Tight Adaptively Secure Broadcast Encryption with Short Ciphertexts and Keys

Romain Gay[1,2], Lucas Kowalczyk[3(✉)], and Hoeteck Wee[1,2]

[1] DIENS, Ecole normale superieure, CNRS, PSL University, 75005 Paris, France
rgay@ens.fr, wee@di.ens.fr
[2] INRIA, Paris, France
[3] Columbia University, New York, USA
luke@cs.columbia.edu

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.

D. Catalano and R. De Prisco (Eds.): SCN 2018, LNCS 11035, pp. 123–139, 2018.
https://doi.org/10.1007/978-3-319-98113-0_7

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, 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 Fig. 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, LL15] 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

Fig. 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 Fig. 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 Sect. 3 and its proof in [GKW18, Sect. 3]. 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 Sect. 4.

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

Fig. 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 we replace the r_i, w_i in the composite-order scheme of Sect. 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 Sect. 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 Sect. 4, 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 Sect. 3 is proven under three static assumptions in composite-order asymmetric bilinear groups. We define the advantage functions referred to in the assumptions in Fig. 3.

Definition 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.*

$$\mathrm{Adv}^{SD1}_{\mathcal{G},\mathcal{A}}(\lambda) := |\Pr[\mathcal{A}(D,T_0)=1] - \Pr[\mathcal{A}(D,T_1)=1]|$$
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 $T_0 := e(g_p, h_p)^{s\alpha},\ T_1 = X \leftarrow_{\mathrm{R}} G_T$

Fig. 3. Advantage functions

2.4 Matrix Diffie-Hellman Assumptions

The security of the prime-order scheme in Sect. 4 is proven under the Matrix Decision Diffie-Hellman (MDDH) Assumption [EHK+13], whose definition we recall here.

Definition 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 $\mathcal{D}_k := \mathcal{D}_{k+1,k}$.*

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 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 $\leftarrow_R \mathcal{G}(1^\lambda)$, $\mathbf{A} \leftarrow_R \mathcal{D}_k, \mathbf{w} \leftarrow_R \mathbb{Z}_p^k, \mathbf{u} \leftarrow_R \mathbb{Z}_p^{\ell}$.

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 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 1 ($\mathcal{D}_{\ell,k}$-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 $([\mathbf{A}]_s, [\mathbf{AW}]_s)$ 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 (Random self-reducibility of $\mathcal{D}_{\ell,k}$-MDDH, [EHK+13]). *Let ℓ, 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

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

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

$\underline{\mathsf{Enc}(\mathsf{pk}, \Gamma \subset [n], M \in G_T)}$:
$s \leftarrow_{\mathrm{R}} \mathbb{Z}_N$
$C_0 := g_p^s$; $C_1 := g_p^{s\left(u + \sum_{j\notin\Gamma} w_j\right)}$; $C_2 := e(g_p,h_p)^{\alpha s} \cdot M$
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 $M = C_2 \cdot D_0 \cdot D_1 \cdot D_2$.

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

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 Sect. 5, 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. The security proof is given in the full version of this paper [GKW18, Sect. 4].

$\underline{\mathsf{Setup}(1^\lambda, 1^n)}$:
$\mathbb{G} \leftarrow_{\mathrm{R}} \mathcal{G}(1^\lambda)$; $\mathbf{A} \leftarrow_{\mathrm{R}} \mathcal{D}_k$; $\mathbf{k} \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{k+1}$; $\{\mathbf{W}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^{(k+1)\times k}, \mathbf{r}_i \leftarrow_{\mathrm{R}} \mathbb{Z}_p^k\}_{i\in[n]}$
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 $D_0 = e(\underbrace{[\mathbf{s}^\top\mathbf{A}^\top]_1}_{=C_0}, \underbrace{[\mathbf{k} + \mathbf{W}_0\mathbf{r}_i]_2}_{=\mathsf{sk}_i}) = [\mathbf{s}^\top\mathbf{A}^\top\mathbf{k} + \mathbf{s}^\top\mathbf{A}^\top\mathbf{W}_0\mathbf{r}_i]_T$.
Compute $D_1 = e(\underbrace{[\mathbf{s}^\top\mathbf{A}^\top(\mathbf{W}_0 + \sum_{j\notin\Gamma}\mathbf{W}_j)]_1}_{=C_1}, \underbrace{[\mathbf{r}_i]_2}_{\in\mathsf{pk}}) = [\mathbf{s}^\top\mathbf{A}^\top\mathbf{W}_0\mathbf{r}_i + \mathbf{s}^\top\mathbf{A}^\top\sum_{j\notin\Gamma}\mathbf{W}_j\mathbf{r}_i]_T$.
Compute $D_2 = e(\underbrace{[\mathbf{s}^\top\mathbf{A}^\top]_1}_{=C_0}, \underbrace{[\sum_{j\notin\Gamma}\mathbf{W}_j\mathbf{r}_i]_2}_{\in\mathsf{pk}\text{ for }i\in\Gamma}) = [\mathbf{s}^\top\mathbf{A}^\top\sum_{j\notin\Gamma}\mathbf{W}_j\mathbf{r}_i]_T$.
Compute and output $M = C_2 \cdot D_0 \cdot D_1^{-1} \cdot D_2$

Fig. 5. $\mathsf{BE}_{\mathsf{prime}}$, an adaptively secure broadcast encryption scheme based on prime-order bilinear groups.

4 Prime Order Construction

Our prime-order construction is detailed in Fig. 5. The security proof is given in the full version of this paper [GKW18, Sect. 6].

5 Tightly Secure, Prime Order Construction

We give the description of our construction and its security proof in the full version of this paper [GKW18, Sects. 7 and 8].

Acknowledgements. Romain Gay is partially supported by a Google Fellowship. Lucas Kowalczyk's work has been done while visiting ENS, Paris. He is 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. Hoeteck Wee is supported in part by ERC Project aSCEND (H2020 639554).

References

[AC16] Agrawal, S., Chase, M.: A study of pair encodings: predicate encryption in prime order groups. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016. LNCS, vol. 9563, pp. 259–288. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49099-0_10

[Att14] Attrapadung, N.: Dual system encryption via doubly selective security: framework, fully secure functional encryption for regular languages, and more. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 557–577. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_31

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

[BBG05] Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005). https://doi.org/10.1007/11426639_26

[BBM00] Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_18

[BGI+01] Barak, B., et al.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44647-8_1

[BGW05] Boneh, D., Gentry, C., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 258–275. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_16

[BKP14] Blazy, O., Kiltz, E., Pan, J.: (Hierarchical) identity-based encryption from affine message authentication. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 408–425. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_23

[BW06] Boneh, D., Waters, B.: A fully collusion resistant broadcast, trace, and revoke system. In: Juels, A., Wright, R.N., di Vimercati, S.C. (eds.) ACM CCS 06, Alexandria, Virginia, USA, 30 October–3 November 2006, pp. 211–220. ACM Press (2006)

[BWZ14] Boneh, D., Waters, B., Zhandry, M.: Low overhead broadcast encryption from multilinear maps. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 206–223. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_12

[BZ14] Boneh, D., Zhandry, M.: Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 480–499. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_27

[CGW15] Chen, J., Gay, R., Wee, H.: Improved dual system ABE in prime-order groups via predicate encodings. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 595–624. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_20

[Che06] Cheon, J.H.: Security analysis of the strong Diffie-Hellman problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 1–11. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_1

[CM14] Chase, M., Meiklejohn, S.: Déjà Q: using dual systems to revisit q-Type assumptions. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 622–639. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_34

[CMM16a] Chase, M., Maller, M., Meiklejohn, S.: Déjà Q all over again: tighter and broader reductions of q-type assumptions. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10032, pp. 655–681. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_22

[CMM16b] Chase, M., Maller, M., Meiklejohn, S.: Déjà Q all over again: tighter and broader reductions of q-type assumptions. Cryptology ePrint Archive, Report 2016/840, (2016). http://eprint.iacr.org/

[Cor00] Coron, J.-S.: On the exact security of full domain hash. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 229–235. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44598-6_14

[CW13] Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_25

[CW14] Chen, J., Wee, H.: Semi-adaptive attribute-based encryption and improved delegation for boolean formula. In: Abdalla, M., De Prisco, R. (eds.) SCN 2014. LNCS, vol. 8642, pp. 277–297. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10879-7_16

[DF02] Dodis, Y., Fazio, N.: Public key broadcast encryption for stateless receivers. In: Feigenbaum, J. (ed.) DRM 2002. LNCS, vol. 2696, pp. 61–80. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-44993-5_5

[EHK+13] Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_8

[FN94] Fiat, A., Naor, M.: Broadcast encryption. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 480–491. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_40

[Fre10] Freeman, D.M.: Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 44–61. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_3

[GKSW10] Garg, S., Kumarasubramanian, A., Sahai, A., Waters, B.: Building efficient fully collusion-resilient traitor tracing and revocation schemes. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM CCS 10, Chicago, Illinois, USA, 4–8 October 2010, pp. 121–130. ACM Press (2010)

[GKW15] Gay, R., Kerenidis, I., Wee, H.: Communication complexity of conditional disclosure of secrets and attribute-based encryption. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9216, pp. 485–502. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48000-7_24

[GKW18] Gay, R., Kowalczyk, L., Wee, H.: Tight adaptively secure broadcast encryption with short ciphertexts and keys. Cryptology ePrint Archive, Report 2018/391 (2018). http://eprint.iacr.org/2018/391

[GST04] Goodrich, M.T., Sun, J.Z., Tamassia, R.: Efficient Tree-based revocation in groups of low-state devices. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 511–527. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28628-8_31

[GSW00] Garay, J.A., Staddon, J., Wool, A.: Long-lived broadcast encryption. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 333–352. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44598-6_21

[GSY99] Gafni, E., Staddon, J., Yin, Y.L.: Efficient methods for integrating traceability and broadcast encryption. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 372–387. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48405-1_24

[Gui13] Guillevic, A.: Comparing the pairing efficiency over composite-order and prime-order elliptic curves. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 2013. LNCS, vol. 7954, pp. 357–372. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38980-1_22

[GW09] Gentry, C., Waters, B.: Adaptive security in broadcast encryption systems (with short ciphertexts). In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 171–188. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01001-9_10

[HJ12] Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 590–607. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_35

[HS02] Halevy, D., Shamir, A.: The LSD broadcast encryption scheme. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 47–60. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_4

[Lew12] Lewko, A.: Tools for simulating features of composite order bilinear groups in the prime order setting. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 318–335. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_20

[LL15] Lee, K., Lee, D.H.: Adaptively secure broadcast encryption under standard assumptions with better efficiency. IET Inf. Secur. **9**, 149–157 (2015)

[LW10] Lewko, A., Waters, B.: New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 455–479. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11799-2_27

[NNL01] Naor, D., Naor, M., Lotspiech, J.: Revocation and tracing schemes for stateless receivers. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 41–62. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44647-8_3

[OT08] Okamoto, T., Takashima, K.: Homomorphic encryption and signatures from vector decomposition. In: Galbraith, S.D., Paterson, K.G. (eds.) Pairing 2008. LNCS, vol. 5209, pp. 57–74. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85538-5_4

[OT09] Okamoto, T., Takashima, K.: Hierarchical predicate encryption for inner-products. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 214–231. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_13

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

[Sv98] Stinson, D.R., van Trung, T.: Some new results on key distribution patterns and broadcast encryption. Des. Codes Cryptograph. **14**(3), 261–279 (1998)

[Wat09] Waters, B.: Dual system encryption: realizing fully secure IBE and HIBE under simple assumptions. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03356-8_36

[Wee14] Wee, H.: Dual system encryption via predicate encodings. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 616–637. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54242-8_26

[Wee16] Wee, H.: Déjà Q: encore! Un petit IBE. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016. LNCS, vol. 9563, pp. 237–258. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49099-0_9