general notion are implied by more standard notions of subgroup decision in prime-order groups.

2.1 Bilinear Groups

In what follows, we refer to a *bilinear group* as a tuple $\mathcal{G} = (N, G, H, G_T, e, \mu)$, where N is either prime or composite, $|G| = |H| = kN$ and $|G_T| = \ell N$ for some $k, \ell \in \mathbb{N}$, and $e : G \times H \to G_T$ is a bilinear map; i.e., e is an efficient map that satisfies both *bilinearity* ($e(x^a, y^b) = e(x, y)^{ab}$ for all $x \in G$, $y \in H$, $a, b \in \mathbb{Z}/N\mathbb{Z}$) and *non-degeneracy* (if $e(x, y) = 1$ for all $x \in G$ then $y = 1$ and if $e(x, y) = 1$ for all $y \in H$ then $x = 1$). In some bilinear groups, we may additionally include generators g and h of G and H respectively (if G and H are cyclic), information about meaningful subgroups of G and H, or some auxiliary information μ that allows for efficient membership testing in G and H (and possibly more). In what follows, we refer to the algorithm that is used to generate such a $\mathcal{G}$ as $\mathsf{BilinearGen}$. Beyond the security parameter, $\mathsf{BilinearGen}$ takes in an additional parameter n that specifies the number of desired subgroups; i.e., for $(N, G, H, G_T, e, \mu) \xleftarrow{\$} \mathsf{BilinearGen}(1^k, n)$, we have $G = \oplus_{i=1}^{n} G_i$ and $H = \oplus_{i=1}^{n} H_i$ (where typically G_i and H_i are cyclic).

In terms of functional properties of bilinear groups, we first define both *projecting* and *canceling*; our definitions are modified versions of the ones originally given by Freeman [18]. We give three flavors of projecting. The first, *weak projecting*, considers projecting into a single subgroup of the source group, without requiring a corresponding map in the target group. The second, which we call simply *projecting*, most closely matches the definition given by Freeman, and considers projecting into a single subgroup in both the source and target groups. Lastly, we define *full* projecting, which considers projecting into every subgroup individually. As we will see in Section 3, we can satisfy all of these flavors by tweaking appropriate parameters in our prime-order construction.

Definition 2.1 (Weak Projecting). *A bilinear group $\mathcal{G} = (N, G, H, G_T, e, \mu)$ is* weakly projecting *if there exist decompositions $G = G_1 \oplus G_2$ and $H = H_1 \oplus H_2$,*

and projection maps π_G *and* π_H *such that* $\pi_G(x_1) = x_1$ *for all* $x_1 \in G_1$ *and* $\pi_G(x_2) = 1$ *for all* $x_2 \in G_2$*, and similarly* $\pi_H(y_1) = y_1$ *for all* $y_1 \in H_1$ *and* $\pi_H(y_2) = 1$ *for all* $y_2 \in H_2$*.*

Definition 2.2 (Projecting). *A bilinear group* $\mathcal{G} = (N, G, H, G_T, e, \mu)$ *is* projecting *if there exist subgroups* $G' \subset G$*,* $H' \subset H$*, and* $G'_T \subset G_T$ *such that there exist non-trivial maps* $\pi_G : G \to G'$*,* $\pi_H : H \to H'$*, and* $\pi_T : G_T \to G'_T$ *such that* $\pi_T(e(x, y)) = e(\pi_G(x), \pi_H(y))$ *for all* $x \in G$*,* $y \in H$*.*

Definition 2.3 (Full projecting). *A bilinear group* $\mathcal{G} = (N, G, H, G_T, e, \mu)$ *is* fully projecting *if there exists some* $n \in \mathbb{N}$ *and decompositions* $G = \oplus_{i=1}^n G_i$*,* $H = \oplus_{i=1}^n H_i$*, and* $G_T = \oplus_{i=1}^n G_{T,i}$*, and non-trivial maps* $\pi_{Gi} : G \to G_i$*,* $\pi_{Hi} : H \to H_i$*, and* $\pi_{Ti} : G_T \to G_{T,i}$ *for all* i *such that* $\pi_{Ti}(e(x, y)) = e(\pi_{Gi}(x), \pi_{Hi}(y))$ *for all* $x \in G$*,* $y \in H$*.*

Definition 2.4 (Canceling). *A bilinear group* $\mathcal{G} = (N, G, H, G_T, e, \mu)$ *is* canceling *if there exists some* $n \in \mathbb{N}$ *and decompositions* $G = \oplus_{i=1}^n G_i$ *and* $H = \oplus_{i=1}^n H_i$ *such that* $e(x_i, y_j) = 1$ *for all* $x_i \in G_i$*,* $y_j \in H_j$*,* $i \neq j$*.*

2.2 Parameter Hiding

Beyond projecting and canceling, we aim to define *parameter hiding*. As mentioned in the introduction, this property roughly says that elements in one subgroup should not reveal anything about related elements in other subgroups, and was previously used, without a formal definition, by Lewko [27]. In essence, parameter hiding in composite-order groups is a simple consequence of the Chinese Remainder Theorem, which tells us that if we sample a random value modulo $N = pq$, its reductions modulo p and q are uncorrelated. In the prime-order setting, a form of parameter hiding can be instantiated from dual pairing vector spaces, leveraging the fact that if one commits to only certain parts of dual orthonormal bases over $\mathbb{F}_p^n$, there is remaining ambiguity in the hidden basis vectors.

The main difficulty in providing a formal definition for parameter hiding is that it is not as self-contained a feature as projecting and canceling: elements within subgroups may be related to elements in other subgroups in a myriad of ways, and their relation to one another may depend both on the form of the element (which can involve any function on the exponents) and on the subgroups. We therefore do not try to consider all types of correlations, but instead focus on one simple type, defined as follows:

Definition 2.5. *For a bilinear group* $\mathcal{G} = (N, G = \oplus_{i=1}^n G_i, H = \oplus_{i=1}^n H_i, G_T, e, \{g_i\}_{i=1}^n, \{h_i\}_{i=1}^n)$*, an element* $x \in \mathbb{Z}/N\mathbb{Z}$*, and indices* $1 \leq i_1, i_2 \leq n$*, an* x-correlated sample *from the subgroup* $G_{i_1} \oplus G_{i_2}$ *is an element of the form* $g_{i_1}^{\alpha} \cdot g_{i_2}^{\alpha x}$ *for* $\alpha \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$*.*

We also consider correlated samples in H, but for convenience we define a y-correlated sample from the subgroup $H_{i_1} \oplus H_{i_2}$ to be an element of the form

$h_{i_1}^{\beta y} \cdot h_{i_2}^{\beta}$ for $\beta \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$. Although we choose this type of correlation mainly for ease of exposition (and because we encounter it in Section 5), our discussion below could be adjusted to accommodate more general types of correlation, which would remain compatible with our prime-order construction in Section 3.

Intuitively then, parameter hiding says that, under certain restrictions about which subgroup elements one is allowed access to, the distributions over x-correlated samples and random samples should in fact be the same, even when x is known. (We need some restrictions because there may be testable relationships between the images of various generators in the target group.) To consider the distributions we can use — i.e., what additional information we might give out besides the samples — we consider distributions $\mathcal{D}$ parameterized by sets $S_G^{\mathsf{ph}} = \{S_{G,\mathrm{gen}}^{\mathsf{ph}}, S_{G,\mathrm{sam}}^{\mathsf{ph}}, S_{G,\mathrm{cor}}^{\mathsf{ph}}\}$, $S_H^{\mathsf{ph}} = \{S_{H,\mathrm{gen}}^{\mathsf{ph}}, S_{H,\mathrm{sam}}^{\mathsf{ph}}, S_{H,\mathrm{cor}}^{\mathsf{ph}}\}$, and C; intuitively, S_G^{ph} and S_H^{ph} tell us which elements to include in the distribution, and C tells us which correlated samples to change to random. Formally, these sets are defined as follows:

- $S_{G,\mathrm{gen}}^{\mathsf{ph}}$ indicates which subgroup generators to include: For all $s_i \in S_{G,\mathrm{gen}}^{\mathsf{ph}}$, include g_{s_i} in $\mathcal{D}$.
- $S_{G,\mathrm{sam}}^{\mathsf{ph}}$ is a multiset that indicates which random samples to include: For all $t_i = (t_{1,i}, \ldots, t_{m_i,i}) \in S_{G,\mathrm{sam}}^{\mathsf{ph}}$, include a random sample from $G_{t_{1,i}} \oplus \ldots \oplus G_{t_{m_i,i}}$ in $\mathcal{D}$.
- $S_{G,\mathrm{cor}}^{\mathsf{ph}}$ is a set that indicates which correlated samples to include: For all $c_i = (x_i, c_{1,i}, c_{2,i}) \in S_{G,\mathrm{cor}}^{\mathsf{ph}}$, include $g_{c_{1,i}}^{a} \cdot g_{c_{2,i}}^{a x_i}$ in $\mathcal{D}$, where $a \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$.
- S_H^{ph} is defined analogously to S_G^{ph}.
- C indicates which correlated samples to change: For all $c_i = (b_i, c_i') \in C$, if $b_i = 0$ then $c_i' \in S_{G,\mathrm{cor}}^{\mathsf{ph}}$ and if $b_i = 1$ then $c_i' \in S_{H,\mathrm{cor}}^{\mathsf{ph}}$; i.e., we require that $C \subseteq \{0 \times S_{G,\mathrm{cor}}^{\mathsf{ph}}\} \cup \{1 \times S_{H,\mathrm{cor}}^{\mathsf{ph}}\}$.

Given all these sets, we now require that they are *well-behaved* in the following two ways: (1) for any changed x-correlated sample, do not reveal the corresponding subgroup generators on either side of the pairing, and (2) do not change correlated samples for the same value x in the same subgroups on opposite sides of the pairing. Formally, we express these requirements as

- Don't include generators for switched samples: For all $(b_i, (x_i, c_{1,i}, c_{2,i})) \in C$, $s_j \in S_{G,\mathrm{gen}}^{\mathsf{ph}}$, and $s_\ell \in S_{H,\mathrm{gen}}^{\mathsf{ph}}$, $s_j \neq c_{1,i}, c_{2,i}$ and $s_\ell \neq c_{1,i}, c_{2,i}$.
- Don't switch x-correlated samples in overlapping subgroups of G and H: For all $(0, (x_i, c_{1,i}, c_{2,i})), (1, (x_j, c_{1,j}, c_{2,j})) \in C$, either $x_i \neq x_j$ or $c_{1,i} \neq c_{1,j}, c_{2,j}$ and $c_{2,i} \neq c_{1,j}, c_{2,j}$.

To see why these restrictions can be necessary, consider trying to establish that an x-correlated sample in $G_1 \oplus G_2$ is identical to a random sample in $G_1 \oplus G_2$, and suppose we are given h_1 and h_2. If we are given $g_1^{\alpha} g_2^{\alpha x}$ (for some random, unknown α), then — assuming we are using a canceling pairing — we can compute $e(g_1, h_1)^{\alpha}$ and $e(g_2, h_2)^{\alpha x}$. When working with specific instantiations,

there may be a known relationship between $e(g_1, h_1)$ and $e(g_2, h_2)$. (In fact, for our IBE construction, $e(g_1, h_1) = e(g_2, h_2)^{-1}$.) In this case, if x is known then we can test for an x-correlation in the target group, and hence distinguish an x-correlated sample from a random one. Similarly, if we have x-correlated samples $g_1^{\alpha} g_2^{\alpha x}$ and $h_1^{\beta x} h_2^{\beta}$, then pairing these yields the identity, which distinguishes them from random.

Definition 2.6 (Parameter Hiding). *We say that a group* $\mathcal{G} = (N, G, H, G_T, e, \mu)$ *satisfies* parameter hiding *with respect to a well-behaved distribution* $\mathcal{D} = (S_G^{\mathsf{ph}}, S_H^{\mathsf{ph}}, C)$ *if* $\mathcal{D}$ *is identical to the distribution in which the correlated samples indicated by* C *are replaced with random samples.*

Example 1. As an example, consider the distribution $\mathcal{D}$ defined by $S_G^{\mathsf{ph}} = \{\{1,2\}, \emptyset, \{(x,1,2),(x,3,4)\}\}$, $S_H^{\mathsf{ph}} = \{\{1,2,5,6\}, \{(3,4),(3,4)\}, \{(y,1,2),(y,3,4)\}\}$, and $C = \{(0,(x,3,4)),(1,(y,3,4)\}$ for any $x, y \in \mathbb{Z}/N\mathbb{Z}$ such that $x \neq y$; we can easily check that these sets are well-behaved in the sense defined above. Then parameter hiding holds for $\mathcal{G} = (N, G, H, G_T, e, \mu)$ if for $a, b, c, d, s, t, u, v, w, z \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$,

$$(N, G, H, G_T, e, \mu, g_1, g_2, h_1, h_2, h_5, h_6, h_3^a h_4^b, h_3^c h_4^d, h_1^{ty} h_2^t, h_3^{zy} h_4^z, g_1^s g_2^{sx}, g_3^w g_4^{wx})$$

is *identical* to

$$(N, G, H, G_T, e, \mu, g_1, g_2, h_1, h_2, h_5, h_6, h_3^a h_4^b, h_3^c h_4^d, h_1^{ty} h_2^t, h_3^v h_4^z, g_1^s g_2^{sx}, g_3^w g_4^u).$$

In our uses of parameter hiding in Section 5, we restrict ourselves to this one example. Again, this is due to the difficulty of providing a fully general definition of parameter hiding, as certain types of correlated samples require more entropy than others. We nevertheless do not find it to be overly limiting to consider this one example, as it keeps our constructions in Section 5 simple and tailored to the requirements that we need. We also use a variant of parameter hiding in the proof for our leakage-resilient BGN variant presented in Section 4. Here, the flexibility in the hidden parameters is leveraged to allow the simulator to a leak on a secret key before fully committing to a complete basis (i.e., before determining how to form an appropriate ciphertext).

2.3 Generalized Correlated Subgroup Decision

Beyond functional properties of bilinear groups, we must also consider the types of security guarantees we can provide. The assumption we define, generalized correlated subgroup decision, considers indistinguishability between subgroups in a very general way: given certain subgroup generators and "correlated" elements across subgroups (i.e., elements in different subgroups that use the same randomness), it should still be hard to distinguish between elements of other subgroups. Formally, we consider sets $S_G^{\mathsf{sgh}} = \{S_{G,\mathrm{gen}}^{\mathsf{sgh}}, S_{G,\mathrm{sam}}^{\mathsf{sgh}}\}$, $S_H^{\mathsf{sgh}} = \{S_{H,\mathrm{gen}}^{\mathsf{sgh}}, S_{H,\mathrm{sam}}^{\mathsf{sgh}}\}$, $T_1 = \{(\ell_1, \lambda_1), \ldots, (\ell_m, \lambda_m)\}$, and $T_2 = \{(\ell'_1, \lambda'_1), \ldots, (\ell'_{m+1}, \lambda'_{m+1})\}$, and an

indicator bit b. (We assume without loss of generality that T_2 is the larger set.) Intuitively, S_G^{sgh} and S_H^{sgh} tell us which group elements an adversary is given, and (T_1, T_2, b) tell us what the challenge terms should look like. We have the following requirements:

- $S_{G,\mathrm{gen}}^{\mathsf{sgh}}$ indicates which subgroup generators to include: Give out g_{s_i} for all $s_i \in S_{G,\mathrm{gen}}^{\mathsf{sgh}}$.
- $S_{G,\mathrm{sam}}^{\mathsf{sgh}}$ indicates which samples to include: For each

$$t_i = ((\ell_{1,i}, \lambda_{1,i}), \ldots, (\ell_{m_i,i}, \lambda_{m_i,i})) \in S_{G,\mathrm{sam}}^{\mathsf{sgh}}$$

  , give out $g_{\ell_{1,i}}^{a_1} \cdot \ldots \cdot g_{\ell_{m_i,i}}^{a_{m_i}}$ and $g_{\lambda_{1,i}}^{a_1} \cdot \ldots \cdot g_{\lambda_{m_i,i}}^{a_{m_i}}$ for $a_1, \ldots, a_{m_i} \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$. These elements are *correlated*, in that the same randomness is used for both.
- The bit b indicates which group the challenge element comes from: $b = 0$ indicates G, and $b = 1$ indicates H.
- The sets T_1 and T_2 must differ in exactly one pair; i.e., there must exist a unique pair P such that $P \notin T_1$ but $P \in T_2$. For this pair $P = (\ell, \lambda)$, we cannot give out the subgroup generators on either side of the pairing, so we require $s_i \neq \ell$ and $s_i \neq \lambda$ for any $s_i \in S_{G,\mathrm{gen}}^{\mathsf{sgh}}$ or $s_i \in S_{H,\mathrm{gen}}^{\mathsf{sgh}}$.

  If $P \in t_i$ for some $t_i \in S_{G,\mathrm{sam}}^{\mathsf{sgh}} \cup S_{H,\mathrm{sam}}^{\mathsf{sgh}}$, then $T_1 \cap t_i \neq \emptyset$; i.e., P can appear only in random samples that also contain another component in the challenge term. Then, assuming $b = 0$ (and replacing g with h if $b = 1$), our challenge elements are of the form $T := (g_{\ell_1}^{a_1} \cdot \ldots \cdot g_{\ell_m}^{a_m}, g_{\lambda_1}^{a_1} \cdot \ldots \cdot g_{\lambda_m}^{a_m})$ and $T' := (g_{\ell_{1'}}^{a_1} \cdot \ldots \cdot g_{\ell'_{m+1}}^{a_{m+1}}, g_{\lambda'_1}^{a_1} \cdot \ldots \cdot g_{\lambda'_{m+1}}^{a_{m+1}})$ for $a_1, \ldots a_{m+1} \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$.

Assumption 2.1 (Generalized Correlated Subgroup Decision). *For all tuples $(S_G^{\mathsf{sgh}}, S_H^{\mathsf{sgh}}, T_1, T_2, b)$ satisfying the requirements specified above and for any $n \in \mathbb{N}$, for any PPT adversary $\mathcal{A}$ given $\mathcal{G} \xleftarrow{\$} \mathsf{BilinearGen}(1^k, n)$ and the elements specified by S_G^{sgh} and S_H^{sgh}, it is hard to distinguish between values T defined by (b, T_1) and values T' defined by (b, T_2).*

As an example, consider the case in which $n = 6$ and $S_G^{\mathsf{sgh}} = \{\{1, 2\}, \{((1, 2), (3, 4))\}\}$, $S_H^{\mathsf{sgh}} = \{\{1, 2, 5, 6\}, \{((1, 2), (3, 4)), ((3, 4), (5, 6))\}\}$, $T_1 = \{(1, 2), (5, 6)\}$, $T_2 = \{(1, 2), (3, 4), (5, 6)\}$, and $b = 0$. In this case, the concrete assumption is: Given $\mathcal{G}$ and generators $g_1, g_2, h_1, h_2, h_5, h_6$, correlated samples from $G_1 \oplus G_3$ and $G_2 \oplus G_4$, correlated samples from $H_1 \oplus H_3$ and $H_2 \oplus H_4$, and correlated samples from $H_3 \oplus H_5$ and $H_4 \oplus H_6$, it should be hard to distinguish correlated samples from $G_1 \oplus G_5$ and $G_2 \oplus G_6$ from correlated samples from $G_1 \oplus G_3 \oplus G_5$ and $G_2 \oplus G_4 \oplus G_6$.

3 A Prime-Order Bilinear Group Satisfying All Features

Our ultimate goal in this section is to define a prime-order bilinear group that satisfies all three of the properties defined in the previous section: projecting,

canceling, and parameter hiding; additionally, we want to require that subgroup decision is hard in this group. Our construction can be viewed as an abstraction of the construction of Seo and Cheon [44], which they prove satisfies (regular) projecting, canceling, and a somewhat restrictive notion of subgroup decision. In contrast, our construction satisfies canceling and parameter hiding, is flexible enough to achieve any of the three flavors of projecting we defined in the previous section (depending on the parameter choices), and comes equipped with reductions for more general instances of subgroup decision.

Notationally, we augment the bilinear groups $\mathcal{G}$ discussed in the previous section: we now focus only on the case when the group order is some prime p, and consider $\mathbb{G} = (p, B_1, B_2, B_T, E, \mu)$ built on top of $\mathcal{G} = (p, G, H, G_T, e)$; this means B_1, B_2, and B_T may contain multiple copies of G, H, and G_T respectively, and that the map E uses e as a component. Because we are moving to bigger spaces, we also include a value μ that allows us to test membership in B_1 and B_2; as an example, consider $B_1 \subset G \times G$. Then, while an efficient membership test for G implies one for $G \times G$, additional information μ may be necessary to allow one to (efficiently) test for membership in B_1.

Our construction crucially uses dual pairing vector spaces, which were introduced by Okamoto and Takashima [38,39] and have been previously used to provide pairings $E : G^n \times H^n \to G_T$, built on top of pairings $e : G \times H \to G_T$, that satisfy the canceling property. As we cannot have a cyclic target space if we want to satisfy projecting, however, we instead need a map whose image is G_T^d for some $d > 1$. Intuitively, we achieve this by piecing together d "blocks," where each block is an instance of a dual pairing vector space; the construction of Seo and Cheon is then obtained as the special case in which $d = n$, and regular dual pairing vector spaces are obtained with $d = 1$. We begin with a key definition:

Definition 3.1 (Dual Orthonormal). *Two bases* $\mathbb{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ *and* $\mathbb{B}^* = (\boldsymbol{b}_1^*, \ldots, \boldsymbol{b}_n^*)$ *of* $\mathbb{F}_p^n$ *are* dual orthonormal *if* $\boldsymbol{b}_j \cdot \boldsymbol{b}_j^* \equiv 1 \bmod p$ *for all* j, $1 \leq j \leq n$, *and* $\boldsymbol{b}_j \cdot \boldsymbol{b}_k^* \equiv 0 \bmod p$ *for all* $j \neq k$.

We note that one can efficiently sample a random pair of dual orthonormal bases $(\mathbb{B}, \mathbb{B}^*)$ by sampling first a random basis $\mathbb{B}$ and then solving uniquely for $\mathbb{B}^*$ using linear algebra over $\mathbb{F}_p$; we denote this sampling process as $(\mathbb{B}, \mathbb{B}^*) \xleftarrow{\$} \mathit{Dual}(\mathbb{F}_p^n)$. By repeating this sampling process d times, we can obtain a tuple $((\mathbb{B}_1, \mathbb{B}_1^*), \ldots, (\mathbb{B}_d, \mathbb{B}_d^*))$ of d pairs of dual orthonormal bases of $\mathbb{F}_p^n$. We denote the vectors of $\mathbb{B}_i$ as $(\boldsymbol{b}_{1,i} \ldots, \boldsymbol{b}_{n,i})$, and the vectors of $\mathbb{B}_i^*$ as $(\boldsymbol{b}_{1,i}^*, \ldots, \boldsymbol{b}_{n,i}^*)$. We then give the following definition:

Definition 3.2 (Concatenation). *The* concatenation *of bases* $(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ *of* $\mathbb{F}_p^n$ *is a collection of* n *vectors* $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n)$ *in* $\mathbb{F}_p^{dn}$, *where each* $\boldsymbol{v}_j := \boldsymbol{b}_{j,1} || \ldots || \boldsymbol{b}_{j,d}$. *Alternatively, we can view each* $\boldsymbol{v}_j$ *as a* $d \times n$ *matrix, where the* i*-th row is* $\boldsymbol{b}_{j,i}$. *We denote the concatenation of* $(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ *as* $\mathsf{Concat}(\mathbb{B}_1, \ldots, \mathbb{B}_d)$.

To begin our construction, we build off $\mathcal{G} = (p, G, H, G_T, e, g, h)$, where g and h are generators of G and H respectively, and consider groups $B_1 \subset G^{dn}$ and $B_2 \subset H^{dn}$. Notationally, we write an element of B_1 as g^A, where $A =$

$(\alpha_{i,j})_{i,j=1}^{d,n}$ is a $d \times n$ matrix and $g^A := (g^{\alpha_{1,1}}, \ldots, g^{\alpha_{1,j}}, \ldots, g^{\alpha_{1,n}}, g^{\alpha_{2,1}}, \ldots, g^{\alpha_{d,n}})$. We similarly write elements of B_2 as h^B for a $d \times n$ matrix $B = (\beta_{ij})_{i,j=1}^{d,n}$, and furthermore define the bilinear map $E : B_1 \times B_2 \to G_T^d$ as

$$E(g^A, h^B) := \left(\prod_{k=1}^{n} e(g^{\alpha_{1,k}}, h^{\beta_{1,k}}), \ldots, \prod_{k=1}^{n} e(g^{\alpha_{d,k}}, h^{\beta_{d,k}}) \right). \tag{1}$$

Observe that the i-th coordinate of the image is equal to $e(g,h)^{A_i \cdot B_i \bmod p}$, where A_i and B_i denote the i-th rows of A and B respectively. Then, to begin to see how our construction will satisfy projecting and canceling, we have the following lemma:

Lemma 3.1. *Let* $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n) = \mathsf{Concat}(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ *and* $(\boldsymbol{v}_1^*, \ldots, \boldsymbol{v}_n^*) = \mathsf{Concat}(\mathbb{B}_1^*, \ldots, \mathbb{B}_d^*)$*, where* $(\mathbb{B}_i, \mathbb{B}_i^*)$ *are dual orthonormal bases of* $\mathbb{F}_p^n$*. Then*

$$E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_j^*}) = (e(g,h), \ldots, e(g,h)) \; \forall j \quad \textit{and} \quad E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_k^*}) = (1_T, \ldots, 1_T) \; \forall j \neq k.$$

Proof. By definition of the pairing,

$$E(g^{\boldsymbol{v_j}}, h^{\boldsymbol{v_k^*}}) = \left(e(g,h)^{\boldsymbol{b}_{j,1} \cdot \boldsymbol{b}_{k,1}^*}, \ldots, e(g,h)^{\boldsymbol{b}_{j,d} \cdot \boldsymbol{b}_{k,d}^*} \right)$$

for any j and k. If $j = k$, then the fact that $(\mathbb{B}_i, \mathbb{B}_i^*)$ are dual orthonormal for all i implies by definition that $\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{j,i}^* \equiv 1 \bmod p$ for all i and j, and thus $E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_j^*}) = (e(g,h), \ldots, e(g,h))$. For the second property, we again use the definition of dual orthonormal bases to see that $\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{k,i}^* \equiv 0 \bmod p$ for all $j \neq k$, and thus $E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_k^*}) = (1_T, \ldots, 1_T)$. □

While Lemma 3.1 therefore shows us directly how to obtain canceling, for projecting we are still mapping into a one-dimensional image. To obtain more dimensions, it turns out we need only perform some additional scalar multiplication. We give the following definition:

Definition 3.3 (Scaling). *Define* $C = (c_{i,j})_{i,j=1}^{d,n}$ *to be a* $n \times d$ *matrix of entries over* $\mathbb{F}_p \setminus \{0\}$*. Given bases* $(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ *of* $\mathbb{F}_p^n$*, we define the* scaling *of these bases by* C *to be new bases* $(\mathbb{D}_1, \ldots, \mathbb{D}_d)$*, where* $\mathbb{D}_i = (c_{1,i}\boldsymbol{b}_{1,i}, \ldots, c_{n,i}\boldsymbol{b}_{n,i})$ *for all* i*,* $1 \leq i \leq d$*. We denote the scaling of* $(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ *by* C *as* $\mathsf{Scale}(C, \mathbb{B}_1, \ldots, \mathbb{B}_d)$*.*

Intuitively then, we use the entries in the i-th column of C to scale the vectors in the basis $\mathbb{B}_i$ and obtain the basis $\mathbb{D}_i$. As we still have $\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{k,i}^* \equiv 0 \bmod p$ for $j \neq k$, multiplication by a scalar will not affect this and we still satisfy canceling. The scalar values do, however, build in extra dimensions into the image of our pairing, as demonstrated by the following lemma:

Lemma 3.2. *Let* $(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ *and* $(\mathbb{B}_1^*, \ldots, \mathbb{B}_d^*)$ *be sets of bases for* $\mathbb{F}_p^n$ *such that* $(\mathbb{B}_i, \mathbb{B}_i^*)$ *are dual orthonormal for all* i*. Define* $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n) := \mathsf{Concat}(\mathbb{D}_1, \ldots, \mathbb{D}_d)$

and $(\boldsymbol{v}_1^*, \ldots, \boldsymbol{v}_n^*) := \mathsf{Concat}(\mathbb{B}_1^*, \ldots, \mathbb{B}_d^*)$*, where* $(\mathbb{D}_1, \ldots, \mathbb{D}_d) = \mathsf{Scale}(C, \mathbb{B}_1, \ldots, \mathbb{B}_d)$ *for some* $C \in M_{n \times d}(\mathbb{F}_p)$*. Then*

$$E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_j^*}) = (e(g,h)^{c_{j,1}}, \ldots, e(g,h)^{c_{j,d}}) \ \forall j \quad \textit{and}$$
$$E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_k^*}) = (1_T, \ldots, 1_T) \ \forall j \neq k.$$

Proof. y definition of the pairing,

$$E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_k^*}) = \left(e(g,h)^{c_{j,1}\boldsymbol{b}_{j,1} \cdot \boldsymbol{b}_{k,1}^*}, \ldots, e(g,h)^{c_{j,d}\boldsymbol{b}_{j,d} \cdot \boldsymbol{b}_{k,d}^*}\right)$$

for any j and k. If $j = k$, then the fact that $(\mathbb{B}_i, \mathbb{B}_i^*)$ are dual orthonormal for all i implies by definition that $\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{j,i}^* \equiv 1 \bmod p$ for all i and j, and thus $c_{j,i}\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{j,i}^* \equiv c_{j,i} \bmod p$ and $E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_j^*}) = (e(g,h)^{c_{j,1}}, \ldots, e(g,h)^{c_{j,d}})$. For the second property, we again use the definition of dual orthonormal bases to see that $\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{k,i}^* \equiv 0 \bmod p$ for all $j \neq k$, and thus $c_{j,i}\boldsymbol{b}_{j,i} \cdot \boldsymbol{b}_{k,i}^* \equiv 0 \bmod p$ and $E(g^{\boldsymbol{v}_j}, h^{\boldsymbol{v}_k^*}) = (1_T, \ldots, 1_T)$. □

We are now ready to give our full construction of an algorithm $\mathsf{BilinearGen}'$, parameterized by integers n and d, and a distribution $\mathcal{D}_{n,d}$ on $n \times d$ matrices, to achieve a setting $\mathbb{G} = (p, B_1, B_2, B_T, E, \mu)$ such that $B_1 \subset G^{dn}$, $B_2 \subset H^{dn}$, and $B_T = G_T^d$. We present this construction in Algorithm 1, and demonstrate that it satisfies projecting, canceling, parameter hiding, and subgroup decision.

The generality of this construction stems from the choices of d, n, and $\mathcal{D}$; in fact, by choosing different values for these parameters, we can satisfy each of the different flavors of projecting from Section 2. To satisfy fully projecting, we choose C from a distribution over matrices of full rank n and use $d \geq n$. If we use a less restrictive distribution, we obtain weaker projection capabilities and a more efficient construction (as we can have $d < n$) when projecting onto all subgroups individually is not needed: to achieve (regular) projecting, we can use $d > 1$ and pick C to be of rank > 1, and to achieve weak projecting we can in fact use $d = 1$ and pick C to be the vector consisting of all 1 entries. (This last case is equivalent to working in regular dual pairing vector spaces.)

Theorem 3.1. *For all values of* $n \geq 2$*, the bilinear group* $\mathbb{G} \xleftarrow{\$} \mathsf{BilinearGen}'(1^k, n, d, \mathcal{D}_{d,n})$ *satisfies canceling, fully projecting as defined in Definition 2.3 for* $d \geq n$ *when* $\mathcal{D}_{d,n}$ *is defined over full-rank matrices, projecting as defined in Definition 2.2 for* $d > 1$ *when* $\mathcal{D}_{d,n}$ *is defined over matrices of rank* > 1*, and weak projecting as defined in Definition 2.1 for* $d = 1$*.*

Proof. Given that our construction was specifically designed to satisfy the conditions for Lemma 3.2, we immediately obtain canceling. To satisfy projecting, we additionally need to construct the projection maps π_{ij} and argue that they satisfy the requirements of Definition 2.3 (in the case that C is full rank). By the way our subgroups are defined, each projection map π_{1i} within the group B_1 must map an arbitrary element $g^{a_1\boldsymbol{v}_1 + \cdots + a_n\boldsymbol{v}_n}$ of B_1 to $g^{a_i\boldsymbol{v}_i} \in B_{1,i}$; similarly, π_{2i} must map $h^{a_1^*\boldsymbol{v}_1^* + \cdots + a_n^*\boldsymbol{v}_n^*} \in B_2$ to $h^{a_i^*\boldsymbol{v}_i^*} \in B_{2,i}$. For π_{1i}, we observe that it

Algorithm 1. $\mathsf{BilinearGen}'$: generate a bilinear group $\mathbb{G}$ that satisfies projecting and canceling

Input: $d, n \in \mathbb{N}$; distribution $\mathcal{D}_{d,n}$ over matrices in $M_{n\times d}(\mathbb{F}_p)$; security parameter 1^k.

1. $(p, G, H, G_T, e) \xleftarrow{\$} \mathsf{BilinearGen}(1^k, 1)$.
2. Pick values g and h such that $G = \langle g\rangle$ and $H = \langle h\rangle$.
3. Sample d pairs $(\mathbb{B}_i, \mathbb{B}_i^*) \xleftarrow{\$} \mathit{Dual}(\mathbb{F}_p^n)$ to obtain two sets $(\mathbb{B}_1, \ldots, \mathbb{B}_d)$ and $(\mathbb{B}_1^*, \ldots, \mathbb{B}_d^*)$ of bases of $\mathbb{F}_p^n$, where $(\mathbb{B}_i, \mathbb{B}_i^*)$ are dual orthonormal.
4. Sample $C = (c_{ij})_{i,j=1}^{d,n} \xleftarrow{\$} \mathcal{D}$ and compute $(\mathbb{D}_1, \ldots, \mathbb{D}_d) := \mathsf{Scale}(C, \mathbb{B}_1, \ldots, \mathbb{B}_d)$.
5. For all i, $1 \le i \le n$, define $B_{1,i} := \langle g^{\boldsymbol{v}_i}\rangle$ and $B_{2,i} := \langle h^{\boldsymbol{v}_i^*}\rangle$, where $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n) := \mathsf{Concat}(\mathbb{D}_1, \ldots, \mathbb{D}_d)$ and $(\boldsymbol{v}_1^*, \ldots, \boldsymbol{v}_n^*) := \mathsf{Concat}(\mathbb{B}_1^*, \ldots, \mathbb{B}_d^*)$.
6. Define $B_1 := \oplus_{i=1}^n B_{1,i} \subset G^{dn}$, $B_2 := \oplus_{i=1}^n B_{2,i} \subset H^{dn}$, and $B_T := G_T^d$. Define the pairing $E : B_1 \times B_2 \to B_T$ as in Equation 1.
7. Finally, to be able to check that an element $g^M \in G^{dn}$ for $M = (m_{ij})_{i,j=1}^{d,n}$ is an element of B_1, we observe that the vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$ span an n-dimensional subspace $\mathbb{V}$ of $\mathbb{F}_p^{dn}$. Thus, there must be another subspace, call it $\mathbb{W}$, of dimension $dn - n$, that contains all vectors in $\mathbb{F}_p^n$ that are orthogonal to vectors in $\mathbb{V}$. Given $\mu_2 := (h^{\boldsymbol{w}_1}, \ldots, h^{\boldsymbol{w}_{(d-1)n}})$, where the $\{\boldsymbol{w}_i\}_{i=1}^{(d-1)n}$ are a basis of $\mathbb{W}$, one can therefore efficiently check if $g^M \in B_1$ by checking if $E(g^M, h^{\boldsymbol{w}_i}) = (1_T, \ldots, 1_T)$ for all i, $1 \le i \le (d-1)n$.
Analogously, given $\mu_1 := (g^{\boldsymbol{w}_1^*}, \ldots, g^{\boldsymbol{w}_{(d-1)n}^*})$, one can check if $h^A \in B_2$ by checking if $E(g^{\boldsymbol{w}_i^*}, h^A) = (1_T, \ldots, 1_T)$, where $\{\boldsymbol{w}_i^*\}_{i=1}^{(d-1)n}$ are a basis for the subspace $\mathbb{W}^*$ of $\mathbb{F}_p^n$ consisting of vectors orthogonal to vectors in the span of $\boldsymbol{v}_1^*, \ldots, \boldsymbol{v}_n^*$.
8. Output $\mathbb{G} := (p, B_1, B_2, B_T, E, (\mu_1, \mu_2))$.

can be computed efficiently by anyone knowing $\boldsymbol{v}_i$ and another vector in $\mathbb{F}_p^{dn}$ that is orthogonal to $\boldsymbol{v}_k$ for all $k \neq i$. The situation for π_{2i} is analogous.

As for the projection maps $\pi_{T,i}$ required for the target space, we define $\pi_{T,i}$ to map an element $e(g,h)^{a_1C_1+\cdots+a_nC_n}$ to $e(g,h)^{a_iC_i}$, where we recall C_i denotes the i-th row of the scaling matrix C (C_i is thus a vector in $\mathbb{F}_p^d$ for all i).

Finally, we show that the required associativity property holds, namely that $E(\pi_{1,i}(g^M), \pi_{2,i}(h^A)) = \pi_{T,i}(E(g^M, h^A))$ for all elements $g^M \in B_1$, $h^A \in B_2$, and for all i, $1 \le i \le d$. To see this, observe that $g^M \in B_1$ implies that $g^M = g^{\alpha_1\boldsymbol{v}_1+\cdots+\alpha_n\boldsymbol{v}_n}$ for some $\alpha_1, \ldots, \alpha_n \in \mathbb{F}_p$, and similarly that $h^A = h^{\beta_1\boldsymbol{v}_1^*+\cdots+\beta_n\boldsymbol{v}_n^*}$. We therefore have that

$$E(\pi_{1,i}(g^M), \pi_{2,i}(h^A)) = E(g^{\alpha_i\boldsymbol{v}_i}, h^{\beta_i\boldsymbol{v}_i^*}) = e(g,h)^{\alpha_i\beta_iC_i},$$

where this last equality follows from Lemma 3.2. On the other hand, we have that

$$\pi_{T,i}(E(g^M, h^A)) = \pi_{T,i}(\prod_{k=1}^{n} e(g,h)^{\alpha_k\beta_kC_k}) = e(g,h)^{\alpha_i\beta_iC_i},$$

and the two quantities are therefore equal.

A similar argument applies to obtaining more limited projections when C has lower rank. □

It remains to prove that our construction also satisfies parameter hiding and subgroup hiding. For the latter property, our definition in Section 2.3 is highly general and we cannot prove that all instances of generalized correlated subgroup decision reduce to any one assumption. Instead, we show that certain "nice" instances of the assumption follow from SXDH.

Before we define a nice instance, we first restrict our attention to the case where $n = 8$, $d = 1$, C is a matrix with all 1 entries. For succinctness here and in later sections, we use $\mathsf{BasicGen}(1^k) = \mathsf{BilinearGen}'(1^k, 8, 1, \mathcal{D})$, where $\mathcal{D}$ produces matrices with all 1 entries; i.e., we use $\mathsf{BasicGen}$ to produce the specific setting in which we are interested in Section 5.

We consider two variants of this setting, which differ only in the auxiliary information μ. For μ as defined above in Algorithm 1, we show that the required instances of the correlated subgroup decision assumption are implied by SXDH. We additionally consider a case where μ is augmented to contain the following three pieces of information: (1) the vectors $\boldsymbol{v}_7$, $\boldsymbol{v}_8$, $\boldsymbol{v}_7^*$, and $\boldsymbol{v}_8^*$; (2) a random basis for the span of $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_6)$ inside $\mathbb{F}_p^8$; and (3) a random basis for the span of $(\boldsymbol{v}_1^*, \ldots, \boldsymbol{v}_6^*)$ inside $\mathbb{F}_p^8$. With this μ, one can then perform a membership test for $G_1 \oplus \ldots \oplus G_6$ on some element $g^{\boldsymbol{v}}$ by computing a basis for the orthogonal space of the span of $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_6)$, pairing against h raised to these vectors, and taking a dot product in $\mathbb{F}_p^8$. While this additional information in μ makes some instances of subgroup decision easy, instances entirely within $G_1 \oplus \ldots \oplus G_6$ and $H_1 \oplus \ldots H_6$ are still implied by SXDH. To refer to this instance with augmented μ in what follows, we call it the *augmented construction*. Now, by "nice," we mean that the instance of the assumption behaves as follows: if the challenge terms are in H (the situation is analogous if they are in G), then there is a single pair in S that is common to the challenge sets T_1 and T_2 that appears in all tuples in $S^{\mathsf{sgh}}_{G,\mathrm{sam}}$ that also contain the differing pair. In other words, the given correlated samples from the opposite side of the challenge that include the differing space must also be attached to a particular space that is guaranteed to be present in the challenge term. As we will see, this feature turns out to be convenient for reducing to SXDH, as demonstrated by the following lemmas. For the augmented construction, we additionally restrict to instances where each correlated sample t_i in $S^{\mathsf{sgh}}_{G,\mathrm{sam}}$ or $S^{\mathsf{sgh}}_{H,\mathrm{sam}}$ is contained within the set $S := \{(1,2),(3,4),(5,6)\}$ (this is to avoid the additional information in μ from compromising the hardness).

Lemma 3.3. *For the augmented construction, the nice instances of the generalized correlated subgroup decision assumption, where additionally each correlated sample t_i in $S^{\mathsf{sgh}}_{G,sam}$ or $S^{\mathsf{sgh}}_{H,sam}$ is contained within the set $\{(1,2),(3,4),(5,6)\}$, are implied by the SXDH assumption.*

Proof. We consider a nice instance of the generalized correlated subgroup decision assumption parameterized by sets S^{sgh}_G and S^{sgh}_H containing singletons and tuples of the pairs $(1,2)$, $(3,4)$, $(5,6)$ and challenge sets T_1 and T_2 differing by one pair. We assume without loss of generality that the differing pair is $(3,4)$, that $(1,2)$ is a common pair to both T_1, T_2, and the challenge terms are in G.

We assume we are given an SXDH challenge of the form (g, h, g^a, g^b, T), where $T = g^{ab}$ or is random in G. We will simulate the specified instance of the

generalized correlated subgroup decision assumption. We first choose a random dual orthonormal bases pair $\mathbb{F}, \mathbb{F}^*$ for $\mathbb{F}_p^8$. We then implicitly define $\mathbb{B}, \mathbb{B}^*$ as follows:

$$\boldsymbol{b}_1 = a\boldsymbol{f}_3 + \boldsymbol{f}_1,\ \boldsymbol{b}_2 = a\boldsymbol{f}_4 + \boldsymbol{f}_2,\ \boldsymbol{b}_3 = \boldsymbol{f}_3,\ \boldsymbol{b}_4 = \boldsymbol{f}_4,$$
$$\boldsymbol{b}_5 = \boldsymbol{f}_5,\ \boldsymbol{b}_6 = \boldsymbol{f}_6,\ \boldsymbol{b}_7 = \boldsymbol{f}_7,\ \boldsymbol{b}_8 = \boldsymbol{f}_8$$
$$\boldsymbol{b}_1^* = \boldsymbol{f}_1^*,\ \boldsymbol{b}_2^* = \boldsymbol{f}_2^*,\ \boldsymbol{b}_3^* = \boldsymbol{f}_3^* - a\boldsymbol{f}_1^*,\ \boldsymbol{b}_4^* = \boldsymbol{f}_4^* - a\boldsymbol{f}_2^*,$$
$$\boldsymbol{b}_5^* = \boldsymbol{f}_5^*,\ \boldsymbol{b}_6^* = \boldsymbol{f}_6^*,\ \boldsymbol{b}_7^* = \boldsymbol{f}_7^*,\ \boldsymbol{b}_8^* = \boldsymbol{f}_8^*.$$

We note that $(\mathbb{B}, \mathbb{B}^*)$ are properly distributed, since applying a linear transformation to randomly sampled dual orthonormal bases while preserving orthonormality produces equivalently distributed bases. We observe that $\boldsymbol{v}_7, \boldsymbol{v}_8, \boldsymbol{v}_7^*, \boldsymbol{v}_8^*$ are known, as are the spans of $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_6\}$ and $\{\boldsymbol{v}_1^*, \ldots, \boldsymbol{v}_6^*\}$. Thus we can produce the specified auxiliary information μ.

Since we have h, g, g^a, we can produce all generators *except* h_3, h_4. Since $(3, 4)$ is the differing pair for the challenges, these generators cannot be required. Since all generators are known on the G side, any correlated samples in G are easy to produce. To produce correlated samples for tuples containing $(1, 2)$ and $(3, 4)$ in H, we simply choose random exponents $t', z \in \mathbb{F}_p$ and implicitly set $t = az + t'$. We can then produce

$$h_1^t h_3^z = h^{t'\boldsymbol{f}_1^* + z\boldsymbol{f}_3^*},\ h_2^t h_4^z = h^{-t'\boldsymbol{f}_2^* - z\boldsymbol{f}_4^*}.$$

To produce the challenge terms, we compute

$$T^{\boldsymbol{f}_3}(g^b)^{\boldsymbol{f}_1},\ T^{\boldsymbol{f}_4}(g^b)^{\boldsymbol{f}_2}.$$

If $(5, 6)$ is also common to T_1, T_2, we can use the generators g_5, g_6 to add on properly distributed terms in these subgroups as well. □

The same proof can also be applied more generally when μ is *not* augmented, resulting in:

Lemma 3.4. *For* $\mathbb{G} \xleftarrow{\$} \mathsf{BasicGen}(1^k)$*, all nice instances of the generalized correlated subgroup decision assumption are implied by SXDH.*

Finally, we prove that parameter hiding holds for the augmented construction as well.

Lemma 3.5. *Parameter hiding, as in Example 1, holds for the augmented construction.*

Proof. This is essentially Lemmas 3 and 4 in [27], and is a consequence of the following observation. We consider sampling a random pair of dual orthonormal bases $\mathbb{F}, \mathbb{F}^*$ of $\mathbb{F}_p^8$, and let A be an invertible 2×2 matrix over $\mathbb{F}_p$. We consider the 8×2 matrix F whose columns are equal to $\boldsymbol{f}_3$ and $\boldsymbol{f}_4$. Then FA is also an 8×2 matrix, and we form a new basis $\mathbb{B}$ from $\mathbb{F}$ and A by taking these columns in place of $\boldsymbol{f}_3, \boldsymbol{f}_4$. To form the dual basis $\mathbb{B}^*$, we similarly multiply the

matrix with columns $\boldsymbol{f}_3^*, \boldsymbol{f}_4^*$ by the transpose of A^{-1}. It is noted in [27] that the resulting distribution of $\mathbb{B}, \mathbb{B}^*$ is equivalent to choosing this pair randomly, and in particular, this distribution is independent of the choice of A. Lemma 4 in [27] observes that if we take $x \neq y$ and define $\boldsymbol{x}$ to be the transpose of $(1, x)$ and $\boldsymbol{y}$ to be the transpose of $(y, -1)$, then choosing random scalars γ, λ in $\mathbb{F}_p$ and a random matrix A over $\mathbb{F}_p$ yields that the joint distribution of $\lambda A^{-1}\boldsymbol{x}$ and $\gamma A^T \boldsymbol{y}$ is negligibly close to the uniform distribution over $\mathbb{F}_p^2 \times \mathbb{F}_p^2$. This is precisely our parameter hiding requirement, where A represents the ambiguity in our precise choice of the generators $\boldsymbol{b}_3, \boldsymbol{b}_4, \boldsymbol{b}_3^*, \boldsymbol{b}_4^*$, conditioned on the span of $\{\boldsymbol{b}_3, \boldsymbol{b}_4\}$ and the span of $\{\boldsymbol{b}_3^*, \boldsymbol{b}_4^*\}$ being known (in addition to the other individual $\boldsymbol{b}_i$ and $\boldsymbol{b}_i^*$ vectors for $i \notin \{3, 4\}$). □

Finally, although we do not use any non-nice instances of the generalized correlated subgroup decision assumption in this work, it is interesting to ask which of the more complex instances can be reduced to SXDH or other static assumptions. For values of $d > 1$, the additional structure required to achieve projecting seems to make directly reducing a large space of assumptions to SXDH difficult. Nonetheless, we are able to rely only on SXDH for our projecting leakage-resilient BGN variant through the use of hybrid transitions that incrementally change the rank of the scaling matrix C. We leave it as an interesting question for future work to further explore the minimal assumptions for supporting a broader class of subgroups decision variants.

4 A Leakage-Resilient BGN Variant

A very elegant use of the projecting property in the composite-order setting is the public key encryption scheme of Boneh, Goh, and Nissim [9], a scheme that is designed to allow arbitrary additions and one multiplication of ciphertexts. The basic group operation is used for ciphertext addition, while the bilinear map is applied during ciphertext multiplication. The secret key is then a projection map (which equates to a factorization of the group order) that allows the decryptor to strip off the blinding factors of the underlying ciphertexts, even after their interaction has migrated to the target group.

While these limited homomorphic properties make the BGN scheme appealing, the rigid structure of keys can be a source of frustration when one attempts to augment its functionality or security guarantees. Having the secret key reveal a factorization of the group order means that different users must generate different groups, and it additionally means that the secret key is uniquely determined (information-theoretically) from the public key. This presents a challenge, for instance, if one wants to design a variant with provable guarantees of leakage resilience.

Proofs of leakage resilience for public key encryption schemes typically follow a strategy inspired by the hash proof paradigm of Naor and Segev [37]. This paradigm starts with a scheme that has many possible secret keys for each public key. A hybrid argument is used, where the first step changes to a malformed — or *invalid* — ciphertext, that decrypts to different messages under the different

secret keys associated to a fixed public key. A bound on the total leakage of the secret key is then used to argue that the adversary cannot tell which of the many possible secret keys the challenger is holding. Thus, even though the challenger may be holding a secret key that decrypts the challenge ciphertext correctly, he may as well be a holding a key that decrypts it to a random message. It is then possible to argue that the scheme remains secure under leakage.

If we wish to apply this kind of proof strategy to a version of the BGN scheme, we first need a way of allowing many secret keys for each public key. The DPVS framework we described in the previous section provides a natural answer. In this framework, the projection map is no longer a factorization, but rather a vector that comes from a suitably high-dimensional space to allow for many possibilities. This makes it rather easy to imagine a BGN variant that preserves the somewhat-homomorphic properties of ciphertexts, yet allows for an exponential number of secret keys per public key.

It is already well-known that applying DPVS and similar techniques for designing vector spaces in the exponent is a useful approach for achieving leakage resilience. For example, Lewko et al. [35] demonstrated that leakage resilience can be incorporated quite easily into dual system encryption proofs by combining mechanisms for canceling, parameter hiding, and the fact that the dot product of sufficiently long vectors over $\mathbb{F}_p$ has convenient information-theoretic properties (roughly, the dot product modulo p is a good two-source extractor). The same high level of compatibility exists between our framework and the pre-existing leakage resilience techniques, thus allowing us to repurpose the same linear algebraic underpinnings that implement projecting and canceling in our framework to achieve leakage resilience for a BGN-type scheme.

4.1 The Scheme

As in the original BGN scheme, we will assume that the message space is small to allow efficient decryption. We use our framework from Section 3 with $n = d = 4$. For the matrix distribution $\mathcal{D}$, we consider all matrices whose second and third rows form a rank-1 submatrix. The setting we then work in is $\mathbb{G} \xleftarrow{\$} \mathsf{BilinearGen}'(1^k, 4, 4, \mathcal{D})$. Rather than use this framework generically, as we do in Section 5, we re-purpose the matrix C and basis vectors $(\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4), (\boldsymbol{v}_1^*, \boldsymbol{v}_2^*, \boldsymbol{v}_3^*, \boldsymbol{v}_4^*) \in \mathbb{F}_p^{16}$ — defined in Step 4 and Step 5 of Algorithm 1 respectively — and use them explicitly in our construction and proofs. Below, we use C_i to denote the i-th row of the scaling matrix C (for $i \in \{1, 2, 3, 4\}$).

- $\mathsf{Setup}(\mathbb{G})$: Pick $r, r^* \xleftarrow{\$} \mathbb{F}_p$ and define $\boldsymbol{u} := \sum_i \boldsymbol{v}_i$, $\boldsymbol{u}^* := \sum_i \boldsymbol{v}_i^*$, $\boldsymbol{w} := r\boldsymbol{v}_2$, and $\boldsymbol{w}^* := r^*\boldsymbol{v}_2^*$. Choose $\boldsymbol{y}$ uniformly at random from the set of vectors in $\mathbb{F}_p^4$ such that $\boldsymbol{y} \cdot C_2 = 0$, noting that $\boldsymbol{y} \cdot C_3 = 0$ then holds automatically as well. Output $pk = (g, g^{\boldsymbol{u}}, g^{\boldsymbol{w}}, h^{\boldsymbol{u}^*}, h^{\boldsymbol{w}^*})$ and $sk = \left(\boldsymbol{y}, sk_T = e(g,h)^{\boldsymbol{y} \cdot (\sum_i C_i)}\right)$. Note that, by construction, $\boldsymbol{y} \cdot (\sum_i C_i) = \boldsymbol{y} \cdot (C_1 + C_4)$ and, by Lemma 3.2, $E(g^{\boldsymbol{u}}, h^{\boldsymbol{u}^*}) = \left(e(g,h)^{\sum_j c_{j,1}}, \ldots, e(g,h)^{\sum_j c_{j,4}}\right)$.

– $\mathsf{Enc}(pk, m)$: We have two types of ciphertexts: Type A and Type B. If we want to be able to perform homomorphic operations on *any* pair of ciphertexts, a single ciphertext could include both types. To form a Type A ciphertext, choose $s \xleftarrow{\$} \mathbb{F}_p$ and compute $\mathsf{ct}_A := g^{m\boldsymbol{u}+s\boldsymbol{w}}$. To form a Type B ciphertext, choose $s^* \xleftarrow{\$} \mathbb{F}_p$ and compute $\mathsf{ct}_B := h^{m\boldsymbol{u}^*+s^*\boldsymbol{w}^*}$. Output $\mathsf{ct} = (\mathsf{ct}_A, \mathsf{ct}_B)$. (Or just ct_A or ct_B, depending on the desired homomorphic properties.)
– $\mathsf{Eval}(pk, \mathsf{ct}_1, \mathsf{ct}_2)$: We describe two evaluation cases: addition of Type A ciphertexts (the operations are analogous for Type B ciphertexts), and multiplication of a Type A and Type B ciphertext (which can then be further added in the target space B_T).

  First pick a random value $t \xleftarrow{\$} \mathbb{F}_p$. If ct_1 and ct_2 are Type A, then return $\mathsf{ct} = \mathsf{ct}_1 \cdot \mathsf{ct}_2 \cdot g^{t\boldsymbol{w}}$. If ct_1 is Type A and ct_2 is Type B, then return $\mathsf{ct} = E(\mathsf{ct}_1, \mathsf{ct}_2) \cdot E(g^{\boldsymbol{w}}, h^{\boldsymbol{w}^*})^t$.
– $\mathsf{Dec}(sk, \mathsf{ct})$: To decrypt a ciphertext $(\mathsf{ct}_1, \mathsf{ct}_2, \mathsf{ct}_3, \mathsf{ct}_4) \in G_T^4$, compute

$$\prod_{i=1}^{4} \mathsf{ct}_i^{y_i} = sk_T^m.$$

  Using knowledge of sk_T, exhaustively search for m (this is possible since we have a small message space). If ct is Type A, then compute $\mathsf{ct}' = E(\mathsf{ct}, \mathsf{Enc}(pk, 1))$ and decrypt ct' (and analogously for a Type B ciphertext).

To see that decryption is correct, observe that

$$\begin{aligned}\prod_i \mathsf{ct}_i^{y_i} &= \prod_i e(g,h)^{m y_i \sum_j c_{j,i}} = e(g,h)^{m \sum_i \sum_j y_i c_{j,i}} \\ &= e(g,h)^{m \sum_j \sum_i y_i c_{j,i}} = e(g,h)^{m \sum_j \boldsymbol{y} \cdot C_j} \\ &= sk_T^m.\end{aligned}$$

To see that evaluation is correct, observe that if ct_1 encrypts m_1 and ct_2 encrypts m_2 then

$$\mathsf{ct} = g^{m_1\boldsymbol{u}+s_1\boldsymbol{w}} \cdot g^{m_2\boldsymbol{u}+s_2\boldsymbol{w}} \cdot g^{t\boldsymbol{w}} = g^{(m_1+m_2)\boldsymbol{u}+(s_1+s_2+t)\boldsymbol{w}},$$

which is a properly distributed Type A encryption of $m_1 + m_2$. Pairing a Type A ct_1 and a Type B ct_2 similarly yields a properly distributed encryption of $m_1 m_2$ in the target space, just as in BGN.

4.2 Security Analysis

The security model we use is leakage against non-adaptive memory attacks, as defined by Akavia et al. [3, Definition 3]. Briefly, the attacker first declares a leakage function f mapping secret keys to $\{0,1\}^\ell$ for a suitably small ℓ. The attacker then receives pk and $f(sk)$, and proceeds as in a standard IND-CPA

game; i.e., it outputs two messages m_0 and m_1, receives an encryption of m_b, and wins if it correctly guesses b. As in the case of the original BGN scheme, it suffices to argue security for challenge ciphertexts generated in G/H, as security for the ciphertexts generated via the multiplicative homomorphism follows from the security of ciphertexts in the base groups. While there are several other interesting models for leakage-resilient PKE security, we choose to work with this one, as it is clean and simple and thus allows us to give a concise demonstration of the use of our framework.

Theorem 4.1. *If SXDH holds in* $\mathbb{G}$ *and* $\ell \leq \log(p-1) - 2k$*, the above construction is leakage resilient with respect to non-adaptive memory attacks.*

As in the typical hash proof system paradigm, we first define invalid ciphertexts that have more blinding randomness than honestly generated ciphertexts. Initially, these are still decrypted consistently by the set of secret keys corresponding to a fixed public key. After having transitioned to a game with an invalid challenge ciphertext, however, we gradually adjust the respective distributions of secret keys and ciphertexts to arrive at a game where, in the adversary's view, it seems that the secret key decrypts the ciphertext randomly.

In the course of these game transitions, we use SXDH in multiple ways. First we use it to change from an honest to an invalid ciphertext by bringing in an additional blinding factor in a new subgroup. This is just a "nice" instance of subgroup decision. We will also use it to make changes to the rank of particular submatrices inside the scaling matrix C. This technique is inspired by the observation in [10] that DDH implies a rank-1 matrix in the exponent is hard to distinguish from a rank-2 matrix. To make the crucial switch from a secret key that properly decrypts the challenge ciphertext to a key that decrypts it incorrectly, we rely on an information-theoretic argument leveraging a form of parameter hiding, along with the leakage bound. Essentially, the simulator uses the remaining ambiguity in the underlying parameters (conditioned on the public key) to help it create an invalid challenge ciphertext after supplying the leakage. The proof of Theorem 4.1 can be found in the full version of our paper [28].

5 An IBE with IND-CCA1 Security

In this section, we discuss how to obtain an IND-CCA1-secure identity-based encryption scheme. Although IND-CCA2-secure IBE schemes have already been constructed, we view this as a demonstration of our techniques rather than an application of independent interest.

Our technique for proving IND-CCA1 security extends from the observation due to Lewko and Waters [33] that dual system proofs can be interpreted as a reduction from a full security game to a weak game in which the attacker does not have access to the public parameters. Using this technique, we first define such a weak game for IND-CCA1 security, and then prove that our IBE construction satisfies it. Next, leveraging a weak form of projection, we reduce the full IND-CCA1 security to this weaker notion by first expanding the system to have extra

components in a space that is not reflected in the public parameters, and then projecting to play the weak game in that space.

In the full version of our paper [28], we formulate our IBE and proof in a unified framework that can be instantiated in either prime-order groups or in composite-order groups. In the prime-order setting, we obtain the following result:

Theorem 5.1. *If SXDH holds in* $\mathbb{G} \xleftarrow{\$} \mathsf{BilinearGen}'(1^k, 8, 1, \mathcal{D})$*, where* $\mathcal{D}$ *produces the vector* $\mathbf{1}$*, then the instantiation of our IBE construction is IND-CCA1 secure.*

References

1. Adj, G., Menezes, A., Oliveira, T., Rodríguez-Henríquez, F.: Weakness of $\mathbb{F}_{3^{6\cdot 509}}$ for discrete logarithm cryptography. In: Cao, Z., Zhang, F. (eds.) Pairing 2013. LNCS, vol. 8365, pp. 20–44. Springer, Heidelberg (2014)
2. Adj, G., Menezes, A., Oliveira, T., Rodríguez-Henríquez, F.: Computing discrete logarithms in $f_{3^{6\cdot 137}}$ and $f_{3^{6\cdot 163}}$ using magma. Cryptology ePrint Archive, Report 2014/057 (2014). http://eprint.iacr.org/2014/057
3. Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous hardcore bits and cryptography against memory attacks. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 474–495. Springer, Heidelberg (2009)
4. Aranha, D.F., Beuchat, J.-L., Detrey, J., Estibals, N.: Optimal eta pairing on supersingular genus-2 binary hyperelliptic curves. In: Dunkelman, O. (ed.) CT-RSA 2012. LNCS, vol. 7178, pp. 98–115. Springer, Heidelberg (2012)
5. Aranha, D.F., Karabina, K., Longa, P., Gebotys, C.H., López, J.: Faster explicit formulas for computing pairings over ordinary curves. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 48–68. Springer, Heidelberg (2011)
6. Ballard, L., Green, M., de Medeiros, B., Monrose, F.: Correlation-resistant storage via keyword-searchable encryption. Cryptology ePrint Archive, Report 2005/417 (2005). http://eprint.iacr.org/
7. 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)
8. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004)
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., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision Diffie-Hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)
11. Boneh, D., Rubin, K., Silverberg, A.: Finding ordinary composite order elliptic curves using the Cocks-Pinch method. Journal of Number Theory **131**(5), 832–841 (2011)
12. Boyen, X., Waters, B.: Compact group signatures without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 427–444. Springer, Heidelberg (2006)

13. Boyen, X., Waters, B.: Full-domain subgroup hiding and constant-size group signatures. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 1–15. Springer, Heidelberg (2007)
14. Brakerski, Z., Kalai, Y.T., Katz, J., Vaikuntanathan, V.: Overcoming the hole in the bucket: public-key cryptography resilient to continual memory leakage. In: 51st FOCS, Las Vegas, Nevada, USA, 23–26 October 2010, pp. 501–510. IEEE Computer Society Press (2010)
15. 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)
16. Dodis, Y., Haralambiev, K., López-Alt, A., Wichs, D.: Cryptography against continuous memory attacks. In: 51st FOCS, Las Vegas, Nevada, USA, 23–26 October 2010, pp. 511–520. IEEE Computer Society Press (2010)
17. Dodis, Y., Lewko, A.B., Waters, B., Wichs, D.: Storing secrets on continually leaky devices. In: Ostrovsky, R. (ed.) 52nd FOCS, Palm Springs, California, USA, 22–25 October 2011, pp. 688–697. IEEE Computer Society Press (2011)
18. 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)
19. Galbraith, S., Paterson, K., Smart, N.: Pairings for cryptographers. Discrete Applied Mathematics **156**(16), 3113–3121 (2008)
20. Gentry, C., Lewko, A., Waters, B.: Witness encryption from instance independent assumptions. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 426–443. Springer, Heidelberg (2014)
21. Göloğlu, F., Granger, R., McGuire, G., Zumbrägel, J.: On the function field sieve and the impact of higher splitting probabilities. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 109–128. Springer, Heidelberg (2013)
22. Guillevic, A.: Comparing the pairing efficiency over composite-order and prime-order elliptic curves. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 2013. LNCS, vol. 7954, pp. 357–372. Springer, Heidelberg (2013)
23. Hayashi, T., Shimoyama, T., Shinohara, N., Takagi, T.: Breaking pairing-based cryptosystems using η_T pairing over $gf(3^{97})$. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 43–60. Springer, Heidelberg (2012)
24. Herold, G., Hesse, J., Hofheinz, D., Ràfols, C., Rupp, A.: Polynomial spaces: a new framework for composite-to-prime-order transformations. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 261–279. Springer, Heidelberg (2014)
25. Joux, A.: Faster index calculus for the medium prime case application to 1175-bit and 1425-bit finite fields. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 177–193. Springer, Heidelberg (2013)
26. 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)
27. Lewko, A.: Tools for simulating features of composite order bilinear groups in the prime order setting. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 318–335. Springer, Heidelberg (2012)
28. Lewko, A., Meiklejohn, S.: A profitable sub-prime loan: obtaining the advantages of composite order in prime-order bilinear groups. Cryptology ePrint Archive, Report 2013/300 (2013). http://eprint.iacr.org/2013/300

29. 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)
30. 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)
31. Lewko, A., Waters, B.: Decentralizing attribute-based encryption. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 568–588. Springer, Heidelberg (2011)
32. 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)
33. Lewko, A., Waters, B.: New proof methods for attribute-based encryption: achieving full security through selective techniques. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 180–198. Springer, Heidelberg (2012)
34. Lewko, A.B., Lewko, M., Waters, B.: How to leak on key updates. In: Fortnow, L., Vadhan, S.P. (eds.) 43rd ACM STOC, San Jose, California, USA, 6–8 June 2011, pp. 725–734. ACM Press (2011)
35. 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)
36. 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)
37. Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009)
38. 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)
39. 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)
40. 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)
41. Okamoto, T., Takashima, K.: Adaptively attribute-hiding (hierarchical) inner product encryption. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 591–608. Springer, Heidelberg (2012)
42. Scott, M.: On the efficient implementation of pairing-based protocols. In: Chen, L. (ed.) IMACC 2011. LNCS, vol. 7089, pp. 296–308. Springer, Heidelberg (2011)
43. Seo, J.H.: On the (im)possibility of projecting property in prime-order setting. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 61–79. Springer, Heidelberg (2012)
44. Seo, J.H., Cheon, J.H.: Beyond the limitation of prime-order bilinear groups, and round optimal blind signatures. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 133–150. Springer, Heidelberg (2012)
45. 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)