Tight Adaptively Secure Broadcast Encryption with Short Ciphertexts and Keys

Romain Gay*
ENS, Paris, France
romain.gay@ens.fr

Lucas Kowalczyk†
Columbia University
luke@cs.columbia.edu

Hoeteck Wee‡
ENS, Paris, France
wee@di.ens.fr

Abstract

We present a new public key broadcast encryption scheme where both the ciphertext and secret keys consist of a constant number of group elements. Our result improves upon the work of Boneh, Gentry and Waters (Crypto '05) as well as several recent follow-ups (TCC '16-A, Asiacrypt '16) in two ways: (i) we achieve adaptive security instead of selective security, and (ii) our construction relies on the decisional k-Linear Assumption in prime-order groups (as opposed to q-type assumptions or subgroup decisional assumptions in composite-order groups); our improvements come at the cost of a larger public key. Finally, we show that our scheme achieves adaptive security in the multi-ciphertext setting with a security loss that is independent of the number of challenge ciphertexts.

1 Introduction

Broadcast encryption schemes [FN94] allow a sender to encrypt messages to a set $\Gamma \subset [n]$ of authorized users such that any user in the set Γ can decrypt, and no (possibly colluding) set of unauthorized users can learn anything about the plaintext. Two key measures of efficiency for broadcast encryption are the size of the secret keys and the ciphertext overhead (beyond description of the recipient set and the symmetric encryption of the message). The early contructions of broadcast encryption schemes achieve ciphertext overhead that grows with the number of either authorized or excluded users [NNL01, HS02, DF02, GST04].

The BGW Cryptosystem. Ideally, we would like a broadcast encryption scheme where the size of secret keys and ciphertext overhead is independent of the number of users. This was first achieved in the break-through work of Boneh, Gentry and Waters [BGW05], which presented a broadcast encryption scheme in bilinear groups where both the secret keys and ciphertext overhead consist of a constant number of group elements. In their scheme, the decryption algorithm needs to know the public key, which is linear in the number of users.

*Supported in part by a Google Fellowship.

†Work done while visiting ENS, Paris. Supported in part by the Defense Advanced Research Project Agency (DARPA) and Army Research Office (ARO) under Contract W911NF-15-C-0236; NSF grants CNS-1445424, CNS-1552932, and CCF-1423306; and an NSF Graduate Research Fellowship DGE-16-44869. Any opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the Defense Advanced Research Projects Agency, Army Research Office, the National Science Foundation, or the U.S. Government.

‡Supported in part by ERC Project aSCEND (H2020 639554).

The BGW cryptosystem has two main limitations, which is the focus of several follow-up works as well as our current one:

- First, the BGW scheme achieves selective security, where an adversary must declare a target set of unauthorized users with which it will attack the scheme *before* even seeing the system parameters. This restriction does not capture the power of many kinds of attackers (for instance: an attacker might choose to corrupt a user after seeing the public parameters, or in response to seeing secret keys for already corrupted parties), so in practice, we would prefer to have schemes that satisfy the more general and stronger notion of adaptive security, which does not place such restrictions on the adversary.

- Next, the BGW scheme relies on parameterized assumptions. Parameterized assumptions (a.k.a q-type assumptions), while in some cases allowing for improvements over the state-of-the-art, are not particularly well understood. The assumptions are often closely related to the schemes which use them. For example, the size of the assumption often scales with the number of oracle queries that can be made in the security reduction. Furthermore, q-type assumptions become stronger as q grows, with the time needed to recover the discrete logarithm and break the assumption scaling inversely with q [Che06]. As a result, it is desirable to design systems that can be proven secure under static assumptions, like the decisional k-Linear Assumption in prime-order bilinear groups (k-Lin).

These limitations were fixed individually by the works of [GW09] and [Wee16, CMM16a] respectively (the latter in composite-order groups), but improving [BGW05] to achieve security that is *both* adaptive and based on a static assumption has remained out of reach.

1.1 Our Results

In this paper we present the first broadcast encryption scheme with constant key and ciphertext overhead size that simultaneously overcomes both of the limitations above. Namely, we achieve adaptive security under a static assumption (k-Lin) in prime-order bilinear groups. Our improvements come at the cost of a larger public key that is quadratic instead of linear in the total number of users. We stress that prior to this work, it was not known how to achieve broadcast encryption with any size public parameters, constant-sized keys and ciphertext overhead, and even just *selective* security under a static assumption in prime-order groups.

As with the BGW cryptosystem and the follow-up works in [Wee16, CMM16a], the decryption algorithm in our scheme needs to know the public key in addition to the secret key. Considering the complications that come with managing user secret keys, which have to be distributed individually and stored securely, we achieve a desirable public/private key size tradeoff that makes sense particularly in applications where decryptors have access to large shared public storage.

We give an additional broadcast encryption scheme with constant key and ciphertext overhead size which is adaptively-secure *in the multi-challenge setting* under static assumptions with a *tight security reduction* (where the security loss is independent of the number of challenge ciphertexts). Tight security reductions, which have been studied previously in the context of encryption [BBM00, HJ12] and signatures [Cor00], are desirable when fixing concrete security parameters, since the security loss directly impacts the size of scheme elements. In the context of advanced encryption schemes, tight constructions were only known for identity-based encryption [CW13]. In this work, we give the first tightly secure broadcast encryption scheme. Note that while our

security loss is independent of the number of challenge ciphertexts, it remains proportional to n: the number of users in the system. In this work, we view n as being not too large since our public key contains $O(n^2)$ group elements, which would be impractical for very large n anyway. Thus, a security loss of a small constant times n is much more desirable than one that is proportional to the number of challenge ciphertexts, which could be much larger for largely deployed systems.

1.2 Related Work

Previous broadcast encryption schemes for n users that are secure in the standard model either carry the baggage of a $(n/t, t)$-tradeoff in key/ciphertext size, use a non-static assumption (i.e q-type assumption), or are only secure in the weaker, selective security setting (see Figure 1). In fact, all known broadcast encryption schemes that are adaptively secure under a static assumption and that use the Dual System Encryption methodology [Att14, Wee14, CGW15, AC16] fall in the scope of the lower bound of $(n/t, t)$ for the (ciphertext overhead, secret key) size proved in [GKW15]. We note that we are able to bypass this lower bound by using the modified definition of broadcast encryption proposed by [BGW05], where decryption is allowed to take public parameters as input in addition to the secret key, as explained above.

Reference	$\lvert\mathsf{ct}\rvert$	$\lvert\mathsf{sk}\rvert$	$\lvert\mathsf{pk}\rvert$	assumption	security	Dec
BGW05 [BGW05]	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(n)$	q-type	selective	pk
GW09 [GW09]	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(n)$	q-type	adaptive	pk
Wee16[Wee16], CMM16[CMM16b]	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(n)$	composite	selective	pk
BW06 [BW06]	$\mathcal{O}(\sqrt{n})$	$\mathcal{O}(\sqrt{n})$	$\mathcal{O}(\sqrt{n})$	composite	adaptive	–
GKSW10 [GKSW10]	$\mathcal{O}(\sqrt{n})$	$\mathcal{O}(\sqrt{n})$	$\mathcal{O}(n)$	2-Lin	adaptive	–
Waters09 [Wat09]	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$	2-Lin	adaptive	–
GKW15 [GKW15]	$\mathcal{O}(n/t)$	$\mathcal{O}(t)$	$\mathcal{O}(n)$	k-Lin	adaptive	–
this work	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(n^2)$	composite	adaptive	pk
this work	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(n^2)$	k-Lin	adaptive	pk

Figure 1: Comparison amongst broadcast encryption schemes in the standard model, where n denotes the number of users, $\lvert\mathsf{ct}\rvert$, $\lvert\mathsf{sk}\rvert$ and $\lvert\mathsf{pk}\rvert$ respectively denote the ciphertext, secret key and public key size (i.e the number of group elements or exponents of group elements). The last column refers to whether or not the decryption algorithm Dec requires the public key pk as input.

Short keys and ciphertext overhead have been accomplished in other schemes by moving outside the standard model: [GW09] gives a construction (different from the one depicted in Figure 1 which uses q-type assumptions) with adaptive security and constant key and ciphertext overhead size, but in the random oracle model; [BWZ14] achieves adaptive security with polylogarithmic (in the number of users) size public parameters, keys, and ciphertext overhead, but is only proven secure in the multilinear generic group model; and [BZ14] achieves adaptive security with linear size public parameters, constant size keys and ciphertext overhead, but relies on strong assumptions, namely, indistinguishability obfuscation [BGI+01]. Lastly, we note that while our constructions harness the power of computational assumptions to achieve their efficiency, the problem of broadcast encryption has been studied in the information-theoretic realm as well [Sv98, SSW00, GSW00, GSY99].

1.3 Our Techniques

We give a construction in the composite-order setting which is secure under standard static decision assumptions to illustrate the main techniques, as well as a construction using prime-order bilinear groups which is secure under k-Lin.

Dual System Proof Methodology. We employ the dual system proof methodology [Wat09] to achieve the adaptive security of our schemes. A dual system encryption scheme is constructed so that an adversary cannot distinguish the distribution of normal keys (or ciphertexts) from special "semi-functional" keys (or ciphertexts). Semi-functional keys are capable of decrypting normal ciphertexts, but semi-functional keys cannot decrypt a semi-functional ciphertext. A typical dual system proof consists of a hybrid where the first step is constructing the challenge ciphertext as a semi-functional ciphertext. The hybrid then runs over each key requested by the adversary, replacing each requested key with a semi-functional key. At the end, only semi-functional keys are given to an adversary whose job is to break the security of a semi-functional ciphertext. Due to the way semi-functional ciphertexts and secret keys are constructed, it is typically easy to argue the game's security at this point (semi-functional secret keys cannot be used to decrypt *any* semi-functional ciphertexts, including the semi-functional challenge ciphertext).

Overview of the Construction Our constructions can be understood by starting with the Boneh-Gentry-Waters construction for broadcast encryption [BGW05], which is selectively-secure under a (non-static) q-type assumption. BGW's public parameters look like:

$$\mathsf{pk} := (g^{\gamma}, g^{\alpha}, g^{\alpha^2}, \ldots, g^{\alpha^n},\ h^{\alpha}, h^{\alpha^2}, \ldots, h^{\alpha^n},\ h^{\alpha^{n+2}}, \ldots, h^{\alpha^{2n}},\ \ e(g,h)^{\alpha^{n+1}})$$

where γ, α are random exponents in $\mathbb{Z}_p$, and g, h respectively generate prime order groups G, H, where $|G| = |H| = p$, and $e : G \times H \to G_T$.

The ciphertext for a subset $\Gamma \subseteq [n]$ and the key for a user $i \in [n]$ are given by:

$$\mathsf{ct}_{\Gamma} := (g^s,\ \ g^{(\gamma + \sum_{j \in \Gamma} \alpha^j)s},\ \ e(g,h)^{s\alpha^{n+1}} \cdot M), \quad \mathsf{sk}_i := h^{\alpha^{n-i+1}\gamma}$$

Decryption works as follows. Note that a message M in a ciphertext is hidden by an encapsulation key $e(g,h)^{s\alpha^{n+1}}$. First, an authorized user of index i pairs $h^{\alpha^{n-i+1}}$ from the public parameters with $g^{(\gamma+\sum_{j\in\Gamma}\alpha^j)s}$ from the ciphertext to get the encapsulation key hidden by a product of $e(g,h)^{s(n+1-i+j)}$ for $j \neq i \in \Gamma$ and $e(g,h)^{s\alpha^{n-i+1}\gamma}$. The former can be removed by performing judicious pairings with elements from pk and g^s from the ciphertext, and the latter can only by removed by computing the pairing of g^s with the (authorized) user's secret key sk_i. The encapsulation key can therefore be computed and used to obtain the message M.

The q-type assumption underlying BGW's security is enabled by the powers of α. These powers prevent a straightforward dual-system proof of adaptive security from static assumptions. To obtain a construction based on static assumptions, we need to remove the powers of α in the scheme. Towards this goal, consider the substitutions:

$$g^{\alpha^j} \mapsto g^{w_j}, \qquad h^{\alpha^{n-j+1}} \mapsto h^{r_j}, \qquad j \in [n]$$

where $w_1, \ldots, w_n, r_1, \ldots, r_n$ are chosen uniformly at random. Correctness of BGW scheme relies on the fact that

$$\{e(g^{\alpha^j s}, h^{\alpha^{n-i+1}})\}_{i,j\in[n], j\neq i}$$

lies in a set of linear size, namely

$$\{e(g^s, h^{\alpha}), \ldots, e(g^s, h^{\alpha^n}),\ e(g^s, h^{\alpha^{n+2}}), \ldots, e(g^s, h^{\alpha^{2n}})\}.$$

With our substitutions, the corresponding collection lies in a set

$$\{e(g^s, h^{w_j r_i})\}_{i,j\in[n],j\neq i}$$

of size $O(n^2)$, and hence the corresponding blow-up in the size of the public key, which needs to additionally contain $\{h^{w_j r_i}\}_{i,j\in[n],i\neq j}$.

Finally, replacing the prime-order pairing group by an composite-order asymmetric bilinear group (G, H, G_T) where $|G| = |H| = N = pq$, so as to use a subgroup membership assumption instead of the q-DBDH assumption used in BGW, and replacing $g \mapsto g_p$, $h \mapsto h_p$, where g_p, h_p respectively generate G_p, H_p: prime order subgroups of groups G, H, we obtain our composite-order scheme.

Alternative Viewpoint. As seen above, we can view our construction as a modification of the broadcast encryption scheme from [BGW05] where we improve the secret key/public key size trade-off. An alternative way to view our construction is to start from the broadcast encryption scheme of Waters [Wat09], which can be proven adaptively secure from static assumptions (using the dual system proof methodology) and features constant size ciphertext overhead, but linear size secret keys. We describe the construction using composite-order asymmetric bilinear groups for simplicity:

$$\begin{aligned}
\mathsf{pk} &:= \left(\{g_p^{w_j}\}_{j\in[n]},\ \ e(g_p, h_p)^{\alpha}\right)\\
\mathsf{ct}_\Gamma &:= \left(g_p^s,\ \ g_p^{s(u+\sum_{j\notin\Gamma} w_j)},\ \ e(g_p, h_p)^{s\alpha}\cdot M\right)\\
\mathsf{sk}_i &:= \left(h_p^{r_i},\ \ \{h_p^{w_j r_i}\}_{\substack{j\in[n],\\ j\neq i}},\ \ h_p^{\alpha+ur_i}\right)
\end{aligned}$$

where s, u, α, w_j, r_i for $i, j \in [n]$ are random exponents in $\mathbb{Z}_N$, and g_p, h_p respectively generate G_p, H_p: prime order subgroups of groups G, H, where $|G| = |H| = N = pq$, and $e : G \times H \to G_T$.

Decryption works as follows. Note that a message M in a ciphertext is again hidden by an encapsulation key $e(g_p, h_p)^{s\alpha}$. To get the encapsulation key $e(g_p, h_p)^{s\alpha}$, decryption pairs g_p^s with $h_p^{\alpha+ur_i}$. To get rid of the extra term $e(g_p, h_p)^{sur_i}$, it pairs $g_p^{s(u+\sum_{j\notin\Gamma} w_j)}$ from the ciphertext together with $h_p^{r_i}$. Doing so, decryption also gets many cross terms of the form $e(g_p, h_p)^{s\sum_{j\notin\Gamma} w_j r_i}$ which can be stripped away, pairing g_p^s with the appropriate $h_p^{w_j r_i}$ from the secret key. Note that these secret key elements are all available only when $i \in \Gamma$ and the key is therefore authorized.

To improve this construction's linear-sized secret keys to constant-size, we pre-compute the values $\{h_p^{r_i}, h_p^{w_j r_i}\}_{j\in[n],j\neq i}$ and include them in the public parameters instead of the secret key. Therefore, the secret key is reduced to the part that contains the encapsulation key α. Note that this crucially takes advantage of our modified model of broadcast encryption where decryption is allowed to use elements from the public key as well as the secret key.

Indeed, the main technical challenge in proving our schemes secure is to carry on the dual-system proof when the values $\{h_p^{r_i}, h_p^{w_j r_i}\}_{j\in[n],j\neq i}$ are public for **every** $i \in [n]$, and only a single group element remains private. This is in contrast to the security proof of previous dual system schemes, such as [Wat09], where the values $h_p^{r_i}, \{h_p^{w_j r_i}\}_{j\in[n],j\neq i}$ are known to the adversary only for queried keys sk_i. We solve it by carefully switching the $h_p^{r_i}, \{h_p^{w_j r_i}\}_{j\in[n],j\neq i}$ for each $i \in [n]$

one by one to semi-functional, thereby changing the distribution of the *public parameters* over the hybrid through the keys. Similar techniques are also found in the selectively secure broadcast encryption of [Wee16, CMM16a], which removed the use of q-type assumptions in [BGW05], using the Déjà Q paradigm introduced by [CM14].

Prime-Order Groups. The scheme we just described in two ways is based on composite-order asymmetric bilinear groups. We give the scheme in detail in Section 3 and its proof in Section 4. For efficiency reasons [Gui13], schemes based on prime-order groups are preferable in practice. As such, we additionally provide a translation of our composite-order scheme to the prime-order setting in Section 5.

Our construction uses a proof paradigm that can be seen as an optimization of known composite to prime-order translation frameworks, such as [Fre10, OT08, OT09, Lew12, CGW15, Att15, AC16]. Roughly speaking, in these frameworks, a random group element g_p^s of a composite order bilinear group G is emulated by a vector of group elements $[\mathbf{As}]_1$, where $\mathbf{s} \in \mathbb{Z}_p^k$, $\mathbf{A} \in \mathbb{Z}_p^{(k+1)\times k}$ is a k-Lin matrix, and we use the bracket notation $[a]_i$ to denote the element g_i^a for $i \in \{1, 2, T\}$ (for a prime order bilinear group $G_1 \times G_2 \to G_T$). Here, k depends on the k-Lin assumption used, i.e: $k = 1$ corresponds to the Symmetric External Diffie-Hellman Assumption, or SXDH. The decision assumption used to argue that $g_p^s \approx g_p^s g_q^s$ in composite order groups is replaced by the k-Lin assumption: $[\mathbf{As}]_1 \approx [\mathbf{u}]_1$, where $\mathbf{A} \in \mathbb{Z}_p^{(k+1)\times k}$ is a k-Lin matrix, $\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$, and $\mathbf{u} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$ is a uniformly random vector over $\mathbb{Z}_p^{k+1}$. Finally, each group element g^{w_i} of the public parameters is mapped to a $(k+1) \times (k+1)$ matrix of group elements.

Our constructions employ an optimization that uses public parameter matrices of size only $(k+1) \times k$, thereby reducing the public parameters and the ciphertext size by a factor of $k+1$ (see Figure 2). This is done by replacing the information theoretic argument at the heart of the dual system encryption methodology (used to switch secret keys to semi-functional secret keys) with a computational argument. Similar techniques are used in [CW14, BKP14, AC16].

In [CGW15]:	
$w_j \to \mathbf{W}_j \in \mathbb{Z}_p^{(k+1)\times(k+1)}$	
$s \to \mathbf{s} \in \mathbb{Z}_p^k$,	$r_i \to \mathbf{r}_i \in \mathbb{Z}_p^k$
$g_p^s \to [\mathbf{s}^\top \mathbf{A}^\top]_1$,	$h_p^{r_i} \to [\mathbf{B}\mathbf{r}_i]_2$
$g_p^{w_j s} \to [\mathbf{s}^\top \mathbf{A}^\top \mathbf{W}_j]_1$,	$h_p^{w_j r_i} \to [\mathbf{W}_j \mathbf{B}\mathbf{r}_i]_2$
In our work:	
$w_j \to \mathbf{W}_j \in \mathbb{Z}_p^{(k+1)\times k}$	
$s \to \mathbf{s} \in \mathbb{Z}_p^k$,	$r_i \to \mathbf{r}_i \in \mathbb{Z}_p^k$
$g_p^s \to [\mathbf{s}^\top \mathbf{A}^\top]_1$,	$h_p^{r_i} \to [\overline{\mathbf{B}}\mathbf{r}_i]_2$
$g_p^{w_j s} \to [\mathbf{s}^\top \mathbf{A}^\top \mathbf{W}_j]_1$,	$h_p^{w_j r_i} \to [\mathbf{W}_j \overline{\mathbf{B}}\mathbf{r}_i]_2$

Figure 2: $\mathbf{A}, \mathbf{B} \in \mathbb{Z}_p^{(k+1)\times k}$ are k-Lin matrices, $\overline{\mathbf{B}} \in \mathbb{Z}_p^{k\times k}$ denotes the k upper rows of $\mathbf{B}$.

Tight security proof in the multi-challenge setting. The security definition of public key encryption schemes typically involves a game where there is only one challenge ciphertext, since this implies security of the scheme when multiple challenge ciphertexts are allowed to be

requested via a standard hybrid argument. However, using such an argument incurs a security loss that is proportional to the number of challenge ciphertexts. This can be problematic since real-life attacks might be performed on many challenge ciphertexts. In particular, for widely deployed schemes, the number of challenge ciphertexts can be as large as 2^{20}, or even 2^{30}. A standard hybrid over the ciphertexts in the latter case results in an increase in the size of the security parameter by 30 compared to the setting where the adversary receives only one challenge ciphertext. For elliptic curve groups eligible to instantiate our scheme in which the SXDH assumption is believed to hold, such an increase would translate to a $2 \cdot 30 = 60$ bit increase in the size of each group element description. Thus, a tight security reduction allows for shorter group element descriptions and increased efficiency. Finally, note that the number of challenge ciphertexts can be unknown during the setup phase, which means that a conservative estimate could assume it to be high during security parameter calculation, thereby resulting in needlessly large group elements used in the scheme. Tight security reductions avoid this problem by allowing the security parameter to be set in a way that is independent of the number of challenge ciphertexts.

To obtain a tightly secure construction, we slightly modify the prime-order scheme mentioned above, so as to allow a different proof strategy. The modification does not incur any increase in the ciphertext size for the most efficient version of the scheme: when $k = 1$ and security holds under 1-Lin a.k.a. the SXDH assumption. In general, the ciphertext size in the tightly secure scheme increases by $k - 1$ group elements when security is based on k-Lin. In the tight-security proof, we simultaneously switch all of the challenge ciphertexts to semi-functional mode using the random self reducibility of the k-Lin assumption. Then, the high-level proof structure is similar to that of previous scheme: we perform a hybrid argument that switches each secret key one by one to a semi-functional version (note that the number of secret keys is upper bounded by n, so this hybrid argument only incurs a security loss that is proportional to n, the number of users). To switch the key sk_ℓ to semi-functional mode, we use entropy from the component $[\mathbf{W}_0\mathbf{r}_\ell]_2$ in the key sk_ℓ to obtain a new random semi-functional component (the component $\gamma_\ell\mathbf{a}^\perp$). Doing so requires analysis of the entropy of $\mathbf{W}_0$ leaked by the public key and the challenge ciphertext(s). When there is only one challenge ciphertext for some set of users Γ, the (non-tight) proof crucially relies on the fact that $\ell \notin \Gamma$ for the challenge Γ, as required by the security game definition and the fact that the adversary queried sk_ℓ. For the tight reduction, we have many challenges Γ_i, so we must deal with potentially more information about $\mathbf{W}_0$ leaked. In fact, this is not the case: the challenge ciphertexts for all sets Γ_i queried to EncO do not leak more information about $\mathbf{W}_0$ than a *single* ciphertext for the set $\bigcup_i \Gamma_i$, which would be an allowed challenge query given the same set of user keys. This allows us to reduce to the argument for the single-ciphertext case.

1.4 Discussion

Prior to this work, it wasn't clear what the bottleneck was in improving a broadcast encryption scheme with constant size secret keys and ciphertext overhead based on q-type assumptions to being based only on static assumptions. More specifically, one might ask: "What exactly is the use of q-type assumptions in [BGW05] buying us?" Our work clarifies that the main bottleneck is to get to linear-size public keys (and not constant-size secret keys or ciphertext overhead). Indeed, as noted earlier, if you replace the r_i, w_i in the composite-order scheme of Section 3 with powers of α ($r_i = \alpha^i, w_i = \alpha^{n-i+1}$), we can compress the public parameters to linear size, and essentially recover the construction of [BGW05]. That is, the role of the q-type assumption is to compress a quadratic number of terms to linear. This is very different from the use of q-type

assumptions in the HIBE of [BBG05], for example, which were replaced with static assumptions by [LW10] without a loss in asymptotic parameters.

2 Preliminaries

2.1 Notation

We denote by $x \leftarrow_{\mathrm{R}} X$ the fact that x is picked uniformly at random from a finite set X. By "PPT", we denote a probabilistic polynomial-time algorithm.

2.2 Bilinear Groups

We instantiate both broadcast encryption schemes using asymmetric bilinear groups. Let $\mathcal{G}$ be a probabilistic polynomial time (PPT) algorithm that on input a security parameter 1^λ returns an asymmetric bilinear group description $\mathbb{G} := (N, G_1, G_2, G_T, e)$, where G_1, G_2 and G_T are cyclic groups of order N, and $e : G_1 \times G_2 \to G_T$ is a non-degenerate bilinear map. We require that the group operations in G_1, G_2 and G_T as well as the bilinear map e are computable in deterministic polynomial time.

Composite-order groups. For the composite-order construction in Section 3, we consider groups of order $N = pq$, where p, q are distinct primes of $\Theta(\lambda)$ bits, and $G_1 = G, G_2 = H$ are asymmetric groups. In this setting, we can write $G = G_p G_q$ and $H = H_p H_q$, where G_p, G_q, H_p, H_q are subgroups of the subscripted order. In addition, we use G_s^*, H_s^* to denote $G_s \setminus \{1\}, H_s \setminus \{1\}$, where $s \in \{p, q\}$. We will often use write g_p, g_q, h_p, h_q to denote random generators for the subgroup G_p, G_q, H_p, H_q.

Prime-order groups. For the prime-order construction in Section 5, we consider groups of order $N = p$ for some prime p of $\Theta(\lambda)$ bits, where G_1 and G_2 are possibly different groups (type 1, 2 or 3 pairing). We write g_1, g_2 to denote random generators of G_1 and G_2 respectively, and $g_T := e(g_1, g_2)$, which is a generator of G_T. We use implicit representation of group elements: for $a \in \mathbb{Z}_p$, define $[a]_s = a g_s \in G_s$ as the implicit representation of a in G_s, for $s \in \{1, 2, T\}$. Given $[a]_1$ and $[b]_2$, one can efficiently compute $[ab]_T$ using the pairing e. For two matrices $\mathbf{A} \in \mathbb{Z}_p^{\ell \times m}$, $\mathbf{B} \in \mathbb{Z}_p^{m \times n}$, define $e([\mathbf{A}]_1, [\mathbf{B}]_2) := [\mathbf{AB}]_T \in G_T^{\ell \times m}$.

2.3 Static Composite-Order Assumptions

The security of the composite-order scheme in Section 3 is proven under three static assumptions in composite-order asymmetric bilinear groups. We define the advantage functions referred to in the assumptions in Figure 3.

Definition 2.1 (Composite-Order Static Decision Assumptions)**.** We say that the Static Decision Assumptions hold relative to $\mathcal{G}$ if for all PPT adversaries $\mathcal{A}$, the advantages $\mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}}(\lambda)$, $\mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}}(\lambda)$, and $\mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}}(\lambda)$ are negligible functions in λ.

2.4 Matrix Diffie-Hellman Assumptions

The security of the prime-order scheme in Section 5 is proven under the Matrix Decision Diffie-Hellman (MDDH) Assumption [EHK+13], whose definition we recall here.

Definition 2.2 (Matrix Distribution)**.** Let $k, \ell \in \mathbb{N}$, with $\ell > k$. We call $\mathcal{D}_{\ell,k}$ a matrix distribution if it outputs matrices in $\mathbb{Z}_p^{\ell \times k}$ of full rank k in polynomial time. We write $\mathcal{D}_k := \mathcal{D}_{k+1,k}$.

$$\mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(D,T_0)=1] - \Pr[\mathcal{A}(D,T_1)=1]|$$
where $\mathbb{G} \leftarrow \mathcal{G}(\lambda),\ D := (g_p, h_p),\ g_p \leftarrow_{\mathrm{R}} G_p^*,\ h_p \leftarrow_{\mathrm{R}} H_p^*$
and $T_0 := g_p^s \leftarrow_{\mathrm{R}} G_p,\ T_1 = g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q$

$$\mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(D,T_0)=1] - \Pr[\mathcal{A}(D,T_1)=1]|$$
where $\mathbb{G} \leftarrow \mathcal{G}(\lambda),\ D := (g_p, h_p, g_p^s g_q^{s'}, h_q^{\alpha'}),$
$g_p \leftarrow_{\mathrm{R}} G_p^*,\ h_p \leftarrow_{\mathrm{R}} H_p^*,\ g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q,\ h_q^{\alpha'} \leftarrow_{\mathrm{R}} H_q$
and $T_0 := h_p^z \leftarrow_{\mathrm{R}} H_p,\ T_1 = h_p^z h_q^{z'} \leftarrow_{\mathrm{R}} H_p H_q$

$$\mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(D,T_0)=1] - \Pr[\mathcal{A}(D,T_1)=1]|$$
where $\mathbb{G} \leftarrow \mathcal{G}(\lambda),\ D := (g_p, h_p, g_p^s g_q^{s'}, h_p^{\alpha} h_q^{\alpha'}),$
$g_p \leftarrow_{\mathrm{R}} G_p^*,\ h_p \leftarrow_{\mathrm{R}} H_p^*,\ g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q,\ h_p^{\alpha} h_q^{\alpha'} \leftarrow_{\mathrm{R}} H_p H_q$
and $T_0 := e(g_p, h_p)^{s\alpha},\ T_1 = X \leftarrow_{\mathrm{R}} G_T$

Figure 3: Advantage functions

Without loss of generality, we assume the first k rows of $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_{\ell,k}$ form an invertible matrix. The $\mathcal{D}_{\ell,k}$-Matrix Diffie-Hellman problem in G_s for $s \in \{1,2,T\}$ is to distinguish the two distributions $([\mathbf{A}]_s, [\mathbf{A}\mathbf{w}]_s)$ and $([\mathbf{A}]_s, [\mathbf{u}]_s)$ where $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_{\ell,k}$, $\mathbf{w} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$ and $\mathbf{u} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^\ell$.

Definition 2.3 ($\mathcal{D}_{\ell,k}$-Matrix Diffie-Hellman Assumption $\mathcal{D}_{\ell,k}$-MDDH)**.** Let $\mathcal{D}_{\ell,k}$ be a matrix distribution. We say that the $\mathcal{D}_{\ell,k}$-Matrix Diffie-Hellman ($\mathcal{D}_{\ell,k}$-MDDH) Assumption holds relative to $\mathcal{G}$ in G_s for $s \in \{1,2,T\}$ if for all PPT adversaries $\mathcal{A}$,

$$\mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{\ell,k},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(,[\mathbf{A}]_s,[\mathbf{A}\mathbf{w}]_s)=1] - \Pr[\mathcal{A}(,[\mathbf{A}]_s,[\mathbf{u}]_s)=1]| = \mathsf{negl}(\lambda),$$

where the probability is taken over $\leftarrow_{\mathrm{R}} \mathcal{G}(1^\lambda)$, $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_k, \mathbf{w} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k, \mathbf{u} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^\ell$.

For each $k \geq 1$, [EHK+13] specifies distributions $\mathcal{L}_k$, $\mathcal{SC}_k$, $\mathcal{C}_k$ (and others) over $\mathbb{Z}_p^{(k+1)\times k}$ such that the corresponding $\mathcal{D}_k$-MDDH assumptions are generically secure in bilinear groups and form a hierarchy of increasingly weaker assumptions. $\mathcal{L}_k$-MDDH is the well known k-Linear Assumption k-Lin with 1-Lin = DDH.

Definition 2.4 (Uniform distribution)**.** Let $\ell, k \in \mathbb{N}$, with $\ell > k$. We denote by $\mathcal{U}_{\ell,k}$ the uniform distribution over all full-rank $\ell \times k$ matrices over $\mathbb{Z}_p$. Let $\mathcal{U}_k := \mathcal{U}_{k+1,k}$.

Among all possible matrix distributions $\mathcal{D}_{\ell,k}$, the uniform matrix distribution $\mathcal{U}_k$ is the hardest possible instance, so in particular k-Lin $\Rightarrow \mathcal{U}_k$-MDDH.

Lemma 2.5 ($\mathcal{D}_{\ell,k}$-MDDH $\Rightarrow \mathcal{U}_k$-MDDH, [EHK+13])**.** *Let $\mathcal{D}_{\ell,k}$ be a matrix distribution. For any PPT adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and $\mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{\ell,k},\mathcal{A}}(\lambda) = \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{U}_k,\mathcal{B}}(\lambda)$.*

Let $Q \geq 1$. For $\mathbf{W} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k\times Q}, \mathbf{U} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{\ell\times Q}$, we consider the Q-fold $\mathcal{D}_{\ell,k}$-MDDH Assumption in G_s for $s \in \{1,2,T\}$ which consists in distinguishing the distributions $([\mathbf{A}]_s, [\mathbf{A}\mathbf{W}]_s)$ from $([\mathbf{A}]_s, [\mathbf{U}]_s)$. That is, a challenge for the Q-fold $\mathcal{D}_{\ell,k}$-MDDH Assumption consists of Q independent challenges of the $\mathcal{D}_{\ell,k}$-MDDH Assumption (with the same $\mathbf{A}$ but different randomness $\mathbf{w}$). In [EHK+13] it is shown that the two problems are equivalent, where (for $Q \geq \ell - k$) the reduction loses a factor $\ell - k$.

Lemma 2.6 (Random self-reducibility of $\mathcal{D}_{\ell,k}$-MDDH, [EHK+13]). *Let $\ell, k, Q \in \mathbb{N}$ with $\ell > k$. For any PPT adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A}) + Q \cdot \mathsf{poly}(\lambda)$ with $\mathsf{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$, and*

$$\mathrm{Adv}^{Q\text{-}\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{\ell,k},\mathcal{A}}(\lambda) \leq (\ell - k) \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{\ell,k},\mathcal{B}}(\lambda) + \frac{1}{p-1}$$

where $\mathrm{Adv}^{Q\text{-}\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{\ell,k},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(\mathbb{G}, [\mathbf{A}]_s, [\mathbf{AW}]_s) = 1] - \Pr[\mathcal{B}(\mathbb{G}, [\mathbf{A}]_s, [\mathbf{U}]_s) = 1]|$ and the probability is over $\mathbb{G} \leftarrow_R \mathcal{G}(1^\lambda)$, $\mathbf{A} \leftarrow_R \mathcal{D}_{\ell,k}, \mathbf{W} \leftarrow_R \mathbb{Z}_p^{k \times Q}, \mathbf{U} \leftarrow_R \mathbb{Z}_p^{\ell \times Q}$.

2.5 Broadcast encryption

A broadcast encryption scheme consists of three randomized algorithms $(\mathsf{Setup}, \mathsf{Enc}, \mathsf{Dec})$, along with a fourth deterministic procedure: KeyGen.

- $\mathsf{Setup}(1^\lambda, 1^n) \to (\mathsf{pk}, \mathsf{msk})$. The setup algorithm gets as input the security parameter 1^λ and the number of users 1^n. It outputs the public parameters pk and master secret key msk.
- $\mathsf{KeyGen}(\mathsf{msk}, i) \to \mathsf{sk}_i$. The key generation algorithm gets as input the master secret key msk and an index $i \in [n]$. It (deterministically) outputs the secret key for user i: sk_i.
- $\mathsf{Enc}(\mathsf{pk}, \Gamma, M) \to \mathsf{ct}_\Gamma$. The encryption algorithm gets as input pk and a subset $\Gamma \subseteq [n]$. It outputs a ciphertext ct_Γ. Here, Γ is public given ct_Γ.
- $\mathsf{Dec}(\mathsf{pk}, \mathsf{sk}_i, \mathsf{ct}_\Gamma) \to M$. The decryption algorithm gets as input $\mathsf{pk}, \mathsf{sk}_i$, and ct_Γ. It outputs a message M.

Correctness

We require that for all $\Gamma \subseteq [n]$, messages M, and $i \in [n]$ for which $i \in \Gamma$,

$$\Pr[\mathsf{ct}_\Gamma \leftarrow \mathsf{Enc}(\mathsf{pk}, \Gamma, M), \mathsf{sk}_i \leftarrow \mathsf{KeyGen}(\mathsf{msk}, i); \mathsf{Dec}(\mathsf{pk}, \mathsf{sk}_i, \mathsf{ct}_\Gamma) = M] = 1$$

where the probability is taken over $(\mathsf{pk}, \mathsf{msk}) \leftarrow \mathsf{Setup}(1^\lambda, 1^n)$ and the coins of Enc.

Security

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

$$\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda) := \left| \Pr_{(b,\mathsf{pk},\mathsf{msk}) \leftarrow \mathsf{SetupO}} \left[\; b' = b \;\middle|\; b' \leftarrow \mathcal{A}^{\mathsf{KeyGenO}(\cdot),\mathsf{EncO}(\cdot,\cdot)}(1^\lambda) \right] - 1/2 \right|$$

where:

- SetupO samples $(\mathsf{pk}, \mathsf{msk}) \leftarrow_{\mathrm{R}} \mathsf{Setup}(1^\lambda, 1^n)$ and $b \leftarrow_{\mathrm{R}} \{0,1\}$, and returns pk. SetupO is called once at the beginning of the game.
- $\mathsf{KeyGenO}(i \in [n])$ returns $\mathsf{KeyGen}(\mathsf{msk}, i)$.
- If M_0 and M_1 are two messages of equal length, and $\Gamma \subset [n]$, $\mathsf{EncO}(\Gamma, M_0, M_1)$ returns $\mathsf{Enc}(\mathsf{pk}, \Gamma, M_b)$.

with the restriction that for all queries $i \in [n]$ that $\mathcal{A}$ makes to $\mathsf{KeyGenO}(\cdot)$ and all queries $\Gamma \subset [n]$ to EncO satisfy $i \notin \Gamma$ (that is, sk_i does not decrypt ct_Γ).

Note that this definition allows the adversary to query EncO multiple times. We call this the *multi-challenge* setting and say that a broadcast encryption scheme is *adaptively secure in the multi-challenge setting* if for all PPT adversaries $\mathcal{A}$, $\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda)$ is a negligible function in λ.

If we only consider adversaries that query EncO once, we have the standard notion of adaptive security. Namely, we say that a broadcast encryption scheme is *adaptively secure* if for all PPT adversaries $\mathcal{A}$ that issue only one query to Enc, $\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda)$ is a negligible function in λ.

Note that a scheme being adaptively secure implies that it is also adaptively secure in the multi-challenge setting via a hybrid argument over the challenge ciphertexts. However, this incurs a security loss proportional to the number of challenge ciphertexts, In Section 7, we present a scheme with a *tight* reduction in the multi-challenge security proof that avoids this loss.

3 Composite-Order Construction

Figure 4 shows our composite order construction.

$\underline{\mathsf{Setup}(1^\lambda, 1^n)}$:
$\mathbb{G} \leftarrow_R \mathcal{G}(1^\lambda)$; $g_p \leftarrow_R G_p^*$, $h_p \leftarrow_R H_p^*$; $\alpha, u \leftarrow_R \mathbb{Z}_N$; $\{w_i, r_i \leftarrow_R \mathbb{Z}_N\}_{i\in[n]}$
Output $\mathsf{pk} = \left(g_p,\ g_p^u,\ \{g_p^{w_i}\}_{i\in[n]},\ \{h_p^{r_i}\}_{i\in[n]},\ \{h_p^{w_i r_j}\}_{i\neq j},\ e(g_p,h_p)^\alpha\right)$ and
$\mathsf{msk} = \left(h_p, \alpha, u, \{r_i\}_{\in[n]}\right)$.

$\underline{\mathsf{KeyGen}(\mathsf{msk}, i \in [n]))}$:
Output $\mathsf{sk}_i = h_p^{\alpha+ur_i} \in H_p$.

$\underline{\mathsf{Enc}(\mathsf{pk}, \Gamma \subset [n], M \in G_T)}$:
$s \leftarrow_R \mathbb{Z}_N$
$C_0 := g_p^s$; $C_1 := g_p^{s\left(u+\sum_{j\notin\Gamma} w_j\right)}$; $C_2 := e(g_p,h_p)^{\alpha s} \cdot M$
Output $\mathsf{ct}_\Gamma := (C_0, C_1, C_2) \in G_p^2 \times G_T$

$\underline{\mathsf{Dec}(\mathsf{ct}_\Gamma, \mathsf{sk}_i)}$:
Compute $D_0 = e(\underbrace{(g_p^s)^{-1}}_{=C_0^{-1}}, \underbrace{h_p^{\alpha+ur_i}}_{=\mathsf{sk}_i}) = e(g_p,h_p)^{-s\alpha+-sur_i}$
Compute $D_1 = e(\underbrace{g_p^{s\left(u+\sum_{j\notin\Gamma} w_j\right)}}_{=C_1}, \underbrace{h_p^{r_i}}_{\text{from pk}}) = e(g_p,h_p)^{sur_i+s\sum_{j\notin\Gamma} w_j r_i}$
Compute $D_2 = e(\underbrace{(g_p^s)^{-1}}_{=C_0^{-1}}, \underbrace{\prod_{j\notin\Gamma} h_p^{w_j r_i}}_{\text{from pk}}) = e(g_p,h_p)^{s\sum_{j\notin\Gamma} w_j r_i}$
Compute and output $M = C_2 \cdot D_0 \cdot D_1 \cdot D_2$.

Figure 4: $\mathsf{BE}_{\mathsf{composite}}$, an adaptively secure broadcast encryption scheme based on composite-order bilinear groups.

4 Security Proof of the Composite-Order Construction

4.1 Hybrid definitions

Our proof will be accomplished through a standard "dual system" series of hybrids over the challenge ciphertext and n key indices, beginning with Game_{real}, the real security game, and ending at Game_{final}, a game in which the adversary has no advantage.

These games will differ in the distribution of the challenge ciphertext, secret keys, and public parameters. We define new semifunctional distributions of ciphertexts and secret keys in Figures 5 and 6.

Semi-functional Ciphertext. A semi-functional ciphertext is formed as follows:

Start with a normal $\mathsf{ct}_\Gamma = M \cdot e(g_p, h_p)^{\alpha s}, \;\; g_p^s, \;\; g_p^{s\left(u + \sum_{j \notin \Gamma} w_j\right)}$

Pick $s', u', w'_1, ..., w'_n \leftarrow_{\mathrm{R}} \mathbb{Z}_N$

Output $\mathsf{ct}'_\Gamma = M \cdot e(g_p, h_p)^{\alpha s}, \;\; g_p^s g_q^{s'}, \;\; g_p^{s\left(u + \sum_{j \notin \Gamma} w_j\right)} g_q^{s'\left(u' + \sum_{j \notin \Gamma} w'_j\right)}$

Figure 5: Semi-functional Ciphertext

To form semi-functional (keys, public parameters) for index t, first the normal public parameter and key generation procedures are performed to get:
$\mathsf{pk} := (g_p, \; g_p^u, \{g_p^{w_i}\}_{i \in [n]}, \; \{h_p^{r_i}\}_{i \in [n]}, \; \{h_p^{w_i r_j}\}_{i \neq j}, \; e(g_p, h_p)^\alpha)$
$\mathsf{sk}_t := h_p^{\alpha + u y_t}$
Draw $\alpha', w'_1, .., w'_n, r'_1, ..., r'_n \leftarrow_{\mathrm{R}} \mathbb{Z}_N$. The remaining steps depend on the particular type of semi-functional key / public parameters:

Type $(t, 1)$ Semi-functional keys:
Semi-functional keys of type $(t, 1)$ are formed as follows:
$\mathsf{sk}_{t,1} := h_p^{\alpha + u r_t} h_q^{u' r'_t}$

Type $(t, 2)$ Semi-functional keys:
Semifunctional keys of type $(t, 2)$ are formed as follows:
$\mathsf{sk}_{t,2} := h_p^{\alpha + u r_t} h_q^{\alpha' + u' r'_t}$

Type $(t, 3)$ Semi-functional keys:
A semi-functional key of type $(t, 3)$ is formed as follows:
$\mathsf{sk}_{t,3} := h_p^{\alpha + u r_t} h_q^{\alpha'}$

Type t public parameters:
Semi-functional public parameters of type t are formed as follows:
$\mathsf{pk}_t := g_p, \; g_p^u, \; \{g_p^{w_i}\}_{i \in [n]}, \; \{h_p^{r_i}\}_{i \neq t \in [n]} \cup \{h_p^{r_t} h_q^{r'_t}\}, \; \{h_p^{w_i r_j}\}_{\substack{i \neq j \\ j \neq t}} \cup \{h_p^{w_i r_t} h_q^{w'_i r'_t}\}_{i \neq t},$
$e(g_p, h_p)^\alpha$

Figure 6: Semi-functional Keys

Game_0
Same as Game_{real}, but challenge ciphertext is semi-functional.

$\mathsf{Game}_{\ell,1}$
Same as $\mathsf{Game}_{(\ell-1),3}$, but the public parameters are semi-functional of type ℓ and the key for index ℓ is semi-functional of type $(\ell, 1)$.

$\mathsf{Game}_{\ell,2}$
Same as $\mathsf{Game}_{\ell,1}$, but the key for index ℓ is semi-functional of type $(\ell, 2)$ (public parameters remain semi-functional of type ℓ).

$\mathsf{Game}_{\ell,3}$
Same as $\mathsf{Game}_{\ell,2}$, but the public parameters are normally formed and the key for index ℓ is semi-functional of type $(\ell, 3)$.

Game_{final}
Same as $\mathsf{Game}_{n,3}$, except the message M_b in the challenge ciphertext is blinded by an independently random group element $X \leftarrow_{\mathrm{R}} G_T$ instead of $e(g_p, h_p)^{\alpha s}$.

Figure 7: Hybrid Games

We use these distributions in Figure 7 to define the following hybrid games, where ℓ ranges from 1 to n.

Note that Game_0 is identical to $\mathsf{Game}_{0,3}$.

We first argue that no adversary can achieve non-negligible difference in advantage between Game_{real} and $\mathsf{Game}_0 \equiv \mathsf{Game}_{0,3}$. We then hybrid over each key index, arguing that no adversary can achieve non-negligible difference in advantage between $\mathsf{Game}_{(\ell-1),3}$ and $\mathsf{Game}_{\ell,1}$, then $\mathsf{Game}_{\ell,1}$ and $\mathsf{Game}_{\ell,2}$, then $\mathsf{Game}_{\ell,2}$ and $\mathsf{Game}_{\ell,3}$ for $\ell = 1, ...n$ until arriving at $\mathsf{Game}_{n,3}$, then finally Game_{final}, at which the adversary has no non-negligible advantage. Namely, we show that

$$\begin{aligned}\mathsf{Game}_{\mathrm{real}} \approx_c \mathsf{Game}_0 \equiv \mathsf{Game}_{0,3} \approx_c \mathsf{Game}_{1,0} \approx_c \mathsf{Game}_{1,1} \equiv \mathsf{Game}_{1,2} \quad &\approx_c \mathsf{Game}_{1,3} \\ \approx_c \mathsf{Game}_{2,0} \qquad\qquad \ldots\ldots \quad &\approx_c \mathsf{Game}_{n,3} \approx_c \mathsf{Game}_{final}\end{aligned}$$

where $\equiv$ denotes statistical equality, and $\approx_c$ denotes computational indistinguishability.

Figure 8 details how the constructions change throughout these games.

Notice that in the hybrid over key requests, all semi-functional keys before the hybrid index t are unable to decrypt a semi-functional ciphertext, even if they were in the authorized set. The key for index t becomes unable to do the same starting in $\mathsf{Game}_{t,2}$.

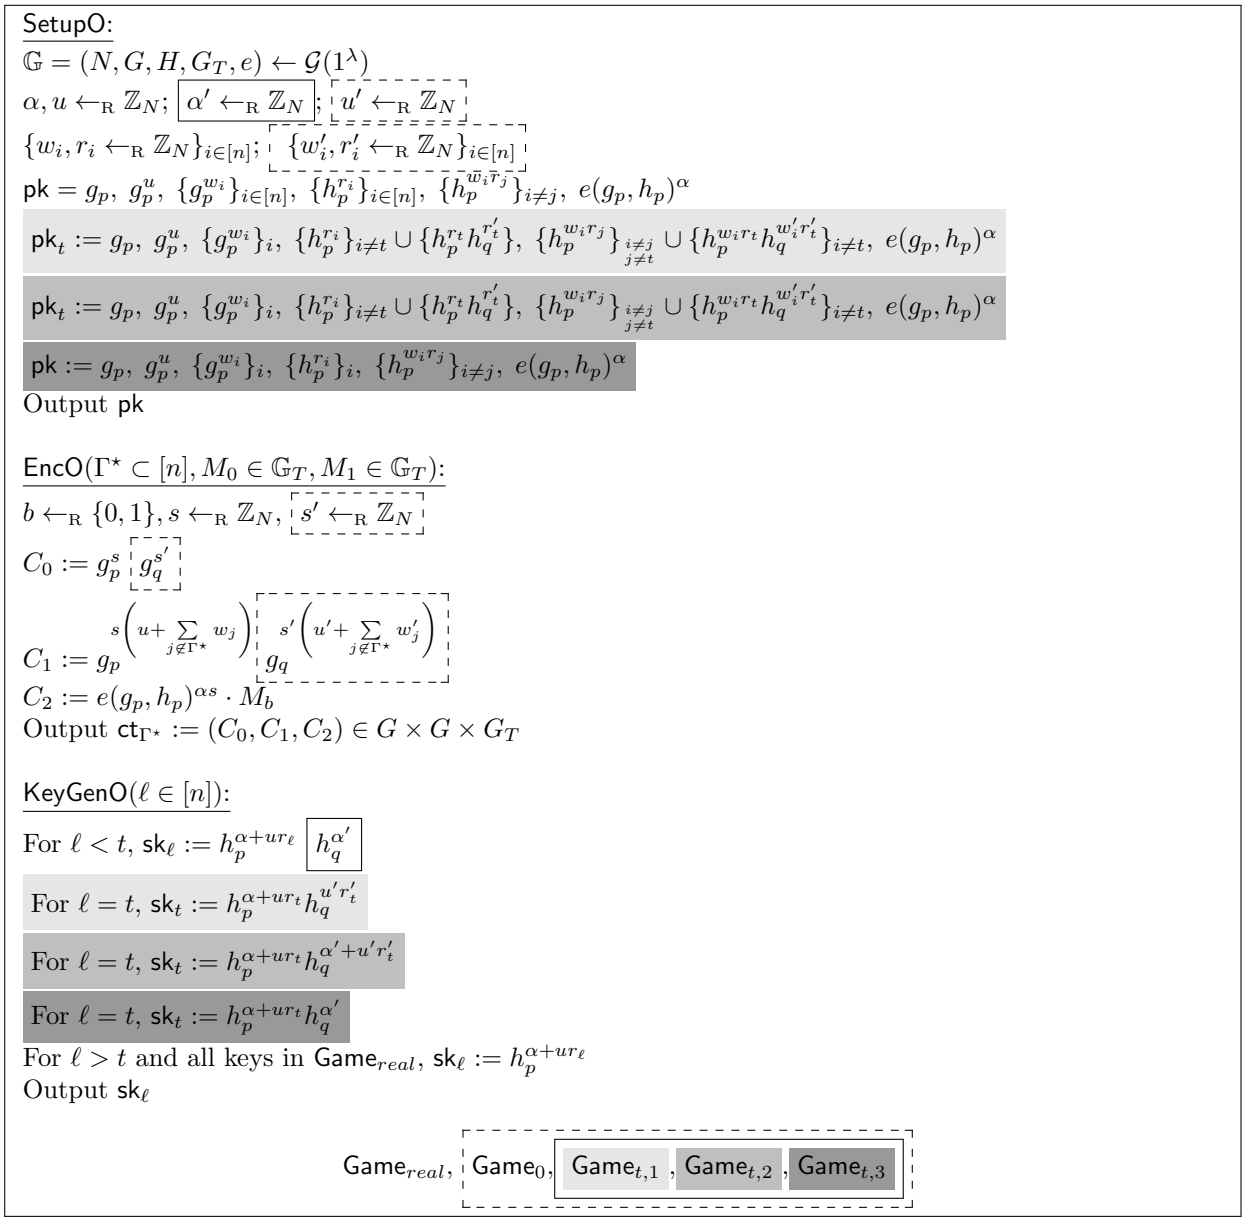

Figure 8: $\mathsf{Game}_{real}, \mathsf{Game}_0, \mathsf{Game}_{t,1}, \mathsf{Game}_{t,2}$(for $1 \le t \le n$), $\mathsf{Game}_{t,3}$ (for $0 \le t \le n$) for the proof of security of $\mathsf{BE}_{\mathsf{composite}}$ defined in Figure 4. In each procedure, the components inside a solid (dotted, light gray, gray) frame are only present in the games marked by a solid (dotted, light gray, gray) frame.

4.2 Hybrid Indistinguishability

We will show that any PPT adversary $\mathcal{A}$'s advantage in the real game, $\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda) = \mathrm{Adv}_{real}$, satisfies the following:

$$\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda) = \mathrm{Adv}_{real} \leq \mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}_1}(\lambda) + n \cdot \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_2}(\lambda) + n \cdot \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_3}(\lambda) + \mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}_4}(\lambda)$$

for adversaries $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4$ whose running times are essentially the same as $\mathcal{A}$'s.

We accomplish this in the following lemmas. Let Adv_i denote the adversary's advantage in Game_i. Then:

Lemma 4.1. $\mathrm{Adv}_{real} - \mathrm{Adv}_0 = \mathrm{Adv}_{real} - \mathrm{Adv}_{0,3} \leq \mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}_1}(\lambda)$

Proof. Given $g_p \leftarrow_{\mathrm{R}} G_p^*$, $h_p \leftarrow_{\mathrm{R}} H_p^*$, and $T = g_p^s \leftarrow_{\mathrm{R}} G_p$ or $g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q$, an adversary $\mathcal{A}_1$ could simulate the security game with $\mathcal{A}$ by running Setup and using KeyGen to respond to all key requests as usual with g_p, h_p.

When the challenge ciphertext is requested for set $\Gamma^\star$, form it as follows:

$$\mathsf{ct}_{\Gamma^\star} = M_b \cdot e(T, h_p)^\alpha, \;\; T, \;\; T^{u + \sum_{j \notin \Gamma^\star} w_j}$$

Notice that when $T = g_p^s$, then this is the same distribution as Game_{real}.

When $T = g_p^s g_q^{s'}$, then this is the same distribution as $\mathsf{Game}_0 = \mathsf{Game}_{0,3}$

(due to the Chinese Remainder theorem, $g_q^{u + \sum_{j \notin \Gamma^\star} w_j}$ is distributed identically to $g_q^{u' + \sum_{j \notin \Gamma^\star} w'_j}$ where u', w'_j are chosen independently at random from $\mathbb{Z}_N$).

It follows that a difference in advantage $\mathrm{Adv}_{real} - \mathrm{Adv}_{0,3}$ of $\mathcal{A}$ could be used by $\mathcal{A}_1$ to achieve the same advantage in the Static Decision Problem 1, so

$$\mathrm{Adv}_{real} - \mathrm{Adv}_{0,3} \leq \mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}_1}(\lambda)$$

□

Lemma 4.2. $\mathrm{Adv}_{(t-1),3} - \mathrm{Adv}_{t,1} \leq \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_2}(\lambda)$ *for* $t = 1, ..., n$.

Proof. Given $g_p \leftarrow_{\mathrm{R}} G_p^*$, $h_p \leftarrow_{\mathrm{R}} H_p^*$, $g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q$, $h_q^{\alpha'} \leftarrow_{\mathrm{R}} H_q$, and $T = h_p^z \leftarrow_{\mathrm{R}} H_p$ or $h_p^z h_q^{z'} \leftarrow_{\mathrm{R}} H_p H_q$, an adversary $\mathcal{A}_2$ could simulate the security game with $\mathcal{A}$ by first forming the public parameters as follows:

$$\alpha, u, w_1, ..., w_n, r_1, ..., r_{t-1}, r_{t+1}, ..., r_n \leftarrow_{\mathrm{R}} \mathbb{Z}_N$$

Output:

$$\mathsf{pk} = g_p, \; g_p^u, \; \{g_p^{w_i}\}_{i \in [n]}, \; \{h_p^{r_i}\}_{i \neq t} \cup \{T\}, \; \{h_p^{w_i r_j}\}_{i \neq j, j \neq t} \cup \{T^{w_i}\}_{i \neq t}, \; e(g_p, h_p)^\alpha$$

To form the (semi-functional) challenge ciphertext for set $\Gamma^\star$, compute:

$$\mathsf{ct}_{\Gamma^\star} = M_b \cdot e(g_p^s g_q^{s'}, h_p)^\alpha, \;\; (g_p^s g_q^{s'}), \;\; (g_p^s g_q^{s'})^{u + \sum_{j \notin \Gamma^\star} w_j}$$

(recall that due to the Chinese Remainder Theorem, the $g_p^u g_q^u$, $g_p^{w_i} g_q^{w_i}$ are distributed identically to $g_p^u g_q^{u'}$, $g_p^{w_i} g_q^{w_i'}$ for independently chosen u', w_i')

To form (semi-functional of type $(\ell, 3)$) keys for indices ℓ less than t, compute:

$$\mathsf{sk}'_{\ell,3} = h_p^{\alpha + u r_\ell}(h_q^{\alpha'})$$

Notice that (normal) keys for indices greater than t can also be computed, since α, u, h_p and r_ℓ are known (for all $\ell \neq t$).

For a key request for index t, compute:

$$\mathsf{sk}_t = h_p^{\alpha}(T)^u$$

Notice that when $T = h_p^z$, then this is the same distribution as $\mathsf{Game}_{(t-1),3}$ (the sk for index t and the public parameters are distributed normally), where $r_t = z$.

When $T = h_p^z h_q^{z'}$, then this is the same distribution as $\mathsf{Game}_{t,1}$ (the sk for index t is semi-functional of type $(t, 1)$ and the public parameters are semi-functional of type t), where $r_t = z$, and $r'_t = z'$.

It follows that a difference in advantage $\mathrm{Adv}_{(t-1),3} - \mathrm{Adv}_{t,1}$ of $\mathcal{A}$ for any $t = 1, ..., n$ could be used by $\mathcal{A}_2$ to achieve the same advantage in the Static Decision Problem 2 , so

$$\mathrm{Adv}_{(t-1),3} - \mathrm{Adv}_{t,1} \leq \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_2}(\lambda) \text{ for } t = 1, ..., n$$

□

Lemma 4.3. $\mathrm{Adv}_{t,1} - \mathrm{Adv}_{t,2} = 0$ *for* $t = 1, ..., n$.

Proof. The distributions of $\mathsf{Game}_{t,1}$ and $\mathsf{Game}_{t,2}$ are actually identical. To see this, note that the only difference between $\mathsf{Game}_{t,1}$ and $\mathsf{Game}_{t,2}$ is that the h_q component of the secret key for index t goes from $h_q^{r'_t u'}$ to $h_q^{\alpha' + r'_t u'}$.

If index t is not queried, then there is obviously no difference in the distribution between games.

Otherwise, if a key for index t is queried, notice that the only place w'_t occurs is in the g_q component of the challenge ciphertext: $g_q^{s'(u' + \sum_{j \notin \Gamma^\star} w'_j)}$ (we know that w'_t occurs in the sum because this key request must be for an index t not in the authorized set $\Gamma^\star$). Therefore, this w'_t in the summation is enough to information-theoretically hide the value of u' given just the challenge ciphertext. The only other place u' occurs is in the h_q component of the secret key for index t: $h_q^{r'_t u'}$. So, u' is enough to make the distribution of the $h_q^{r'_t u'}$ uniformly random (identical to $h_q^{\alpha' + r'_t u'}$ for an independent random α').

Either way, the two distributions are identical, and therefore $\mathrm{Adv}_{t,1} - \mathrm{Adv}_{t,2} = 0$ for $t = 1, ..., n$. □

Lemma 4.4. $\mathrm{Adv}_{t,2} - \mathrm{Adv}_{t,3} \leq \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_3}(\lambda)$ *for* $t = 1, ..., n$.

Proof. Given $g_p \leftarrow_{\mathrm{R}} G_p^*$, $h_p \leftarrow_{\mathrm{R}} H_p^*$, $g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q$, $h_q^{\alpha'} \leftarrow_{\mathrm{R}} H_q$, and $T = h_p^z \leftarrow_{\mathrm{R}} H_p$ or $h_p^z h_q^{z'} \leftarrow_{\mathrm{R}} H_p H_q$, an adversary $\mathcal{A}_3$ could simulate the security game with $\mathcal{A}$ by first forming the public parameters as follows:

$$\alpha, u, w_1, ..., w_n, r_1, ..., r_{t-1}, r_{t+1}, ..., r_n \leftarrow_{\mathrm{R}} \mathbb{Z}_N$$

Output:

$$\mathsf{pk} = g_p,\ g_p^u,\ \{g_p^{w_i}\}_{i\in[n]},\ \{h_p^{r_i}\}_{i\neq t} \cup \{T\},\ \{h_p^{w_i r_j}\}_{i\neq j, j\neq t} \cup \{T^{w_i}\}_{i\neq t},\ e(g_p, h_p)^\alpha$$

To form the (semi-functional) challenge ciphertext for set $\Gamma^\star$, compute:

$$\mathsf{ct}_{\Gamma^\star} = M_b \cdot e(g_p^s g_q^{s'}, h_p)^\alpha,\ \ (g_p^s g_q^{s'}),\ \ (g_p^s g_q^{s'})^{u + \sum_{j\notin\Gamma^\star} w_j}$$

(recall that due to the Chinese Remainder Theorem, the $g_p^u g_q^u$, $g_p^{w_i} g_q^{w_i}$ are distributed identically to $g_p^u g_q^{u'}$, $g_p^{w_i} g_q^{w_i'}$ for independently chosen u', w_i')

To form (semi-functional of type $(\ell, 3)$) keys for indices less than t, compute:

$$\mathsf{sk}'_{\ell,3} = h_p^{\alpha + u r_\ell}(h_q^{\alpha'})$$

Notice that (normal) keys for indices greater than t can also be computed, since α, u, h_p, and r_ℓ are known (for all $\ell \neq t$).

For a key request for index t compute:

$$\mathsf{sk}_t = h_p^\alpha (T)^u (g_q^{\alpha'})$$

Notice that when $T = h_p^z h_q^{z'}$, then this is the same distribution as $\mathsf{Game}_{t,2}$ (The sk for index t is semi-functional of type $(t, 2)$ and the public parameters are semi-functional of type t), where $r_t = z$, and $r'_t = z'$.

When $T = h_p^z$, then this is the same distribution as $\mathsf{Game}_{t,3)}$ (the public parameters are distributed normally and the tth sk is semi-functional of type $(t, 3)$), where $r_t = z$.

It follows that a difference in advantage $\mathrm{Adv}_{t,2} - \mathrm{Adv}_{t,3}$ of $\mathcal{A}$ for any $t = 1, ..., n$ could be used by $\mathcal{A}_3$ to achieve the same advantage in the Static Decision Problem 2, so

$$\mathrm{Adv}_{t,2} - \mathrm{Adv}_{t,3} \leq \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_3}(\lambda) \text{ for } t = 1, ..., n.$$

□

The preceding three lemmas take us all the way up to $\mathsf{Game}_{n,3}$, where the public parameters are normally formed, the challenge ciphertext is semi-functional, and all keys are semi-functional of type $(n, 3)$. We argue that any difference in advantage of $\mathcal{A}$ between this game and Game_{final}, which is the same game except the message M_b is blinded by a independently random target group element, can be used to achieve the same advantage in the Static Decision Problem 3:

Lemma 4.5. $\mathrm{Adv}_{n,3} - \mathrm{Adv}_{final} \leq \mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}_4}(\lambda)$

Proof. Given $g_p \leftarrow_{\mathrm{R}} G_p^*$, $h_p \leftarrow_{\mathrm{R}} H_p^*$, $g_p^s g_q^{s'} \leftarrow_{\mathrm{R}} G_p G_q$, $h_p^\alpha h_q^{\alpha'} \leftarrow_{\mathrm{R}} H_p H_q$, and $T = e(g_p, h_p)^{\alpha s}$ or $X \leftarrow_{\mathrm{R}} G_T$, an adversary $\mathcal{A}_4$ could simulate the security game with $\mathcal{A}$ by forming the public parameters as follows:

$$u, w_1, ..., w_n, r_1, ..., r_n \leftarrow_{\mathrm{R}} \mathbb{Z}_N$$

$$\mathsf{pk} = g_p,\ g_p^u,\ \{g_p^{w_i}\}_{i\in[n]},\ \{h_p^{r_i}\}_{i\in[n]},\ \{h_p^{w_i r_j}\}_{i\neq j},\ e(g_p, h_p^\alpha h_q^{\alpha'})$$

To form the (semi-functional) challenge ciphertext for set $\Gamma^\star$, compute:

$$\mathsf{ct}_{\Gamma^\star} = M_b \cdot T, \;\; (g_p^s g_q^{s'}), \;\; (g_p^s g_q^{s'})^{u + \sum_{j \notin \Gamma^\star} w_j}$$

To form (semi-functional of type $(\ell, 3)$) keys, compute:

$$\mathsf{sk}_{\ell,3} = (h_p^\alpha h_q^{\alpha'}) h_p^{u r_\ell}$$

for any $\ell \in [n]$ requested.

Note that if $T = e(g_p, h_p)^{\alpha s}$, then this game is distributed exactly as in $\mathsf{Game}_{n,3}$.

If $T = X$ for a uniformly random X, then we are in Game_{final}.

It follows that a difference in advantage $\mathrm{Adv}_{n,3} - \mathrm{Adv}_{final}$ of $\mathcal{A}$ could be used by $\mathcal{A}_4$ to achieve the same advantage in the Static Decision Problem 3, so $\mathrm{Adv}_{n,3} - \mathrm{Adv}_{final} \leq \mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}_4}(\lambda)$. □

Theorem 4.6. *If the Static Decision Assumptions of Definition 2.1 hold, then the broadcast encryption scheme* $\mathsf{BE}_{\mathsf{composite}}$ *defined in Figure 4 is adaptively secure.*

Proof. Summing the statements of the previous lemmas gives us:

$$\mathrm{Adv}_{real} - \mathrm{Adv}_{final} \leq \mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}_1}(\lambda) + n \cdot \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_2}(\lambda) + n \cdot \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_3}(\lambda) + \mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}_4}(\lambda)$$

In Game_{final} the challenge message M_b is information theoretically hidden by X, so it is obvious that no PPT adversary can achieve nonzero advantage in this game (that is, $\mathrm{Adv}_{final} = 0$). So, we have:

$$\mathrm{Adv}_{real} \leq \mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}_1}(\lambda) + n \cdot \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_2}(\lambda) + n \cdot \mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_3}(\lambda) + \mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}_4}(\lambda)$$

Our static decision assumptions state that $\mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}_1}(\lambda)$, $\mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_2}(\lambda)$, $\mathrm{Adv}^{SD2}_{\mathcal{G},\mathcal{A}_3}(\lambda)$, $\mathrm{Adv}^{SD3}_{\mathcal{G},\mathcal{A}_4}(\lambda)$ are negligible functions of λ (and n is a polynomial function of λ), so the advantage $\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda) = \mathrm{Adv}_{real}$ is a negligible function of λ, and therefore our scheme is adaptively secure. □

5 Prime Order Construction

Our prime-order construction is detailed in Figure 9.

$\underline{\mathsf{Setup}(1^\lambda, 1^n)}$:
$\mathbb{G} \leftarrow_{\mathrm{R}} \mathcal{G}(1^\lambda)$; $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_k$; $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$; $\{\mathbf{W}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}, \mathbf{r}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k\}_{i\in[n]}$
Output $\mathsf{pk} := \Big([\mathbf{A}]_1, [\mathbf{A}^\top \mathbf{W}_0]_1 \{[\mathbf{A}^\top \mathbf{W}_i]_1, [\mathbf{r}_i]_2\}_{i\in[n]}, [\mathbf{A}^\top \mathbf{k}]_T, \{[\mathbf{W}_j \mathbf{r}_i]_2\}_{i,j\in[n], i\neq j}\Big)$ and
$\mathsf{msk} := \big([\mathbf{k}]_2, \{[\mathbf{W}_0 \mathbf{r}_i]_2\}_{i\in[n]}\big)$.

$\underline{\mathsf{KeyGen}(\mathsf{msk}, i \in [n])}$:
Output $\mathsf{sk}_i := [\mathbf{k} + \mathbf{W}_0 \mathbf{r}_i]_2 \in G_2^{(k+1)}$.

$\underline{\mathsf{Enc}(\mathsf{pk}, \Gamma \subset [n], M \in G_T)}$:
$\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$
$C_0 := [\mathbf{s}^\top \mathbf{A}^\top]_1$; $C_1 := [\mathbf{s}^\top \mathbf{A}^\top (\mathbf{W}_0 + \sum_{j\notin\Gamma^\star} \mathbf{W}_j)]_1$; $C_2 := [\mathbf{s}^\top \mathbf{A}^\top \mathbf{k}]_T \cdot M$
Output $\mathsf{ct}_\Gamma := (C_0, C_1, C_2) \in G_1^{2k+1} \times G_T$

$\underline{\mathsf{Dec}(\mathsf{ct}_\Gamma, \mathsf{sk}_i)}$: // ct_Γ and sk_i implictly contain a description of Γ and i
Compute $e(\underbrace{[\mathbf{s}^\top \mathbf{A}^\top]_1}_{=C_0}, \underbrace{[\mathbf{k} + \mathbf{W}_0 \mathbf{r}_i]_2}_{=\mathsf{sk}_i}) / e(\underbrace{[\mathbf{s}^\top \mathbf{A}^\top (\mathbf{W}_0 + \sum_{j\notin\Gamma} \mathbf{W}_j)]_1}_{=C_1}, \underbrace{[\mathbf{r}_i]_1}_{\in\mathsf{pk}}) = [\mathbf{s}^\top \mathbf{A}^\top \mathbf{k} - \mathbf{s}^\top \mathbf{A}^\top \textstyle\sum_{j\notin\Gamma} \mathbf{W}_j \mathbf{r}_i]_T$.
Multiply the previous term by $e(\underbrace{[\mathbf{s}^\top \mathbf{A}^\top]_1}_{=C_0}, \underbrace{[\sum_{j\notin\Gamma} \mathbf{W}_j \mathbf{r}_i]_2}_{\in\mathsf{pk} \text{ for } i\in\Gamma})$ to get the encapsulation key $[\mathbf{s}^\top \mathbf{A}^\top \mathbf{k}]_T$ and to output the message M.

Figure 9: $\mathsf{BE}_{\mathsf{prime}}$, an adaptively secure broadcast encryption scheme based on prime-order bilinear groups.

6 Security Proof of the Prime-Order Construction

We now give the security proof of the scheme $\mathsf{BE}_{\mathsf{prime}}$, presented in Figure 9.

Theorem 6.1. *If the $\mathcal{D}_k$-MDDH Assumption holds in G_1 and G_2, then the broadcast encryption scheme $\mathsf{BE}_{\mathsf{prime}}$ defined in Figure 9 is adaptively secure (as defined in section 2.5). Namely, for any adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and*

$$\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A}}(\lambda) \leq (2n+1) \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_k,\mathcal{B}}(\lambda) + 2^{-\Omega(\lambda)},$$

where n is the number of users.

We prove Theorem 6.1 via a series of games described in Figure 10 and we use Adv_i to denote the advantage of $\mathcal{A}$ in game Game_i. Namely, we show that: $\mathsf{Game}_{\mathrm{real}} \approx_c \mathsf{Game}_0 \approx_c \mathsf{Game}_1 \approx_c \ldots \approx_c \mathsf{Game}_n$, where $\approx_c$ denotes computational indistinguishability. $\mathsf{Game}_{\mathrm{real}}$ is the security game as defined in Section 2.5, and the other Game_i are defined in Figure 10. Theorem 6.1 follows from Lemma 6.2, 6.3 and 6.4 below.

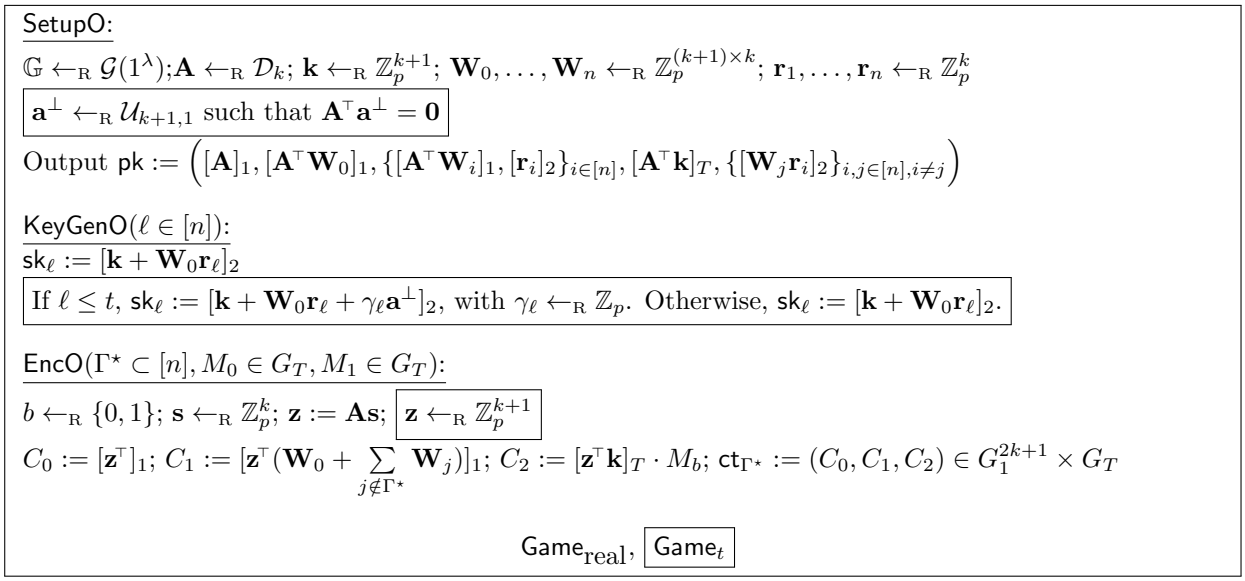

Figure 10: $\mathsf{Game}_{\mathrm{real}}, \mathsf{Game}_t$ (for $0 \leq t \leq n$) for the proof of security of $\mathsf{BE}_{\mathsf{prime}}$ defined in Figure 9. Here n denotes the number of users. In each procedure, the components inside a solid frame are only present in the games marked by a solid frame.

Lemma 6.2 ($\mathsf{Game}_{\mathrm{real}} \approx_c \mathsf{Game}_0$)**.** *There exists an adversary $\mathcal{B}_0$ such that $\mathbf{T}(\mathcal{B}_0) \approx \mathbf{T}(\mathcal{A})$ and*

$$|\mathrm{Adv}_{real} - \mathrm{Adv}_0| \leq \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_k,\mathcal{B}_0}(\lambda).$$

Here, we use the MDDH assumption to switch the distribution of the challenge ciphertext.

Proof. To go from $\mathsf{Game}_{\mathrm{real}}$ to Game_0, we switch the distribution of the vector $[\mathbf{z}]_1$ in the challenge ciphertext, using the $\mathcal{D}_k$-MDDH Assumption on $[\mathbf{A}]_1$ (see Definition 2.3).

Upon receiving a challenge $(\mathbb{G}, [\mathbf{A}]_1, [\mathbf{v}]_1)$ for the $\mathcal{D}_k$-MDDH Assumption, $\mathcal{B}_0$ picks $b \leftarrow_{\mathrm{R}} \{0,1\}$, $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$; $\mathbf{W}_0, \ldots, \mathbf{W}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}$; $\mathbf{r}_1, \ldots, \mathbf{r}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$, sets $[\mathbf{z}]_1 := [\mathbf{v}]_1$, and simulates the public parameters, the secret keys and the challenge ciphertext as defined in Figure 10. Note that when $[\mathbf{v}]_1$ is a proper MDDH sample, $\mathcal{B}_0$ simulates $\mathsf{Game}_{\mathrm{real}}$, and when $[\mathbf{v}]_1$ is uniformly random over G_1^{k+1}, it simulates Game_0. □

Lemma 6.3 ($\mathsf{Game}_{t-1} \approx_c \mathsf{Game}_t$)**.** *For all $t \in [n]$, there exists an adversary $\mathcal{B}_{t-1}$ such that $\mathbf{T}(\mathcal{B}_{t-1}) \approx \mathbf{T}(\mathcal{A})$ and*

$$|\mathrm{Adv}_{t-1} - \mathrm{Adv}_t| \leq 2 \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_k,\mathcal{B}_{t-1}}(\lambda).$$

Here, we embed an MDDH challenge in pk and sk_t. More precisely, the simulator sets $\mathbf{r}_t := \overline{\mathbf{B}}\mathbf{v}_t \in \mathbb{Z}_p^k$, where $\mathbf{B} \leftarrow_{\mathrm{R}} \mathcal{D}_k$ and $\mathbf{v}_t \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$, i.e. the upper part of an MDDH challenge. The lower part $\underline{\mathbf{B}}\mathbf{v}_t \in \mathbb{Z}_p$ is embedded in sk_t, if sk_t is queried by the adversary (it may not be the case, in particular if $t \in \Gamma^\star$). Note that the simulator needs to know if sk_t is going to be queried by the adversary when simulating pk.

Proof. Upon receiving a challenge $(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{v}]_2)$ for the $\mathcal{D}_k$-MDDH Assumption, $\mathcal{B}_{t-1}$ simulates $\mathcal{A}$'s view as follows.

- SetupO:
  $\mathcal{B}_{t-1}$ guesses if $\mathcal{A}$ is going to query sk_t (by picking a random $\beta \leftarrow_{\mathrm{R}} \{0,1\}$). If so ($\beta = 1$), $\mathbf{W}_0$ and $\mathbf{W}_t$ are implicitly defined as $\mathbf{W}_0 := \widehat{\mathbf{W}}_0 - \mathbf{a}^\perp \mathbf{T_B}$ and $\mathbf{W}_t := \widehat{\mathbf{W}}_t + \mathbf{a}^\perp \mathbf{T_B}$, where $\widehat{\mathbf{W}}_0, \widehat{\mathbf{W}}_t \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}$, and $\mathbf{T_B} := \underline{\mathbf{B}}\overline{\mathbf{B}}^{-1} \in \mathbb{Z}_p^{1\times k}$ (recall that wlog., $\overline{\mathbf{B}}$ is an invertible matrix). Otherwise ($\beta = 0$), they are defined as $\mathbf{W}_0 := \widehat{\mathbf{W}}_0 \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}$ and $\mathbf{W}_t := \widehat{\mathbf{W}}_t \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}$.

  Then, $\mathcal{B}_{t-1}$ picks $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_k$; $\mathbf{a}^\perp \leftarrow \mathcal{U}_{k+1,1}$ such that $\mathbf{A}^\top \mathbf{a}^\perp = \mathbf{0}$; $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$; $\mathbf{W}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}$; $\mathbf{v}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$ and sets $[\mathbf{r}_i]_2 := [\overline{\mathbf{B}}\mathbf{v}_i]_2$, for $i \in [n], i \neq t$. Finally, $\mathcal{B}_{t-1}$ embeds the upper part of the MDDH challenge in $\mathbf{r}_t$ by setting $[\mathbf{r}_t]_2 := [\overline{\mathbf{v}}]_2$. It outputs

$$\mathsf{pk} := \Big([\mathbf{A}^\top]_1, [\mathbf{A}^\top \mathbf{W}_0]_1, \ldots, \underbrace{[\mathbf{A}^\top \mathbf{W}_t]_1}_{=[\mathbf{A}^\top \widehat{\mathbf{W}}_t]_1}, \ldots, [\mathbf{A}^\top \mathbf{W}_n]_1, [\mathbf{A}^\top \mathbf{k}]_T, [\mathbf{r}_1]_2, \ldots, \underbrace{[\mathbf{r}_t]_2}_{=[\overline{\mathbf{v}}]_2}, \ldots, [\mathbf{r}_n]_2,$$

$$\{[\mathbf{W}_i \mathbf{r}_j]_2\}_{i,j\in[n],i\neq j,i\neq t}, \{ \underbrace{[\mathbf{W}_t \mathbf{r}_j]_2}_{\substack{=[\widehat{\mathbf{W}}_t \overline{\mathbf{B}}\mathbf{v}_i]_2 \text{ if } \beta = 0 \\ =[\widehat{\mathbf{W}}_t \overline{\mathbf{B}}\mathbf{v}_i + \mathbf{a}^\perp \underline{\mathbf{B}}\mathbf{v}_i]_2 \text{ if } \beta = 1}} \}_{j\in[n],j\neq t}\Big).$$

  Note that the simulated pk is identically distributed (independently of β) as the pk is Game_{t-1} and Game_t (pk is identically distributed in these two games).

- KeyGenO($\ell \in [n]$):
  For each query $\ell \in [n]$, $\mathcal{B}_{t-1}$ picks $\gamma_\ell \leftarrow_{\mathrm{R}} \mathbb{Z}_p$ and outputs $\mathsf{sk}_\ell := [\mathbf{k} + \mathbf{W}_0 \mathbf{r}_\ell + \gamma_\ell \cdot \mathbf{a}^\perp]_2$ if $\ell \leq t-1$, and $\mathsf{sk}_\ell := [\mathbf{k} + \mathbf{W}_0 \mathbf{r}_\ell]_2$ if $\ell > t$, where $\mathbf{W}_0$ is implicitly set to $\widehat{\mathbf{W}}_0$ if $\beta = 0$, and to $\widehat{\mathbf{W}}_0 - \mathbf{a}^\perp \mathbf{T_B}$ if $\beta = 1$. If $\ell = t$, then, β should be 1. If this is not the case, $\mathcal{B}_{t-1}$ aborts the simulation, since the guess was incorrect. Otherwise, it outputs $\mathsf{sk}_t := [\mathbf{k} + \widehat{\mathbf{W}}_0 \overline{\mathbf{v}} + \mathbf{a}^\perp \underline{\mathbf{v}}]_2$. Note that when $[\mathbf{v}]_2$ is a real MDDH challenge, i.e $\underline{\mathbf{v}} = \mathbf{T_B} \overline{\mathbf{v}}$, then $\widehat{\mathbf{W}}_0 \overline{\mathbf{v}} + \mathbf{a}^\perp \underline{\mathbf{v}} = \mathbf{W}_0 \mathbf{r}_t$, that is, sk_t is distributed as in Game_{t-1}. When it is a uniformly random vector, i.e $\underline{\mathbf{v}} = \mathbf{T_B} \overline{\mathbf{v}} + \boxed{\gamma_t}$, where $\gamma_t \leftarrow_{\mathrm{R}} \mathbb{Z}_p$, then $\widehat{\mathbf{W}}_0 \overline{\mathbf{v}} + \mathbf{a}^\perp \underline{\mathbf{v}} + \boxed{\gamma_t \mathbf{a}^\perp}$, that is, sk_t is distributed as in Game_t.

- EncO($\Gamma^\star \subset [n], M_0 \in G_T, M_1 \in G_T$):
  $\mathcal{B}_{t-1}$ picks $b \leftarrow_{\mathrm{R}} \{0,1\}$; $\mathbf{z} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$; sets $C_0 := [\mathbf{z}^\top]_1$ and $C_2 := [\mathbf{z}^\top \mathbf{k}]_T \cdot M_b$. Then,

  - If $\beta = 0$:
    Then, $\mathcal{B}_{t-1}$ sets
    $$C_1 := [\mathbf{z}^\top (\mathbf{W}_0 + \sum_{j \notin \Gamma^\star} \mathbf{W}_j)]_1,$$
    where $\mathbf{W}_0 = \widehat{\mathbf{W}}_0$ and $\mathbf{W}_t = \widehat{\mathbf{W}}_t$.
  - If $\beta = 1$:
    If $t \in \Gamma^\star$, then in particular sk_t cannot be queried by $\mathcal{A}$, by defintion of the security game. Therefore, $\mathcal{B}_{t-1}$ aborts the simulation: the guess was incorrect. Otherwise, it sets
    $$C_1 := [\mathbf{z}^\top (\widehat{\mathbf{W}}_0 + \sum_{i \notin \Gamma^\star, i \neq t} \mathbf{W}_i + \widehat{\mathbf{W}_t})]_1.$$

Note that $\mathcal{B}_{t-1}$ is correctly simulating sk_t, since $\mathbf{W}_0 + \mathbf{W}_t = \widehat{\mathbf{W}}_0 - \mathbf{a}^\perp \mathbf{T}_\mathbf{B} + \widehat{\mathbf{W}_t} + \mathbf{a}^\perp \mathbf{T}_\mathbf{B} = \widehat{\mathbf{W}}_0 + \widehat{\mathbf{W}_t}$, i.e the extra terms cancel out. Recall that these terms must cancel out since $\mathcal{B}_{t-1}$ only knows $\widehat{\mathbf{W}}_0$ and $\widehat{\mathbf{W}}_t$, and not $\mathbf{T}_\mathbf{B}$.

Finally, $\mathcal{B}_{t-1}$ outputs $\mathsf{ct}_{\Gamma^\star} := (C_0, C_1, C_2)$.

We have shown that when the guess β is correct: if $[\mathbf{v}]_2$ is an MDDH challenge, $\mathcal{B}_{t-1}$ simulates Game_{t-1}, otherwise, it simulates Game_t. The guess β is correct with probability exactly $1/2$, since $\beta \leftarrow_\mathrm{R} \{0,1\}$ is independent from $\mathcal{A}$'s view, and independent, in particular, from $\mathcal{A}$'s secret key queries. When β is correct, if the adversary $\mathcal{A}$ guesses correctly b, then $\mathcal{B}_{t-1}$ outputs 1; otherwise, it outputs 0. When the guess β is chosen incorrectly, $\mathcal{B}_{t-1}$ aborts the simulation and outputs 0. We call E the event: guess β is successful, $\neg E$ its complement. We have:

$$
\begin{aligned}
&\mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_k,\mathcal{B}_{t-1}}(\lambda) = |\Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{Br}]_2) = 1] - \Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{u}]_2) = 1]| \\
&= \Pr[E] \cdot \Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{Br}]_2) = 1|E] + \Pr[\neg E] \cdot \Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{Br}]_2) = 1|\neg E] \\
&- \Pr[E] \cdot \Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{u}]_2) = 1|E] - \Pr[\neg E] \cdot \Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{u}]_2) = 1|\neg E]| \\
&\geq 1/2 \cdot |\Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{Br}]_2) = 1|E] - \Pr[\mathcal{B}_{t-1}(\mathbb{G}, [\mathbf{B}]_2, [\mathbf{u}]_2) = 1|E]| \\
&= 1/2 \cdot |\mathrm{Adv}_{t-1} - \mathrm{Adv}_t|
\end{aligned}
$$

where $\mathbf{r} \leftarrow_\mathrm{R} \mathbb{Z}_p^k$, $\mathbf{u} \leftarrow_\mathrm{R} \mathbb{Z}_p^{k+1}$, and the probabilities are taken over the random coins of $\mathcal{A}$ and $\mathcal{B}_{t-1}$. □

Lemma 6.4 (Game_n)**.**

$$|\mathrm{Adv}_n| \leq 1/p.$$

Proof. Here, we argue that the component of $\mathbf{k}$ in $\mathsf{span}(\mathbf{a}^\perp)$ is masked by γ_ℓ in sk_ℓ, and it is hidden from the pk since only $\mathbf{A}^\top \mathbf{k}$ appears. Therefore, if $\mathbf{z}$ used in the challenge ciphertext is not in the span of $\mathbf{A}$, then the value $[\mathbf{z}^\top \mathbf{k}]_T$ is random and completely masks the message M_b. We use the fact that the following two distributions are the same:

$$(\mathbf{k}, \{\gamma_i\}_{\ell \in [n]}) \text{ and } (\mathbf{k} + \mu \cdot \mathbf{a}^\perp, \{\gamma_i - \mu\}_{i \in [n]}),$$

where $\mathbf{k} \leftarrow_\mathrm{R} \mathbb{Z}_p^{k+1}$, $\gamma_\ell, \mu \leftarrow_\mathrm{R} \mathbb{Z}_p$. Note that the extra component $\mu \cdot \mathbf{a}^\perp$ does not appear in the pk since $\mathbf{A}^\top(\mathbf{k} + \mu \cdot \mathbf{a}^\perp)) = \mathbf{A}^\top \mathbf{k}$ and it does not appear in the secret keys since $\mathbf{k} + \mu \cdot \mathbf{a}^\perp + \gamma_\ell \cdot \mathbf{a}^\perp - \mu \cdot \mathbf{a}^\perp = \mathbf{k} + \gamma_\ell \cdot \mathbf{a}^\perp$, i.e the extra terms cancel out. The challenge ciphertext contains $C_2 := [\mathbf{z}^\top \mathbf{k} + \boxed{\mu \cdot \mathbf{z}^\top \mathbf{a}^\perp}]_T \cdot M_b$. If $\mathbf{z}^\top \mathbf{a}^\perp \neq 0$ (which happens with probability $1 - 1/p$ over the choice of $\mathbf{z} \leftarrow_\mathrm{R} \mathbb{Z}_p^{k+1}$), then, C_2 is uniformly random over G_T, and $\mathcal{A}$ cannot guess b with probability more than $1/2$. □

7 Tightly Secure, Prime Order Construction

Figure 11 contains our tightly secure prime-order construction. It is similar to the scheme of Figure 9 except that matrices $\mathbf{A}$ and $\mathbf{B}$ are sampled from the distribution $\mathcal{D}_{2k,k}$, and not $\mathcal{D}_k$. In fact, the scheme of Figure 9 is already tightly-secure under the SXDH assumption (corresponding to the case $k = 1$). We generalize the construction here to work under the $\mathcal{D}_{2k,k}$-MDDH Assumption for any $k \in \mathbb{N}$ and any matrix distribution $\mathcal{D}_{2k,k}$.

$\underline{\mathsf{Setup}(1^\lambda, 1^n)}$:
$\mathbb{G} \leftarrow_{\mathrm{R}} \mathcal{G}(1^\lambda)$;$\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_{2k,k}$; $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k}$; $\{\mathbf{W}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k\times k}, \mathbf{r}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k\}_{i\in[n]}$
Output $\mathsf{pk} := \Big([\mathbf{A}^\top]_1, [\mathbf{A}^\top\mathbf{W}_0]_1\{[\mathbf{A}^\top\mathbf{W}_i]_1, [\mathbf{r}_i]_2\}_{i\in[n]}, [\mathbf{A}^\top\mathbf{k}]_T, \{[\mathbf{W}_j\mathbf{r}_i]_2\}_{i,j\in[n],i\neq j}\Big)$ and
$\mathsf{msk} := \big([\mathbf{k}]_2, \{[\mathbf{W}_0\mathbf{r}_i]_2\}_{i\in[n]}\big)$.

$\underline{\mathsf{KeyGen}(\mathsf{msk}, i \in [n])}$:
Output $\mathsf{sk}_i := [\mathbf{k} + \mathbf{W}_0\mathbf{r}_i]_2 \in G_2^{2k}$.

$\underline{\mathsf{Enc}(\mathsf{pk}, \Gamma \subset [n], M \in G_T)}$:
$\mathbf{s} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$
$C_0 := [\mathbf{s}^\top\mathbf{A}^\top]_1$; $C_1 := [\mathbf{s}^\top\mathbf{A}^\top(\mathbf{W}_0 + \sum_{j\notin\Gamma^\star} \mathbf{W}_j)]_1$; $C_2 := [\mathbf{s}^\top\mathbf{A}^\top\mathbf{k}]_T \cdot M$
Output $\mathsf{ct}_\Gamma := (C_0, C_1, C_2) \in G_1^{3k} \times G_T$

$\underline{\mathsf{Dec}(\mathsf{ct}_\Gamma, \mathsf{sk}_i)}$: // ct_Γ and sk_i implictly contain a description of Γ and i
Compute $e(\underbrace{[\mathbf{s}^\top\mathbf{A}^\top]_1}_{=C_0}, \underbrace{[\mathbf{k} + \mathbf{W}_0\mathbf{r}_i]_2}_{=\mathsf{sk}_i}) / e(\underbrace{[\mathbf{s}^\top\mathbf{A}^\top(\mathbf{W}_0 + \sum_{j\notin\Gamma} \mathbf{W}_j)]_1}_{=C_1}, \underbrace{[\mathbf{r}_i]_1}_{\in\mathsf{pk}}) = [\mathbf{s}^\top\mathbf{A}^\top\mathbf{k} - \mathbf{s}^\top\mathbf{A}^\top \textstyle\sum_{j\notin\Gamma} \mathbf{W}_j\mathbf{r}_i]_T$.
Multiply the previous term by $e(\underbrace{[\mathbf{s}^\top\mathbf{A}^\top]_1}_{=C_0}, \underbrace{[\sum_{j\notin\Gamma} \mathbf{W}_j\mathbf{r}_i]_2}_{\in\mathsf{pk} \text{ for } i\in\Gamma})$ to get the encapsulation key $[\mathbf{s}^\top\mathbf{A}^\top\mathbf{k}]_T$ used to obtain the message M.

Figure 11: $\mathsf{BE}_{\mathsf{tight}}$, a tightly, adaptively-secure broadcast encryption scheme based on prime-order bilinear groups.

We prove the tight security of this scheme in the multichallenge setting in Section 8.

8 Security Proof of the Tightly Secure, Prime-Order Construction

We now prove the security of the scheme $\mathsf{BE}_{\mathsf{tight}}$, presented in Figure 11.

Theorem 8.1. *If the $\mathcal{D}_k$-MDDH Assumption holds in G_1 and G_2, then the broadcast encryption scheme $\mathsf{BE}_{\mathsf{tight}}$ defined in Figure 11 is tightly, adaptively secure in the multi-challenge setting (as defined in section 2.5). Namely, for any adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\mathsf{ct}} + Q_{\mathsf{sk}}) \cdot \mathsf{poly}(\lambda)$, where Q_{ct}, Q_{sk} denote the number of calls to EncO and $\mathsf{KeyGenO}$ respectively, $\mathsf{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$, and*

$$\mathrm{Adv}^{\mathsf{BE}}_{\mathcal{A},\mathsf{multiChal}}(\lambda) \leq 2(n+k) \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{2k,k},\mathcal{B}}(\lambda) + 2^{-\Omega(\lambda)},$$

where n is the number of users.

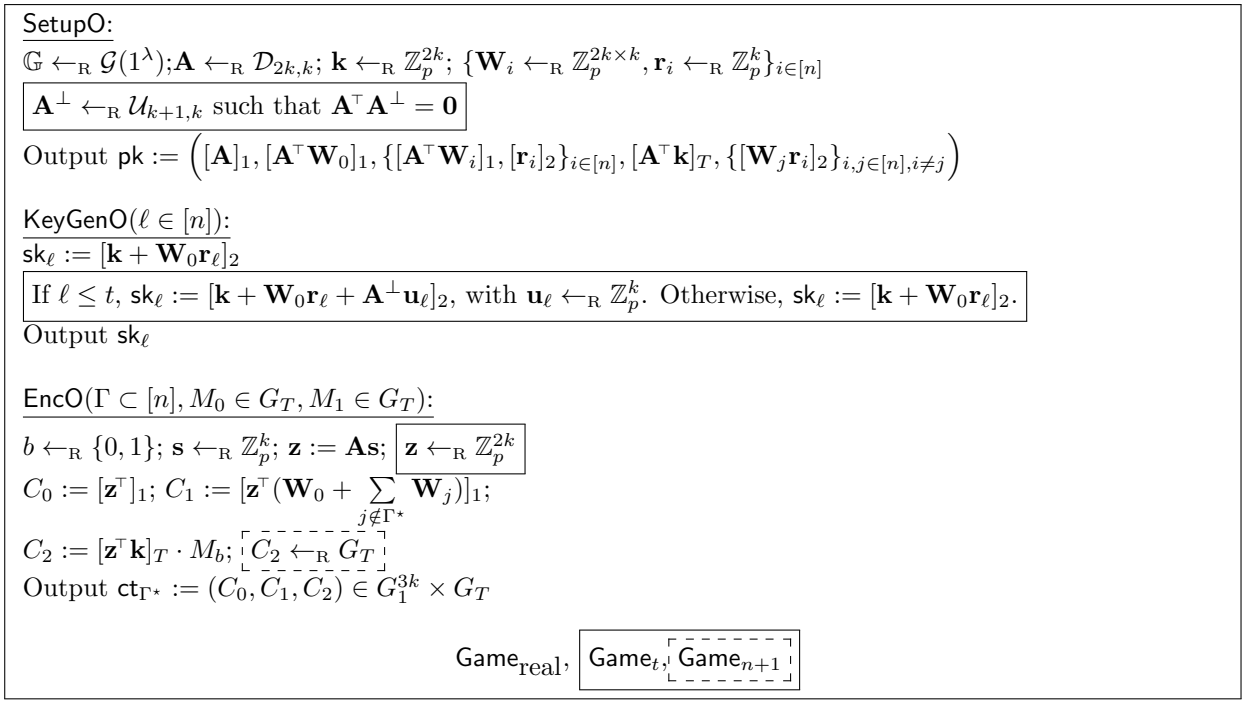

Figure 12: $\mathsf{Game}_{\mathrm{real}}, \mathsf{Game}_t$ (for $0 \leq t \leq n$), Game_{n+1} for the proof of security of $\mathsf{BE}_{\mathsf{tight}}$ defined in Figure 11. Here n denotes the number of users. In each procedure, the components inside a solid (dotted) frame are only present in the games marked by a solid (dotted) frame.

We prove Theorem 8.1 via a series of games described in Figure 12 and we use Adv_i to denote the advantage of $\mathcal{A}$ in game Game_i. Namely, we show that:

$$\mathsf{Game}_{\mathrm{real}} \approx_c \mathsf{Game}_0 \approx_c \mathsf{Game}_1 \approx_c \ldots \approx_c \mathsf{Game}_{n+1}$$

where $\approx_c$ denotes computational indistinguishability. Game_{real} is the security game as defined in Section 2.5, and the other Game_i are defined in Figure 12. Theorem 8.1 follows from the Lemma 8.2, 8.3, 8.4 and 8.5 below.

Lemma 8.2 ($\mathsf{Game}_{real} \approx_c \mathsf{Game}_0$). *There exists an adversary $\mathcal{B}_0$ such that $\mathbf{T}(\mathcal{B}_0) \approx \mathbf{T}(\mathcal{A}) + (Q_{\mathsf{ct}} + Q_{\mathsf{sk}}) \cdot \mathsf{poly}(\lambda)$, where Q_{ct}, Q_{sk} denote the number of calls to EncO and $\mathsf{KeyGenO}$ respectively, $\mathsf{poly}(\lambda)$ is independent of $\mathbf{T}(\mathcal{A})$, and*

$$|\mathrm{Adv}_{real} - \mathrm{Adv}_0| \leq k \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{2k,k},\mathcal{B}_0}(\lambda) + \frac{1}{p-1}.$$

Here, we use the MDDH assumption to tightly switch the distribution of all challenge ciphertexts.

Proof. To go from $\mathsf{Game}_{\mathrm{real}}$ to Game_0, we switch the distribution of the vector $[\mathbf{z}]_1$ in all challenge ciphertexts, using the Q_{ct}-fold $\mathcal{D}_k$-MDDH Assumption on $[\mathbf{A}]_1$ (see Definition 2.3).

We build an adversary $\mathcal{B}'_0$ against the Q_{ct}-fold $\mathcal{D}_{2k,k}$-MDDH Assumption, such that $\mathbf{T}(\mathcal{B}'_0) \approx \mathbf{T}(\mathcal{A}) + (Q_{\mathsf{ct}} + Q_{\mathsf{sk}}) \cdot \mathsf{poly}(\lambda)$ with $\mathsf{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$, and

$$|\mathrm{Adv}_0 - \mathrm{Adv}_1| \leq \mathrm{Adv}^{Q_{\mathsf{ct}}\text{-}\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{2k,k},\mathcal{B}'_0}(\lambda).$$

This implies the lemma by Lemma 2.6 (self-reducibility of $\mathcal{D}_{2k,k}$-MDDH). Upon receiving a challenge $(\mathbb{G}, [\mathbf{A}]_1 \in G_1^{2k\times k}, [\mathbf{H}]_1 := [\mathbf{h}_1|\mathbf{h}_2|\ldots|\mathbf{h}_{Q_{\mathsf{ct}}}]_1 \in G_1^{2k\times Q_{\mathsf{ct}}})$ for the Q_{ct}-fold $\mathcal{D}_k$-MDDH Assumption, $\mathcal{B}'_0$ picks $b \leftarrow_{\mathrm{R}} \{0,1\}$, $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k}$; $\mathbf{W}_0, \ldots, \mathbf{W}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k\times k}$; $\mathbf{r}_1, \ldots, \mathbf{r}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$, thanks to which it can simulate SetupO and $\mathsf{KeyGenO}$ as described in Figure 12. To simulate $\mathsf{EncO}(\Gamma, M_0, M_1)$ on its i'th query, for $i = 1, \ldots, Q_{\mathsf{ct}}$, $\mathcal{B}'_0$ sets $[\mathbf{z}]_1 := [\mathbf{h}_i]_1$, and returns $([\mathbf{z}^\top]_1, [\mathbf{z}^\top(\mathbf{W}_0 + \sum_{j\notin\Gamma} \mathbf{W}_j)]_1, [\mathbf{z}^\top\mathbf{k}]_T \cdot M_b$. Note that when $[\mathbf{H}]_1$ is a proper MDDH sample, $\mathcal{B}'_0$ simulates $\mathsf{Game}_{\mathrm{real}}$, and when $[\mathbf{H}]_1$ is uniformly random over $G_1^{2k\times Q_{\mathsf{ct}}}$, it simulates Game_0. □

Lemma 8.3 ($\mathsf{Game}_{t-1} \approx_c \mathsf{Game}_t$). *For all $t \in [n]$, there exists an adversary $\mathcal{B}_{t-1}$ such that $\mathbf{T}(\mathcal{B}_{t-1}) \approx \mathbf{T}(\mathcal{A}) + (Q_{\mathsf{ct}} + Q_{\mathsf{sk}}) \cdot \mathsf{poly}(\lambda)$, with $\mathsf{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$; and*

$$|\mathrm{Adv}_{t-1} - \mathrm{Adv}_t| \leq 2 \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{2k,k},\mathcal{B}_{t-1}}(\lambda).$$

Here, we embed an MDDH challenge in pk and sk_t. More precisely, the simulator sets $\mathbf{r}_t := \overline{\mathbf{B}}\mathbf{v}_t \in \mathbb{Z}_p^k$, where $\mathbf{B} \leftarrow_{\mathrm{R}} \mathcal{D}_k$ and $\mathbf{v}_t \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k$, i.e. the upper part of an MDDH challenge. The lower part $\underline{\mathbf{B}}\mathbf{v}_t \in \mathbb{Z}_p$ is embedded in sk_t, if sk_t is queried by the adversary (it may not be the case, in particular if $t \in \Gamma^\star$). Note that the simulator needs to know if sk_t is going to be queried by the adversary when simulating pk. The proof of this lemma is exactly as the proof of Lemma 6.3, in Section 5. See the later for further details.

Lemma 8.4 ($\mathsf{Game}_n \approx_c \mathsf{Game}_{n+1}$). *There exists an adversary $\mathcal{B}_n$ such that $\mathbf{T}(\mathcal{B}_n) \approx \mathbf{T}(\mathcal{A}) + (Q_{\mathsf{ct}} + Q_{\mathsf{sk}}) \cdot \mathsf{poly}(\lambda)$, with $\mathsf{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$; and*

$$|\mathrm{Adv}_n - \mathrm{Adv}_{n+1}| \leq k \cdot \mathrm{Adv}^{\mathsf{MDDH}}_{\mathcal{G},\mathcal{D}_{2k,k},\mathcal{B}_n}(\lambda) + \frac{1}{p-1}.$$

Here, we use the MDDH Assumption to increase the entropy in all challenge ciphertexts, thereby hiding the underlying plaintexts. Namely, we use the fact that the vector $\mathbf{k}$ has some entropy that is not revealed by pk and the queried sk, which can be used together with the randomness of the challenge ciphertexts (the vector $\mathbf{z} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k}$) to embed an MDDH challenge.

Proof. We use the three following facts:

1. we can write $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k} := \mathbf{k}' + \mathbf{A}^{\perp}\mathbf{u}$, where $\mathbf{k}' \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k}$ and $\mathbf{u} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k}$, without changing the distribution of $\mathbf{k}$. Note that the term $\mathbf{A}^{\perp}\mathbf{u}$ does not appear in the pk since $\mathbf{A}^{\top}(\mathbf{k}' + \mathbf{A}^{\perp}\mathbf{u}) = \mathbf{A}^{\top}\mathbf{k}'$, and is hidden from the queried sk_{ℓ} thanks to the $\mathbf{A}^{\perp}\mathbf{u}_{\ell}$ components in them.

2. we write $\mathbf{z} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k}$ sampled in challenge ciphertexts as $\mathbf{z} := \mathbf{A}\mathbf{v} + \mathbf{B}\mathbf{w}$, where $\mathbf{v}, \mathbf{w} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k}$, $(\mathbf{A}|\mathbf{B})$ forms a basis of $\mathbb{Z}_p^{2k}$, and $\mathbf{B}^{\top}\mathbf{A}^{\perp}$ is invertible.

3. we write the vector $\mathbf{u}$ used in $\mathbf{k}' + \mathbf{A}^{\perp}\mathbf{u}$ as $\left(\mathbf{B}^{\top}\mathbf{A}^{\perp}\right)^{-1}\mathbf{u}'$, where $\mathbf{u}' \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k}$ which does not change the distribution of $\mathbf{u}$.

From these three facts, we deduce we can write challenge ciphertexts as $\mathsf{ct}_{\Gamma} := \big([\mathbf{z}]_1 := [\mathbf{v}^{\top}\mathbf{A}^{\top} + \mathbf{w}^{\top}\mathbf{B}^{\top}]_1, [\mathbf{z}^{\top}(\mathbf{W}_0 + \sum_{j \notin \Gamma} \mathbf{W}_j)]_1, [\mathbf{z}^{\top}\mathbf{k}']_T + [\mathbf{w}^{\top}\mathbf{u}']_T + M_b\big)$. We can argue that $[\binom{\mathbf{w}}{\mathbf{w}^{\top}\mathbf{u}'}]_1 \approx_c [\mathbf{w}']_1 \leftarrow_{\mathrm{R}} G_1^{k+1}$ by the $\mathcal{U}_k$-MDDH Assumption. More concretely, we build an adversary $\mathcal{B}'_n$, against the Q_{ct}-fold $\mathcal{U}_k$-MDDH Assumption, such that $\mathbf{T}(\mathcal{B}'_n) \approx \mathbf{T}(\mathcal{A}) + (Q_{\mathsf{ct}} + Q_{\mathsf{sk}}) \cdot \mathsf{poly}(\lambda)$ with $\mathsf{poly}(\lambda)$ independent of $\mathbf{T}(\mathcal{A})$, and

$$|\mathrm{Adv}_n - \mathrm{Adv}_{n+1}| \leq \mathrm{Adv}^{Q_{\mathsf{ct}}\text{-}\mathsf{MDDH}}_{\mathcal{G},\mathcal{U}_k,\mathcal{B}'_n}(\lambda).$$

This implies the lemma by Lemma 2.5 ($\mathcal{D}_{2k,k}$-$\mathsf{MDDH} \Rightarrow \mathcal{U}_k$-$\mathsf{MDDH}$) and Lemma 2.6 (self-reducibility of $\mathcal{D}_{2k,k}$-MDDH).

Upon receiving a challenge $(\mathbb{G}, [\mathbf{U}]_1 \leftarrow_{\mathrm{R}} G_1^{(k+1)\times k}, [\mathbf{H}]_1 := [\mathbf{h}_1|\mathbf{h}_2|\ldots|\mathbf{h}_{Q_{\mathsf{ct}}}]_1 \in G_1^{(k+1)\times Q_{\mathsf{ct}}})$ for the Q_{ct}-fold $\mathcal{D}_k$-MDDH Assumption, $\mathcal{B}'_n$ picks $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_{2k,k}$, $b \leftarrow_{\mathrm{R}} \{0,1\}$, $\mathbf{k}' \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k}$; $\mathbf{W}_0, \ldots, \mathbf{W}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k\times k}$; $\mathbf{r}_1, \ldots, \mathbf{r}_n \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k}$, thanks to which it can simulate SetupO and $\mathsf{KeyGenO}$ as described in Figure 12. To simulate $\mathsf{EncO}(\Gamma, M_0, M_1)$ on its i'th query, for $i = 1, \ldots, Q_{\mathsf{ct}}$, $\mathcal{B}'_n$ picks $\mathbf{v} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k}$, $\mathbf{B} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{2k\times k}$ and sets $[\mathbf{z}]_1 := [\mathbf{A}\mathbf{v}]_1 + [\mathbf{B}\overline{\mathbf{h}}_i]_1$, and returns $\mathsf{ct}_{\Gamma} := ([\mathbf{z}^{\top}]_1, [\mathbf{z}^{\top}(\mathbf{W}_0 + \sum_{j \notin \Gamma} \mathbf{W}_j)]_1, [\mathbf{z}^{\top}\mathbf{k}']_T + [\underline{\mathbf{h}}_i]_T + M_b$. Note that when $[\mathbf{H}]_1$ is a proper $\mathcal{U}_k$-MDDH sample, that is, when $\mathbf{h}_i := \binom{\mathbf{w}}{\mathbf{w}^{\top}\mathbf{u}'}$ with $\mathbf{w}, \mathbf{u}' \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k}$, $\mathcal{B}'_n$ simulates $\mathsf{Game}_{\mathrm{n}}$, and when $[\mathbf{H}]_1$ is uniformly random over $G_1^{(k+1)\times Q_{\mathsf{ct}}}$, it simulates Game_{n+1}. □

Lemma 8.5 (Game_{n+1})**.** *For all $\mathcal{A}$,* $\mathrm{Adv}_{n+1} = 0$.

Proof. In this game, plaintexts are completely masked by uniformly random values in the challenge ciphertexts, so nonzero advantage is impossible to achieve. □

References

[AC16] Shashank Agrawal and Melissa Chase. A study of pair encodings: Predicate encryption in prime order groups. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 259–288, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.

[Att14] Nuttapong Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 557–577, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.

[Att15] Nuttapong Attrapadung. Dual system encryption framework in prime-order groups. *IACR Cryptology ePrint Archive*, 2015:390, 2015.

[BBG05] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 440–456, Aarhus, Denmark, May 22–26, 2005. Springer, Heidelberg, Germany.

[BBM00] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany.

[BGI+01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.

[BGW05] Dan Boneh, Craig Gentry, and Brent Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 258–275, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Heidelberg, Germany.

[BKP14] Olivier Blazy, Eike Kiltz, and Jiaxin Pan. (Hierarchical) identity-based encryption from affine message authentication. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 408–425, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.

[BW06] Dan Boneh and Brent Waters. A fully collusion resistant broadcast, trace, and revoke system. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 211–220, Alexandria, Virginia, USA, October 30 – November 3, 2006. ACM Press.

[BWZ14] Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 206–223, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.

[BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.

[CGW15] Jie Chen, Romain Gay, and Hoeteck Wee. Improved dual system ABE in prime-order groups via predicate encodings. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 595–624, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.

[Che06] Jung Hee Cheon. Security analysis of the strong Diffie-Hellman problem. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 1–11, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany.

[CM14] Melissa Chase and Sarah Meiklejohn. Déjà Q: Using dual systems to revisit q-type assumptions. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 622–639, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.

[CMM16a] Melissa Chase, Mary Maller, and Sarah Meiklejohn. Déjà Q all over again: Tighter and broader reductions of q-type assumptions. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 655–681, Hanoi, Vietnam, December 4–8, 2016. Springer, Heidelberg, Germany.

[CMM16b] Melissa Chase, Mary Maller, and Sarah Meiklejohn. Deja q all over again: Tighter and broader reductions of q-type assumptions. Cryptology ePrint Archive, Report 2016/840, 2016. http://eprint.iacr.org/.

[Cor00] Jean-Sébastien Coron. On the exact security of full domain hash. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 229–235, Santa Barbara, CA, USA, August 20–24, 2000. Springer, Heidelberg, Germany.

[CW13] Jie Chen and Hoeteck Wee. Fully, (almost) tightly secure IBE and dual system groups. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 435–460, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

[CW14] Jie Chen and Hoeteck Wee. Semi-adaptive attribute-based encryption and improved delegation for Boolean formula. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 277–297, Amalfi, Italy, September 3–5, 2014. Springer, Heidelberg, Germany.

[DF02] Y. Dodis and N. Fazio. Public key broadcast encryption for stateless receivers. In *DRM*, pages 61–80, 2002.

[EHK+13] Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

[FN94] Amos Fiat and Moni Naor. Broadcast encryption. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 480–491, Santa Barbara, CA, USA, August 22–26, 1994. Springer, Heidelberg, Germany.

[Fre10] David Mandell Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 44–61, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.

[GKSW10] Sanjam Garg, Abishek Kumarasubramanian, Amit Sahai, and Brent Waters. Building efficient fully collusion-resilient traitor tracing and revocation schemes. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 10*, pages 121–130, Chicago, Illinois, USA, October 4–8, 2010. ACM Press.

[GKW15] Romain Gay, Iordanis Kerenidis, and Hoeteck Wee. Communication complexity of conditional disclosure of secrets and attribute-based encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 485–502, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.

[GST04] Michael T. Goodrich, Jonathan Z. Sun, and Roberto Tamassia. Efficient tree-based revocation in groups of low-state devices. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 511–527, Santa Barbara, CA, USA, August 15–19, 2004. Springer, Heidelberg, Germany.

[GSW00] Juan A. Garay, Jessica Staddon, and Avishai Wool. Long-lived broadcast encryption. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 333–352, Santa Barbara, CA, USA, August 20–24, 2000. Springer, Heidelberg, Germany.

[GSY99] Eli Gafni, Jessica Staddon, and Yiqun Lisa Yin. Efficient methods for integrating traceability and broadcast encryption. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 372–387, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany.

[Gui13] Aurore Guillevic. Comparing the pairing efficiency over composite-order and prime-order elliptic curves. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *ACNS 13*, volume 7954 of *LNCS*, pages 357–372, Banff, AB, Canada, June 25–28, 2013. Springer, Heidelberg, Germany.

[GW09] Craig Gentry and Brent Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 171–188, Cologne, Germany, April 26–30, 2009. Springer, Heidelberg, Germany.

[HJ12] Dennis Hofheinz and Tibor Jager. Tightly secure signatures and public-key encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 590–607, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.

[HS02] Dani Halevy and Adi Shamir. The LSD broadcast encryption scheme. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 47–60, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Heidelberg, Germany.

[Lew12] Allison B. Lewko. Tools for simulating features of composite order bilinear groups in the prime order setting. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 318–335, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.

[LW10] Allison B. Lewko and Brent Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 455–479, Zurich, Switzerland, February 9–11, 2010. Springer, Heidelberg, Germany.

[NNL01] Dalit Naor, Moni Naor, and Jeffery Lotspiech. Revocation and tracing schemes for stateless receivers. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 41–62, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.

[OT08] Tatsuaki Okamoto and Katsuyuki Takashima. Homomorphic encryption and signatures from vector decomposition. In Steven D. Galbraith and Kenneth G. Paterson, editors, *PAIRING 2008*, volume 5209 of *LNCS*, pages 57–74, Egham, UK, September 1–3, 2008. Springer, Heidelberg, Germany.

[OT09] Tatsuaki Okamoto and Katsuyuki Takashima. Hierarchical predicate encryption for inner-products. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 214–231, Tokyo, Japan, December 6–10, 2009. Springer, Heidelberg, Germany.

[SSW00] J.N. Staddon, D.R. Stinson, and R. Wei. Combinatorial properties of frameproof and traceability codes. Cryptology ePrint Archive, Report 2000/004, 2000. http://eprint.iacr.org/2000/004.

[Sv98] Douglas R. Stinson and Tran van Trung. Some new results on key distribution patterns and broadcast encryption. *Designs, Codes and Cryptography*, 14(3):261–279, 1998.

[Wat09] Brent Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 619–636, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Heidelberg, Germany.

[Wee14] Hoeteck Wee. Dual system encryption via predicate encodings. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 616–637, San Diego, CA, USA, February 24–26, 2014. Springer, Heidelberg, Germany.

[Wee16] Hoeteck Wee. Déjà Q: Encore! Un petit IBE. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 237–258, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.