The n-Diffie-Hellman Problem and Its Applications

Liqun Chen[1] and Yu Chen[2,3]

[1] Hewlett-Packard Laboratories, Bristol, UK
liqun.chen@hp.com
[2] School of Computer Science, Peking University, Beijing, China
[3] Institute of Information Engineering, Chinese Academy of Sciences
cycosmic@gmail.com

Abstract. The main contributions of this paper are twofold. On the one hand, the twin Diffie-Hellman (twin DH) problem proposed by Cash, Kiltz and Shoup is extended to the n-Diffie-Hellman (n-DH) problem for an arbitrary integer n, and this new problem is shown to be at least as hard as the ordinary DH problem. Like the twin DH problem, the n-DH problem remains hard even in the presence of a decision oracle that recognizes solution to the problem. On the other hand, observe that the double-size key in the Cash et al. twin DH based encryption scheme can be replaced by two separated keys each for one entity, that results in a 2-party encryption scheme which holds the same security feature as the original scheme but removes the key redundancy. This idea is further extended to an n-party case, which is also known as n-out-of-n encryption. As examples, a variant of ElGamal encryption and a variant of Boneh-Franklin IBE have been presented; both of them have proved to be CCA secure under the computational DH assumption and the computational bilinear Diffie-Hellman (BDH) assumption respectively, in the random oracle model. The two schemes are efficient, due partially to the size of their ciphertext, which is independent to the value n.

Keywords: the (strong) n-DH assumption, the (strong) n-BDH assumption, multiple public key encryption, multiple identity-based encryption.

1 Introduction

In EUROCRYPT 2008 [6], Cash, Kiltz and Shoup proposed a new computational problem and named it the *twin Diffie-Hellman* (twin DH) problem with the meaning that given a random triple of the form $(X_1, X_2, Y) \in \mathbb{G}^3$ for a cyclic group $\mathbb{G}$, compute $\mathrm{dh}(X_1, Y)$ and $\mathrm{dh}(X_2, Y)$, where dh is the DH function. They also proposed the *strong twin DH* problem, which is the twin DH problem under the condition that an adversary is given access to a corresponding decision twin DH oracle. They proved that the strong twin DH problem is as hard as the (ordinary) DH problem, i.e., given a random pair of the form $(X, Y) \in \mathbb{G}^2$, compute $\mathrm{dh}(X, Y)$.

X. Lai, J. Zhou, and H. Li (Eds.): ISC 2011, LNCS 7001, pp. 119–134, 2011.

The motivation of their introducing the (strong) twin DH problem is the following: it is well-known that there exist many cryptographic constructions (e.g., the Diffie-Hellman non-interactive key exchange protocol [17] and the Cramer-Shoup encryption scheme [13]) which are based on the DH problem, but security of these constructions can only be proved under the strong DH problem, i.e., the adversary is given access to a decision DH oracle. The reason is that in the security proof, the simulator need the help of the decision oracle to keep the simulation coherent throughout the game. By employing the strong twin DH problem in these constructions, they can successfully prove that the modified constructions are secure under the DH problem, since the strong twin DH problem implies the DH problem. This is a clever trick.

However, their method is not cost free. In order to employ the twin DH problem, their modified construction is "a bit less efficient" than the original one; specifically, the modified construction doubles the key of the original one. For example, in their twin Identity-Based Encryption (IBE) scheme, a master key of a Key Generation Center (KGC) is twin private/public key pairs, written as $((x_1, X_1), (x_2, X_2))$, instead of one (x, X) in the original IBE scheme, and accordingly, an user's secret key associated with this user's identity id (served as a public key of the user) is also two secret values written as (S_1, S_2), each of which is computed under one master key pair. Therefore, a key redundancy is the cost of tighter security reduction.

Can we use this key redundancy to achieve some extra useful function without imposing an efficiency penalty? Observe that in their twin IBE scheme, the identity value id in computing S_1 does not have to be the same as in computing S_2; the two private/public master key pairs (x_1, X_1) and (x_2, X_2) can each belong to an individual KGC. With this slight modification, a user can have two independent identities each associated with one KGC. For example, Alice has her working email address associated with her employer as one KGC and her passport number associated with the government of her country as another KGC. These two KGCs are independent authorities, and do not necessarily have any trust relation or communication between them. Furthermore, the number of the identities and KGCs in the IBE scheme may not be restricted to two[1].

This observation leads to the main contributions of our paper that the twin DH problem can be extended to the n-DH problem for an arbitrary number n, which enables us to build an efficient encryption scheme with multiple public keys and an efficient IBE scheme with multiple KGCs and identities. This type of encryption is also known as n-out-of-n encryption, in which a given message is encrypted under a set of n individual public keys, and the associated decryption operation makes use of the n corresponding secret keys. It is relevant to other well-known encryption primitives with multi-receivers, such as broadcast encryption [5, 16] (known as 1-out-of-n encryption) and threshold cryptosystem [15] (known as t-out-of-n encryption). The latter has an attractive

[1] The multi-KGC IBE is not an unsolved problem and could be implemented from extending an existing IBE scheme, but we want to show how we can do it *efficiently* using n-out-of-n encryption.

application, namely attribute-based encryption (ABE) [20, 3]. Compared with the well-explored t-out-of-n threshold encryption or ABE schemes, e.g. using a secret sharing technique [24], an n-out-of-n encryption scheme seems a naive solution. But we think it is worthy studying this solution properly since it has the advantage of simplicity in both algorithm implementation and security analysis.

More specifically, there are a number of contributions in this paper. Here we describe a brief overview of each contribution individually.

THE n-DH PROBLEM. We present a modification of the twin DH problem [6] by extending the number of the (ordinary) DH instances from 2 to an arbitrary integer n, and name it the *n-DH problem*. Intuitively, the n-DH problem is that given a random $n+1$ tuple of the form $(X_1, \ldots, X_n, Y) \in \mathbb{G}^{n+1}$ for a cyclic group $\mathbb{G}$, compute $(\mathrm{dh}(X_1, Y), \ldots, \mathrm{dh}(X_n, Y))$ where dh is the DH function. We also present the *strong n-DH problem* which is the n-DH problem under the condition that an adversary is given access to a corresponding decision n-DH oracle. We prove that the strong n-DH problem is just as hard as the DH problem.

THE n-BDH PROBLEM. We present a modification of the twin Bilinear-DH (twin BDH) problem [6, 12]. by extending the number of the (ordinary) BDH instances from 2 to an arbitrary integer n, and name it the *n-BDH problem*. Intuitively, the n-BDH problem is that given a random $2n+1$ tuple of the form $(X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n) \in \mathbb{G}^{2n+1}$ for a cyclic group $\mathbb{G}$, compute $(\mathrm{bdh}(X_1, Y, Z_1), \ldots, \mathrm{bdh}(X_n, Y, Z_n))$ where bdh is the BDH function. We also present the *strong n-BDH problem* which is the n-BDH problem under the condition that an adversary is given access to a corresponding decision n-BDH oracle. We prove that the strong n-BDH problem is just as hard as the BDH problem.

CONCEPT AND EXAMPLE OF AN MPKE SCHEME. We formalize the concept of an n-out-of-n public key encryption scheme and call it a Multiple Public Key Encryption (MPKE) scheme. MPKE schemes can be used in those applications, which requires that either a decryptor must be in the possession of n private keys (e.g., each can be bound with an particular attribute) or that n decryptors (each with an individual key) must work together, in order to decrypt a given ciphertext. As a concrete MPKE example, we present a new modification of the hashed ElGamal encryption scheme [1], and name it the n-ElGamal encryption scheme. Based on the strong n-DH assumption (that implies based on the ordinary DH assumption), we prove that the n-ElGamal encryption scheme has chosen ciphertext security in the random oracle [2].

CONCEPT AND EXAMPLE OF AN MIBE SCHEME. We formalize the concept of a Multiple Identity-Based Encryption (MIBE) scheme, which is an MPKE scheme with the identity-based key setting under the condition that the n KGCs, each generating a private key from an identity value, can be independent to each other. This type of IBE schemes has already been introduced in the literature, e.g. [7, 10, 11]. To the best of our knowledge, the security of the schemes in [7, 10, 11] have not been rigorously analyzed. As a concrete MIBE example, we present a new modification of the Boneh-Franklin IBE scheme [4] and name it the n-IBE scheme. Based on the strong n-BDH assumption (that implies based

on the ordinary BDH assumption), we prove that the n-IBE scheme has chosen ciphertext security in the random oracle [2].

The rest of this paper is organized as follows. We describe definitions of the (strong) n-BDH assumption in Section 2 and of the (strong) n-BDH assumption in Section 3. After that, we present definitions of security models for MPKE schemes and MIBE schemes in Section 4, followed by a concrete MPKE scheme with a rigorous security analysis in Section 5, and a concrete MIBE scheme in Section 6 (due to the limited space, its rigorous security analysis is in the full paper [8]). We end the paper with conclusions and some open questions for future work in Section 7.

2 The n-DH Assumption

Let $\mathbb{G}$ be a cyclic group of prime order p and with generator g, and let dh be the DH function defined as

$$\mathrm{dh}(X, Y) := Z, \text{ where } X = g^x, Y = g^y \text{ and } Z = g^{xy}.$$

Recall that the DH assumption states it is hard to compute $\mathrm{dh}(X, Y)$ given random $X, Y \in \mathbb{G}$. We define the *n-DH function* function by

$$\mathrm{ndh} : \mathbb{G}^{n+1} \to \mathbb{G}^n, (X_1, \ldots, X_n, Y) \mapsto (\mathrm{dh}(X_1, Y), \ldots, \mathrm{dh}(X_n, Y)).$$

We also define a corresponding *n-DH predicate* by

$$\mathrm{ndhp}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n) := \mathrm{ndh}(X_1, \ldots, X_n, \hat{Y}) \stackrel{?}{=} (\hat{Z}_1, \ldots, \hat{Z}_n).$$

The n-DH assumption states that it is hard to compute $\mathrm{ndh}(X_1, \ldots, X_n, Y)$ given random $X_1, \ldots, X_n, Y \in \mathbb{G}$. Accordingly, the *strong n-DH assumption* states that it is hard to compute $\mathrm{ndh}(X_1, \ldots, X_n, Y)$ given random $X_1, \ldots, X_n, Y \in \mathbb{G}$ along with access to the predicate $\mathrm{ndhp}(X_1, \ldots, X_n, \cdot, \cdot, \ldots, \cdot)$, which returns $\mathrm{ndhp}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n)$ on input $(\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n)$. We have the following theorem to address the relation between the DH assumption and the (strong) n-DH assumption:

Theorem 2.1 (DH via strong n-DH). *The (ordinary) DH assumption holds if and only if the strong n-DH assumption holds.*

It is clear that the DH assumption implies the n-DH assumption. We now prove that the DH assumption implies the strong n-DH assumption. To do this, by following the trapdoor test technique of [6], we first create a trapdoor test.

Theorem 2.2 (Trapdoor Test for n-DH). *Let $\mathbb{G}$ be a cyclic group of prime order p with generator g. Let $I = \{2, \ldots, n\}$, and suppose X_1, r_i, s_i for all $i \in I$ are mutually independent random variables, where X_1 is randomly taken in $\mathbb{G}$, and each of r_i and s_i is uniformly distributed over $\mathbb{Z}_p$, and define the random variables $X_i := g^{s_i}/X_1^{r_i}$. Further suppose that $\hat{Y}, \hat{Z}_1, \cdots, \hat{Z}_n$ are random variables taking values in $\mathbb{G}$, each of which is defined as some function of X_i for all $i \in \{1\} \cup I$. Then we have:*

1. *Each X_i for $i \in I$ is uniformly distributed over $\mathbb{G}$;*
2. *All X_i for $i \in \{1\} \cup I$ are mutually independent;*
3. *If $X_i = g^{x_i}$ for $i \in \{1\} \cup I$, then the probability that the truth value of*

$$\hat{Z_1}^{r_2} \hat{Z}_2 = \hat{Y}^{s_2} \wedge \cdots \wedge \hat{Z_1}^{r_i} \hat{Z}_i = \hat{Y}^{s_i} \wedge \cdots \wedge \hat{Z_1}^{r_n} \hat{Z}_n = \hat{Y}^{s_n} \tag{1}$$

does not agree with the truth value of

$$\hat{Z}_1 = \hat{Y}^{x_1} \wedge \cdots \wedge \hat{Z}_i = \hat{Y}^{x_i} \wedge \cdots \wedge \hat{Z}_n = \hat{Y}^{x_n} \tag{2}$$

is at most $(1/p)^{n-1}$; moreover if (2) holds, then (1) certainly holds.

Proof. Observe that $s_i = r_i x_1 + x_i$ for $i \in I$ where $I = \{2, \ldots, n\}$. It is not difficult to verify that each X_i for $i \in I$ is uniformly distributed over $\mathbb{G}$, and that all X_i for $i \in \{1\} \cup I$ and r_i for $i \in I$ are mutually independent, from which the items 1 and 2 follow. To prove the item 3, condition on fixed values of X_i for $i \in \{1\} \cup I$. In the resulting conditional probability space, each r_i for $i \in I$ is uniformly distributed over $\mathbb{Z}_p$, while all x_i, $\hat{Y}$, $\hat{Z}_i$ for $i \in \{1\} \cup I$ are fixed. If (2) holds, (1) certainly holds, because $s_i = r_i x_1 + x_i$ for $i \in I$. Conversely, if (2) does not hold, we show that (1) holds with probability at most $(1/p)^{n-1}$. We take the $n-1$ equations of (1) separately. Each of them uses the same argument as in the proof of the trapdoor test of [6]. Observe that (1) is equivalent to

$$(\hat{Z_1}/\hat{Y}^{x_1})^{r_2} = \hat{Y}^{x_2}/\hat{Z}_2 \wedge \cdots \wedge (\hat{Z_1}/\hat{Y}^{x_1})^{r_i} = \hat{Y}^{x_i}/\hat{Z}_i \wedge \cdots \wedge (\hat{Z_1}/\hat{Y}^{x_1})^{r_n} = \hat{Y}^{x_n}/\hat{Z}_n. \tag{3}$$

Let us take a look at the $(i-1)^{th}$ equation of (3). We can see that if $\hat{Z}_1 = \hat{Y}^{x_1}$ and $\hat{Z}_i \neq \hat{Y}^{x_i}$ no matter whether the other equations of (2) holds or not, then this equation certainly does not hold. This leaves us with the case $\hat{Z}_1 \neq \hat{Y}^{x_1}$. In this case, the left hand side of the equation is a random element of $\mathbb{G}$ (since r_i is uniformly distributed over $\mathbb{Z}_p$), but the right hand side is a fixed element of $\mathbb{G}$. So this equation holds with probability $1/p$. (3) holds if and only if $n-1$ different equations all hold. Now, we argue that these $n-1$ equations are mutually independent, because each r_i for $i \in I$ is uniformly distributed over $\mathbb{Z}_p$, therefore, the probability that (3) holds is at most $(1/p)^{n-1}$. □

Using this trapdoor test as a tool, we can prove Theorem 2.1. Let $\mathcal{B}$ be a DH adversary. Denote its advantage by $\mathsf{AdvDH}_{\mathcal{B},\mathbb{G}}$ with the meaning of the probability that $\mathcal{B}$ computes $\mathrm{dh}(X, Y)$, given random $X, Y \in \mathbb{G}$. Let $\mathcal{A}$ be a strong n-DH adversary. Denote its advantage by $\mathsf{AdvnDH}_{\mathcal{A},\mathbb{G}}$ with the meaning of the probability that $\mathcal{A}$ computes $\mathrm{ndh}(X_1, \ldots, X_n, Y)$, given random $X_i, Y \in \mathbb{G}$ for $i \in \{1, \ldots, n\}$, along with access to the predicate $\mathrm{ndhp}(X_1, \ldots, X_n, \cdot, \cdot, \ldots, \cdot)$, which on input $(\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n)$, returns $\mathrm{ndhp}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n)$. Theorem 2.1 is a special case of the following:

Theorem 2.3. *Suppose $\mathcal{A}$ is a strong n-DH adversary that makes at most Q_d queries to its decision oracle, and runs in time at most τ. Then there exists a DH adversary $\mathcal{B}$ with the following properties: $\mathcal{B}$ runs in time at most τ, plus*

the time to perform $O(Q_d \log q)$ group operations and some minor bookkeeping; moreover,

$$\left(1 - \frac{Q_d}{p^{n-1}}\right) \mathsf{AdvnDH}_{\mathcal{A},\mathbb{G}} \leq \mathsf{AdvDH}_{\mathcal{B},\mathbb{G}}.$$

In addition, if $\mathcal{B}$ does not output "failure", then its output is correct with probability at least $1 - Q_d/p^{n-1}$.

Proof. The DH adversary $\mathcal{B}$ works as follows, given a challenge instance (X, Y) of the DH problem. First, $\mathcal{B}$ chooses r_i, $s_i \in \mathbb{Z}_p$ for $i \in I$ and $I = \{2, ..., n\}$ at random, sets $X_1 := X$ and $X_i := g^s / X_1^{r_i}$, and gives $\mathcal{A}$ the challenge instance $(X_1, \ldots, X_n, Y)$. Second, $\mathcal{B}$ processes each decision query $(\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n)$ by testing if

$$\hat{Z}_1^{r_2} \hat{Z}_2 = \hat{Y}^{s_2} \wedge \cdots \wedge \hat{Z}_1^{r_i} \hat{Z}_i = \hat{Y}^{s_i} \wedge \cdots \wedge \hat{Z}_1^{r_n} \hat{Z}_n = \hat{Y}^{s_n}$$

holds. Finally, if and when $\mathcal{A}$ outputs $(Z_1, \ldots, Z_n)$, $\mathcal{B}$ tests if this output is correct by testing whether

$$Z_1^{r_2} Z_2 = Y^{s_2} \wedge \cdots \wedge Z_1^{r_i} Z_i = Y^{s_i} \wedge \cdots \wedge Z_1^{r_n} Z_n = Y^{s_n}$$

holds; if this does not hold, $\mathcal{B}$ outputs "failure", and otherwise, $\mathcal{B}$ outputs Z_1.

Provide the oracle simulation is perfect, adversary $\mathcal{A}$'s view is identical to its view in the real environment. It remains to calculate the accuracy of the trapdoor test. Note that the probability of the trapdoor test returning a wrong decision result for a query is at most $(1/p)^{n-1}$, and this happens at most Q_d times. Therefore the trapdoor test can simulate the decision oracle perfectly with probability at least $1 - Q_d/p^{n-1}$. Theorem 2.3 follows immediately. □

3 The n-BDH Assumption

In groups equipped with a pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ where $\mathbb{G}$ and $\mathbb{G}_T$ are cyclic groups of prime order p and $\mathbb{G}$ is with generate g, we recall that the BDH function is defined as

$$\mathrm{bdh}(X, Y, Z) := W, \text{ where } X = g^x, Y = g^y, Z = g^z, \text{ and } W = e(g, g)^{xyz}.$$

The BDH assumption states that computing $\mathrm{bdh}(X, Y, Z)$ for random $X, Y, Z \in \mathbb{G}$ is a hard problem. The strong BDH assumption [21] states that the BDH problem remains hard even with the help of a corresponding decision oracle.

Note that for the purpose of describing our main results as simply as possible, without loss of the generality, we make use of symmetric pairings (also called Type-1 pairings). It does not mean that our proposed assumptions and schemes only work with symmetric pairings. Without changing the main results of this paper, this symmetric pairing representation can be modified to the asymmetric pairing one (i.e., $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ where $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ are cyclic groups of prime order p). More specifically, one may use Type-2 pairings, where there is

an efficiently computable group isomorphism $\psi : \mathbb{G}_2 \to \mathbb{G}_1$ mapping $g_2 \in \mathbb{G}_2$ to $g_1 \in \mathbb{G}_1$, or Type-3 pairings, where there is no known efficiently computable group isomorphism $\psi : \mathbb{G}_2 \to \mathbb{G}_2$ mapping g_2 to g_1. We refer readers to [19] for the details of these three types of pairings.

We define the *n-BDH* function by

$$\mathrm{nbdh} : \mathbb{G}_n \to \mathbb{G}_T^n,$$

$$(X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n) \mapsto (\mathrm{bdh}(X_1, Y, Z_1), \ldots, \mathrm{bdh}(X_n, Y, Z_n)).$$

We also define a corresponding *n-BDH predicate* by

$$\mathrm{nbdhp}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n, \hat{W}_1, \ldots, \hat{W}_n) :=$$
$$\mathrm{nbdh}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n) \stackrel{?}{=} (\hat{W}_1, \ldots, \hat{W}_n).$$

The *n-BDH assumption* states that it is hard to compute $\mathrm{nbdh}(X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n)$ given random $X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n \in \mathbb{G}$. The *strong n-BDH assumption* states that it is hard to compute $\mathrm{nbdh}(X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n)$, given random $X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n \in \mathbb{G}$, along with the access to the predicate $\mathrm{nbdh}(X_1, \ldots, X_n, \cdot, \cdot, \ldots, \cdot, \cdot, \ldots, \cdot)$, which on input $(\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n, \hat{W}_1, \ldots, \hat{W}_n)$, returns $\mathrm{nbdhp}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n, \hat{W}_1, \ldots, \hat{W}_n)$.

We have the following result to address the relation between the BDH assumption and the (strong) n-BDH assumption:

Theorem 3.1 (BDH via strong n-BDH). *The (ordinary) BDH assumption holds if and only if the strong n-BDH assumption holds.*

It is clear that the BDH assumption implies the n-BDH assumption. We prove that the BDH assumption implies the strong n-BDH assumption. Again, by following the technique developed in [6], we first create a trapdoor test.

Theorem 3.2 (Trapdoor Test for n-BDH). *Let $\mathbb{G}$ be a cyclic group of prime order p with a generator g and a pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$, where $\mathbb{G}_T$ is another cyclic group of order p. Let $I = \{2, \ldots, n\}$, and suppose X_1, r_i, s_i for $i \in I$ are all mutually independent random variables, where X_1 is randomly taken in $\mathbb{G}$, and each of r_i and s_i is uniformly distributed over $\mathbb{Z}_p$, and define the random variables $X_i := g^{s_i}/X_1^{r_i}$ for $i \in I$. Further suppose that $(\hat{Y}_1, \ldots, \hat{Y}_n, \hat{Z}, \hat{W}_i, \ldots, \hat{W}_n)$ are random variables taking values in $\mathbb{G}$, each of which is defined as some function of X_i for all $i \in \{1\} \cup I$. Then we have:*

1. *Each X_i for $i \in I$ is uniformly distributed over $\mathbb{G}$;*
2. *All X_i for $i \in \{1\} \cup I$ are mutually independent;*
3. *If $X_i = g^{x_i}$ for $i \in \{1\} \cup I$, the probability that the truth value of*

$$\hat{W}_1^{\,r_2}\hat{W}_2 = e(\hat{Y}_2, \hat{Z})^{s_2} \wedge \cdots \wedge \hat{W}_1^{\,r_i}\hat{W}_i = e(\hat{Y}_i, \hat{Z})^{s_i} \wedge \cdots \wedge \hat{W}_1^{\,r_n}\hat{W}_n = e(\hat{Y}_n, \hat{Z})^{s_n} \tag{4}$$

does not agree with the truth value of

$$\hat{W}_1 = e(\hat{Y}_1, \hat{Z})^{x_1} \wedge \cdots \wedge \hat{W}_i = e(\hat{Y}_i, \hat{Z})^{x_i} \wedge \cdots \wedge \hat{W}_n = e(\hat{Y}_n, \hat{Z})^{x_n} \tag{5}$$

is at most $(1/p)^{n-1}$; moreover if (5) holds, then (4) certainly holds.

The proof of this theorem is similar to the proof of Theorem 2.2. Due to the limited space, we have put this proof in the full paper [8].
Using this trapdoor test as a tool, we can prove Theorem 3.1. Let $\mathcal{B}$ be a BDH adversary. Denote its BDH advantage by $\mathsf{AdvBDH}_{\mathcal{B},\mathbb{G}}$ with the meaning of the probability that $\mathcal{B}$ computes $\mathrm{bdh}(X, Y, Z)$, given random $X, Y, Z \in \mathbb{G}$. Let $\mathcal{A}$ be a strong nbdh adversary. Denote its advantage by $\mathsf{AdvnBDH}_{\mathcal{A},\mathbb{G}}$ with the meaning of the probability that $\mathcal{A}$ computes $\mathrm{ndh}(X_1, \ldots, X_n, Y, Z_1, \ldots, Z_n)$, given random $X_i, Y, Z_i \in \mathbb{G}$ for $i \in \{1, \ldots, n\}$, along with access to a decision oracle for the predicate $\mathrm{nbdhp}(X_1, \ldots, X_n, \cdot, \cdot, \ldots, \cdot, \cdot, \ldots, \cdot)$, which on input $(\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n, \hat{W}_1, \ldots, \hat{W}_n)$, returns $\mathrm{nbdhp}(X_1, \ldots, X_n, \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n, \hat{W}_1, \ldots, \hat{W}_n)$. Theorem 3.1 is a special case of the following:

Theorem 3.3. *Suppose $\mathcal{A}$ is a strong n-BDH adversary that makes at most Q_d queries to its decision oracle, and runs in time at most τ. Then there exists a BDH adversary $\mathcal{B}$ with the following properties: $\mathcal{B}$ runs in time at most τ, plus the time to perform $O(Q_d \log q)$ group operations and some minor bookkeeping; moreover,*

$$\left(1 - \frac{Q_d}{p^{n-1}}\right) \mathsf{AdvnBDH}_{\mathcal{A},\mathbb{G}} \leq \mathsf{AdvBDH}_{\mathcal{B},\mathbb{G}}.$$

In addition, if $\mathcal{B}$ does not output "failure", then its output is correct with probability at least $1 - Q_d/p^{n-1}$.

The proof of this theorem is similar to the proof of Theorem 2.3. Again, due to the limited space, we have put this proof in the full paper [8].

4 Definitions of MPKE and MIBE

In this section we present formal definitions of a Multiple Public Key Encryption (MPKE) scheme and of a Multiple Indentity-Based Encryption (MIBE) scheme, including their security notion: chosen ciphertext security, which are based on the usual definitions of chosen ciphertext security for a public key encryption scheme [22] and an identity-based encryption scheme [4]. Recall that these two types of encryption schemes are n-out-of-n encryption schemes. In the security model an adversary is not allowed to corrupt any decryption key from the entirely n set of the keys.

4.1 Multiple Public Key Encryption

A Multiple Public Key Encryption scheme (say MPKE), with a security parameter 1^κ and associated system parameters params (include descriptions of a finite key space $\mathcal{K}$, a finite message space $\mathcal{M}$, and a finite ciphertext space $\mathcal{C}$), is specified by three algorithms: KeyGen, Encrypt, and Decrypt:
KeyGen: takes 1^κ and params as input, and generates a set n of public and secret key pairs, written as $(pk_i, sk_i) \in \mathcal{K}$ for $i = 1, \ldots, n$. We also denote the n public keys by $\mathbf{pk} = (pk_1, \ldots, pk_n)$ and the n secret keys by $\mathbf{sk} = (sk_1, \ldots, sk_n)$.

Encrypt: takes as input params, $\mathbf{pk}$, and a message $M \in \mathcal{M}$. It returns a ciphertext $C \in \mathcal{C}$.

Decrypt: takes as input params, a ciphertext $C \in \mathcal{C}$ and $\mathbf{sk}$, and returns M.

These algorithms must satisfy the standard consistency constraint, namely when $(\mathbf{pk}, \mathbf{sk}) \leftarrow \mathsf{KeyGen}(1^\kappa, \mathsf{params})$, then

$$\forall M \in \mathcal{M} : \mathsf{Decrypt}(\mathsf{params}, C, \mathbf{sk}) = M \text{ where } C = \mathsf{Encrypt}(\mathsf{params}, \mathbf{pk}, M).$$

Chosen ciphertext security of the scheme MPKE is defined by the following attack game, played between a challenger $\mathcal{CH}$ and an adversary $\mathcal{A}$:

Setup. The challenger takes a security parameter 1^κ and associated params, and runs the KeyGen algorithm. It gives the resulting $\mathbf{pk}$ together with params to $\mathcal{A}$, and keeps the corresponding $\mathbf{sk}$ to itself.

Phase 1. $\mathcal{A}$ makes a number of decryption queries to the challenger, where the input to each query is a ciphertext, say $\hat{C}$. To answer such a query, the challenger decrypts $\hat{C}$ and sends the result to $\mathcal{A}$. These queries may be asked adaptively, that is, each query may depend on the replies to previous queries.

Challenge. Once the adversary decides that Phase 1 is over, it outputs two equal length plaintexts $M_0, M_1 \in \mathcal{M}$ on which it wishes to be challenged. The challenger picks a random bit $\beta \in \{0, 1\}$, encrypts M_β, and sends the resulting ciphertext C^* as the challenge to $\mathcal{A}$.

Phase 2. $\mathcal{A}$ issues more decryption queries as in Phase 1, but with the restriction that $\hat{C} \neq C^*$. These queries may be asked adaptively as in Phase 1.

Guess. Finally, $\mathcal{A}$ outputs a guess $\beta' \in \{0, 1\}$ and wins the game if $\beta = \beta'$.

We refer to such an adversary $\mathcal{A}$ as an IND-CCA adversary. We define adversary $\mathcal{A}$'s advantage over the scheme MPKE by $\mathsf{AdvCCA}_{\mathcal{A},\mathsf{MPKE}}(\kappa) = \left|\Pr[\beta = \beta'] - \frac{1}{2}\right|$. The probability is over the random bits used by the challenger and the adversary.

Definition 4.1. *We say that a multiple public key encryption scheme* MPKE *is* IND-CCA *secure if for any probabilistic polynomial time* IND-CCA *adversary* $\mathcal{A}$ *the advantage* $\mathsf{AdvCCA}_{\mathcal{A},\mathsf{MPKE}}(\kappa)$ *is negligible*[2]*.*

When we analyze the scheme MPKE in the random oracle model, then hash functions are modeled as random oracles, and both the challenger and adversary are given access to the random oracles in the above attack game. In that case, we write $\mathsf{AdvCCA}^{\mathrm{ro}}_{\mathcal{A},\mathsf{MPKE}}(\kappa)$ for the corresponding advantage.

4.2 Multiple Identity-Based Encryption

A Multiple Identity-Based Encryption scheme, denoted by MIBE, is specified by four algorithms: Setup, Extract, Encrypt and Decrypt:

[2] We say that a function $f(\kappa)$ is negligible if for every $c > 0$ there exists a value κ_c such that $f(\kappa) < 1/\kappa^c$ for all $\kappa < \kappa_c$.

Setup: takes a security parameter 1^κ, and returns system parameters params and a set n of master public and secret key pairs, written as (mpk_i, msk_i) for $i = 1, \ldots, n$; without loss of generality, each key pair (mpk_i, msk_i) is associated with the i-th of a set n KGCs. We denote the n master public keys by $\mathbf{mpk} = (mpk_1, \ldots, mpk_n)$ and the n master secret keys by $\mathbf{msk} = (msk_1, \ldots, msk_n)$. The parameters params include a description of a finite message space $\mathcal{M}$, and a description of a finite ciphertext space $\mathcal{C}$.

Extract: takes as input params, a master key msk_i and an arbitrary identity $id_i \in \{0,1\}^*$ for $i \in \{1, \ldots, n\}$. It returns a secret key sk_i. By repeating the Extract algorithm n times with different i values, one can obtain $\mathbf{sk} = (sk_1, \ldots, sk_n)$ associated with $\mathbf{id} = (id_1, \ldots, id_n)$. Note that msk_i and id_i do not have to uniquely match to each other. Theoretically speaking, any arbitrary identity can bind with any master key, and therefore, the case $id_i = id_j$ for $i \neq j$ is allowed.

Encrypt: takes as input params, **pk**, **id** and a message $M \in \mathcal{M}$. It returns a ciphertext $C \in \mathcal{C}$.

Decrypt: takes as input params, a ciphertext $C \in \mathcal{C}$ and **sk**, and returns M.

These algorithms must satisfy the standard consistency constraint, namely when $(\mathbf{mpk}, \mathbf{msk}, \mathsf{params}) \leftarrow \mathsf{Setup}(1^\kappa)$ and $\mathbf{sk} \leftarrow \mathsf{Extract}(\mathsf{params}, \mathbf{msk}, \mathbf{id})$, then

$$\forall m \in \mathcal{M} : \mathsf{Decrypt}(\mathsf{params}, C, \mathbf{sk}) = M \text{ where } C = \mathsf{Encrypt}(\mathsf{params}, \mathbf{mpk}, \mathbf{id}, M).$$

Chosen ciphertext security of scheme MIBE is defined by the following attack game, played between a challenger $\mathcal{CH}$ and an adversary $\mathcal{A}$:

Setup. The challenger runs the Setup algorithm. It gives the adversary the resulting params and **mpk**, and keeps the associated **msk** to itself.

Phase 1. The adversary issues queries $q_1, \ldots, q_m$ where query q_i is one of:

- Extraction query $\langle i, \hat{id}_i \rangle$. The challenger responds by running algorithm Extract to generate the private key $\hat{sk}_i$ associated with $\hat{id}_i$ and msk_i. It sends $\hat{sk}_i$ to $\mathcal{A}$.
- Decryption query $\langle \hat{\mathbf{id}}, \hat{C} \rangle$. The challenger responds by running algorithm Extract n times to generate the private key $\hat{\mathbf{sk}}$ corresponding to $\hat{\mathbf{id}}$. It then runs algorithm Decrypt to decrypt the ciphertext $\hat{C}$. It sends the resulting plaintext to $\mathcal{A}$.

These queries may be asked adaptively, that is, each query q_i may depend on the replies to $q_1, \ldots, q_{i-1}$.

Challenge. Once the adversary decides that Phase 1 is over it outputs two equal length plaintexts $M_0, M_1 \in \mathcal{M}$ and a set of identities $\hat{\mathbf{id}}^*$ on which it wishes to be challenged. The only constraint is that each element id_i^* of $\hat{\mathbf{id}}^*$ did not appear in any private key extraction query associated with msk_i in Phase 1. The challenger picks a random bit $\beta \in \{0,1\}$ and set $C^* = \mathsf{Encrypt}(\mathsf{params}, \mathbf{mpk}, \hat{\mathbf{id}}^*, M_\beta)$. It sends C^* as the challenge to the adversary.

Phase 2. The adversary issues more queries $q_{m+1}, \ldots, q_r$ where q_i is one of:

– Extraction query $\langle i, \hat{id}_i \rangle$, where $\hat{id}_i \neq$ the i-th element of $\hat{\mathbf{id}}^*$. Challenger responds as in Phase 1.

– Decryption query $\langle \hat{\mathbf{id}}, \hat{C} \rangle \neq \langle \hat{\mathbf{id}}^*, C^* \rangle$. Challenger responds as in Phase 1.

These queries may be asked adaptively as in Phase 1.

Guess. The adversary outputs a guess $\beta' \in \{0, 1\}$ and wins the game if $\beta = \beta'$.

We refer to such an adversary $\mathcal{A}$ as an IND-ID-CCA adversary. We define $\mathcal{A}$'s advantage over the scheme MIBE by $\mathsf{AdvCCA}_{\mathcal{A},\mathsf{MIBE}}(\kappa) = |\Pr[\beta = \beta'] - \frac{1}{2}|$. The probability is over the random bits used by the challenger and the adversary.

Definition 4.2. *We say that a Multiple IBE scheme* MIBE *is* IND-ID-CCA *secure if for any probabilistic polynomial time* IND-ID-CCA *adversary* $\mathcal{A}$ *the advantage* $\mathsf{AdvCCA}_{\mathcal{A},\mathsf{MIBE}}(\kappa)$ *is negligible.*

When we analyze such a scheme MIBE in the random oracle model, we write $\mathsf{AdvCCA}^{\mathrm{ro}}_{\mathcal{A},\mathsf{MIBE}}(\kappa)$ for the corresponding advantage.

5 The n-ElGamal Encryption Scheme

In this section, we present details of the n-ElGamal encryption scheme. The scheme makes use of a hash function H and a symmetric cipher $\mathsf{SE} = (\mathsf{E}, \mathsf{D})$. Let $\mathbb{G}$ be a cyclic group of prime order p and with generator g. A set of public keys for this scheme is denoted by a n-tuple of random group elements $\mathbf{pk} = (X_1, \ldots, X_n)$, with a set of corresponding secret keys denoted by $\mathbf{sk} = (x_1, \ldots, x_n)$, where $X_i = g^{x_i}$ for $i \in I$ and $I = (1, \ldots, n)$. To encrypt a message $m \in \mathcal{M}$, one chooses a random $y \in \mathbb{Z}_p$, and computes

$$Y := g^y, Z_i := X_i^y \text{ for } i \in I, k := \mathsf{H}(Y, Z_1, \ldots, Z_n), C := \mathsf{E}(k, M).$$

The ciphertext is (Y, c). Decryption works accordingly: given (Y, c) and secret key $\mathbf{sk}$, one computes

$$Z_i := Y^{x_i} \text{ for } i \in I, k := \mathsf{H}(Y, Z_1, \ldots, Z_n), M := \mathsf{D}(k, C).$$

As mentioned earlier, the size of the ciphertext in this scheme is independent to the number of public and secret keys n. Like the twin ElGamal encryption scheme [6], the scheme does not add redundancy in the ciphertext in order to achieve CCA security, as in the Fujisaki-Okamoto transformation [18]. Following the arguments in [1, 6, 14], we now show that the n-ElGamal encryption scheme is secure against chosen ciphertext attack, under the strong n-DH assumption. By Theorem 2.1, the same holds under the (ordinary) DH assumption. Formally speaking, we denote the n-ElGamal encryption scheme $\mathsf{MPKE}_{\mathrm{ndh}}$, and analyze security of this scheme with the following theorem, under the security model previously defined in Section 4.1.

Theorem 5.1. *Suppose* H *is modeled as a random oracle,* SE *is secure against chosen ciphertext attack, and the DH assumption holds in* $\mathbb{G}$. *The* $\mathsf{MPKE}_{\mathrm{ndh}}$ *is*

secure against chosen ciphertext attack. In particular, suppose $\mathcal{A}$ is an adversary that carries out a chosen ciphertext attack against $\mathsf{MPKE}_{\mathrm{ndh}}$ in the random oracle model, and $\mathcal{A}$ runs in time τ, and makes at most Q_h hash queries and Q_d decryption queries. Then there exists an adversary $\mathcal{B}_{\mathrm{dh}}$ against the DH problem and an adversary $\mathcal{B}_{\mathrm{sym}}$ against the chosen ciphertext security of SE, such that both $\mathcal{B}_{\mathrm{dh}}$ and $\mathcal{B}_{\mathrm{sym}}$ run in time at most τ, plus the time to perform $O((Q_h + Q_d)\log p)$ group operations; moreover,

$$\mathsf{AdvCCA}^{\mathrm{ro}}_{\mathcal{A},\mathsf{MPKE}_{\mathrm{ndh}}} \leq \left(\frac{p^{n-1}}{p^{n-1}-Q_h}\right)\mathsf{AdvDH}_{\mathcal{B}_{\mathrm{dh}},\mathbb{G}} + \mathsf{AdvCCA}_{\mathcal{B}_{\mathrm{sym}},\mathsf{SE}}.$$

Proof. We proceed with a sequence of games.

Game 0. This is the original chosen ciphertext attack game for a MPKE scheme as defined in Section 4.1. Let S_0 be the event that $\beta' = \beta$ in this game.

Setup: To start the game, the challenger generates the secret key set $\mathbf{sk} = (x_1, \ldots, x_n)$ and their corresponding public key set $\mathbf{pk} = (X_1, \ldots, X_n)$. The challenger gives $\mathbf{pk}$ to the adversary.

Hash oracle query $\langle \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n \rangle$: The challenger maintains a list of tuples $(Y, Z_1, \ldots, Z_n, k)$ as explained below. We refer to this list as the L list, which is initially empty and indexed by elements of $\mathbb{G}^{n+1}$. Whenever the adversary makes a query $\langle \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n \rangle$, if there is already a tuple on the L list indexed by it then the challenger responds with $L[\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n] = \hat{k}$. Otherwise, the challenger picks a random symmetric key $\hat{k}$, adds the tuple $\langle \hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n, \hat{k} \rangle$ to the L list and responds the adversary with $\hat{k}$.

Phase 1 - Decryption query $\langle \hat{Y}, \hat{C} \rangle$: The challenger answers the decryption queries using $\mathbf{sk}$. The challenger need to call the H query in this operation.

Challenge: Once the adversary decides that Phase 1 is over it outputs two messages M_0, M_1 on which it wishes to be challenged. The challenger chooses a random $y \in \mathbb{Z}_p$, sets $Y := g^y$, $Z_i = X_i^y$ for $i = 1, \ldots, n$, then fetches the symmetric key k by querying H with $\langle Y, Z_1, \ldots Z_n \rangle$, and computes $c := \mathsf{E}_k(M_\beta)$, and returns the ciphertext (Y, C) to $\mathcal{A}$.

Phase 2. The decryption queries in Phase 2 are processed just as in Phase 1.

Guess: The adversary $\mathcal{A}$ outputs its guess β' for β.

That finishes the description of Game 0. Despite the syntactic difference, it is clear that

$$\mathsf{AdvCCA}^{\mathrm{ro}}_{\mathcal{A},\mathsf{MPKE}_{\mathrm{ndh}}} = |\Pr[S_0] - 1/2|. \tag{6}$$

Game 1. We now describe Game 1, which is the same as Game 0, but with the following difference: the challenger will abort the game if the adversary query H at $\langle Y, Z_1, \ldots, Z_n \rangle$ either in Phase 1 or Phase 2. Everything else remains exactly the same as Game 0. Let S_1 be the event that $\beta' = \beta$ in Game 1 and F be the event that the adversary queries the random oracle at $\langle Y, Z_1, \ldots Z_n \rangle$ in Game 1. Since Game 0 and Game 1 proceed identically unless F occurs, we have

$$|\Pr[S_1] - \Pr[S_0]| \leq \Pr[F]. \tag{7}$$

We claim that

$$\Pr[F] \leq \mathsf{AdvnDH}_{\mathcal{B}_{\mathrm{ndh},\mathbb{G}}}, \tag{8}$$

where $\mathcal{B}_{\mathrm{ndh}}$ is an efficient strong n-DH adversary that makes at most Q_h decison oracle queries. Next we detail how $\mathcal{B}_{\mathrm{ndh}}$ plays the role of the challenger in Game 1 to gain the advantage as claimed.

Setup: $\mathcal{B}_{\mathrm{ndh}}$ is given $(X_1, \ldots, X_n, Y)$ as the n-DH challenge instance. $\mathcal{B}_{\mathrm{ndh}}$ gives the adversary $\mathbf{pk} = (X_1, \ldots, X_n)$. Note that the only difference between $\mathcal{B}_{\mathrm{ndh}}$ and the challenger in Game 1 is that the former does not know the $\mathbf{sk} = (x_1, \ldots, x_n)$.

Hash oracle queries: Except processes the queries the same way as the challenger does in Game 1, for every random oracle query $(\hat{Y}, \hat{Z}_1, \ldots, \hat{Z}_n)$, $\mathcal{B}_{\mathrm{ndh}}$ sends this tuple to its own decision oracle, and marks it "good" or "bad" accordingly.

Phase 1 - Decryption queries: $\mathcal{B}_{\mathrm{ndh}}$ can process the decryption queries without using the secret key: given a ciphertext $(\hat{Y}, \hat{c})$, it checks if it has already seen a "good" tuple of the form $(\hat{Y}, \cdot, \ldots, \cdot)$ in L; if so, it uses the key associated with that tuple; if not, it generates a random key, and it will stay on the lookout for a "good" tuple of the form $(\hat{Y}, \cdot, \ldots, \cdot)$ in future random oracle queries, associating this key with that tuple to keep things consistent.

Challenge: Once the adversary decides that Phase 1 is over it outputs two messages M_0, M_1 on which it wishes to be challenged. $\mathcal{B}_{\mathrm{ndh}}$ checks if there is a "good" tuple of the form $(Y, \cdot, \ldots, \cdot)$, if so, it aborts; if not, it generates a random key k (it will stay on the lookout for a "good" tuple of the form $(\hat{Y}, \cdot, \ldots, \cdot)$ in future random oracle queries, associating this key with that tuple to keep things consistent), and computes $c := \mathsf{E}_k(M_\beta)$, and returns the ciphertext (Y, c) to $\mathcal{A}$.

Phase 2 - Decryption queries: The decryption queries in Phase 2 are processed just as in Phase 1. If the adversary issues a "good" tuple of the form $(Y, \cdot, \ldots, \cdot)$, $\mathcal{B}_{\mathrm{ndh}}$ aborts.

Guess: The adversary $\mathcal{A}$ outputs its guess β' for β.

At the end of the game, $\mathcal{B}_{\mathrm{ndh}}$ checks if it has seen a "good" tuple of the form $(Y, \cdot, \ldots, \cdot)$; if so, it outputs the last n components. According to the definition of event F, Equation (8) follows immediately. Theorem 2.3 gives us an efficient DH adversary $\mathcal{B}_{\mathrm{dh}}$ with

$$\mathsf{AdvnDH}_{\mathcal{B}_{\mathrm{ndh},\mathbb{G}}} \leq \frac{p^{n-1}}{p^{n-1} - Q_h} \mathsf{AdvDH}_{\mathcal{B}_{\mathrm{dh},\mathbb{G}}}.$$

Finally, it is easy to see that in Game 1, the adversary is essentially playing the chosen ciphertext attack game against SE. Thus, there is an efficient adversary $\mathcal{B}_{\mathrm{sym}}$ such that

$$|\Pr[S_1] - 1/2| = \mathsf{AdvCCA}_{\mathcal{B}_{\mathrm{sym}},\mathsf{SE}}. \tag{9}$$

Theorem 5.1 now follows by combining (6)-(9). □

6 The n-IBE Scheme

We now present details of the n-IBE scheme. Let $\mathbb{G}$ and $\mathbb{G}_T$ be two cyclic groups of prime order p and $\mathbb{G}$ with generator g, and further let the two groups be equipped with a pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. A master public key set is a tuple of n group elements $\mathbf{mpk} = (X_1, \ldots, X_n)$, where $X_i = g^{x_i}$ for $i \in I$ and $I = \{1, \ldots, n\}$. The corresponding master private key set is a tuple $\mathbf{msk} = (x_1, \ldots, x_n)$, which are selected at random from $\mathbb{Z}_p$. We treat the secret/public master key set $(\mathbf{msk}, \mathbf{mpk})$ as n separate key pairs $(x_1, X_1), \ldots, (x_n, X_n)$, which belong to n Key Generation Centers (KGCs) respectively. This scheme uses a symmetric cipher $\mathsf{SE} = (\mathsf{E}, \mathsf{D})$ and two hash functions H and G, where G is defined as $\mathbb{G} \times \{0,1\}^* \rightarrow \mathbb{G}$, and H is defined as $(\{0,1\}^*)^n \times \mathbb{G} \times \mathbb{G}_T^n \times \rightarrow \{0,1\}^\lambda$ (λ is the length of a symmetric key in algorithm SE).

A private key set associated with n individual identities, denoted by $\mathbf{id} = (id_1, \ldots, id_n)$ for $id_i \in \{0,1\}^*$ and $i \in I$, is a tuple of n group elements $\mathbf{sk} = (S_1, \ldots, S_n)$. The i-th element of $\mathbf{sk}$ is $S_i = \mathsf{G}(X_i, id_i)^{x_i}$. To encrypt a message $M \in \mathcal{M}$ for $\mathbf{id}$, one chooses $y \in \mathbb{Z}_p$ at random and sets

$$Y := g^y, W_i := e(\mathsf{G}(X_i, id_i), X_i)^y \text{ for } i \in I,$$

$$k := \mathsf{H}(id_1, \ldots, id_n, Y, W_1, \ldots, W_n), C := \mathsf{E}(k, M).$$

The ciphertext is (Y, C). To decrypt using $\mathbf{sk}$ for $\mathbf{id}$, one computes

$$W_i := e(S_i, Y) \text{ for } i \in I, k := \mathsf{H}(id_1, \ldots, id_n, Y, W_1, \ldots, W_n), M := \mathsf{D}(k, C).$$

Similar to the n-ElGamal encryption scheme in Section 5, the length of the ciphertext in the n-IBE scheme is independent to the number of KGCs and identities n. Like the twin IBE scheme of [6], the n-IBE scheme does not add redundancy to the ciphertext as in the Fujisaki-Okamoto transformation [18], which, e.g., is used in the Boneh-Franklin IBE scheme [4] and the Sakai-Kasahara IBE scheme [9, 23]. Now we denote our n-IBE scheme by $\mathsf{MIBE}_{\text{nbdh}}$. It holds chosen ciphertext attack security under the strong n-BDH assumption, as shown in Theorem 6.1. By Theorem 3.1, it also means to be secure under the BDH assumption. The theorem can be proved by following the security analysis approach for the twin IBE scheme in [6] (the approach was originally proposed in [21]). Due to the limited space, we have put this proof in the full paper [8].

Theorem 6.1. *Suppose* H *and* G *are modeled as random oracles. Further, suppose the BDH assumption holds with* $(\mathbb{G}, \mathbb{G}_T, e)$*, and that the symmetric cipher* $\mathsf{SE} = (\mathsf{E}, \mathsf{D})$ *is secure against chosen ciphertext attack. Then* $\mathsf{MIBE}_{\text{nbdh}}$ *is secure against the chosen ciphertext attack. In particular, suppose* $\mathcal{A}$ *is an adversary that carries out a chosen ciphertext attack against* $\mathsf{MIBE}_{\text{nbdh}}$*, and that* $\mathcal{A}$ *runs in time* τ*, and makes at most* Q_h *hash* H *queries,* Q_g *hash* G *queries,* Q_d *decryption queries, and* Q_e *secret key* sk_i *extraction queries associated with* id_i*, where* sk_i *(*id_i*) is an element of* $\mathbf{id}$ *(*$\mathbf{sk}$*). Then there exist a BDH adversary* $\mathcal{B}_{\text{bdh}}$ *and an adversary* $\mathcal{B}_{\text{sym}}$ *against the chosen ciphertext security of* SE*,*

such that both $\mathcal{B}_{\mathrm{bdh}}$ *and* $\mathcal{B}_{\mathrm{sym}}$ *run in time at most* τ*, plus that time to perform* $O((Q_e + Q_h + Q_g + Q_d) \log p)$ *group operations; moreover*[3]

$$\mathsf{AdvCCA}^{\mathrm{ro}}_{\mathcal{A},\mathsf{MIBE}_{\mathrm{nbdh}}} \leq \left(\frac{eQ_e}{n}\right)^n \cdot \left(\frac{q^{n-1}}{q^{n-1} - Q_h} \cdot \mathsf{AdvBDH}_{\mathcal{B}_{\mathrm{bdh},\mathbb{G}}} + \mathsf{AdvCCA}_{\mathcal{B}_{\mathrm{sym},\mathsf{SE}}}\right).$$

7 Conclusions

We have proposed a new computational problem called the n-DH problem, which is an extension of the twin DH problem of [6], and also proposed the associated strong n-DH problem and the (strong) n-BDH problem. We have shown that the strong n-DH (n-BDH) problem is as hard as the ordinary DH (BDH) problem. We have introduced a formal definition of n-out-of-n encryption which has two versions, namely MPKE and MIBE for the conventional public key setting and identity-based key setting respectively. We have also proposed an efficient MPKE (MIBE) scheme and proved it is CCA secure under the DH (BDH) assumption.

In our security model for an MPKE (MIBE) scheme, the adversary is not allowed to corrupt any individual key in the whole set of n keys, which is used in the challenge phase. This security model suits our target applications of multiple key encryption very well, where the decryption process requires that either a decryptor must holds n keys or that n decryptors much work together. However, whether this model can be strengthened and whether there is any practical motivation to any enhancement of the model might be an interesting topic for further investigation. Whether there are other applications which can benefit from the (strong) n-DH/n-BDH problem is another question which could lead to some future research.

References

1. Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (2001)
2. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: The 1st ACM Conference on Computer and Communications Security, pp. 62–73. ACM Press, New York (1993)
3. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: IEEE Symposium on Security and Privacy (SP 2007), pp. 321–334 (2007)
4. Boneh, D., Franklin, M.: Identity-based encryption from the weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
5. 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)
6. Cash, D.M., Kiltz, E., Shoup, V.: The twin Diffie-Hellman problem and applications. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer, Heidelberg (2008)

[3] Here $e \approx 2.71$ is the base of the natural logarithm.

7. Chen, L.: An interpretation of identity-based cryptography. In: Aldini, A., Gorrieri, R. (eds.) FOSAD 2007. LNCS, vol. 4677, pp. 183–208. Springer, Heidelberg (2007)
8. Chen, L., Chen, Y.: The n-Diffie-Hellman problem and its applications, Cryptology ePrint Archive, Report 2011/397 (2011)
9. Chen, L., Cheng, Z.: Security proof of sakai-kasahara's identity-based encryption scheme. In: Smart, N.P. (ed.) Cryptography and Coding 2005. LNCS, vol. 3796, pp. 442–459. Springer, Heidelberg (2005)
10. Chen, L., Harrison, K.: Multiple trusted authorities in identifier based cryptography from pairings on elliptic curves, HP Labs Technical Reports, HPL-2003-48
11. Chen, L., Harrison, K., Soldera, D., Smart, N.: Applications of multiple trust authorities in pairing based cryptosystems. In: Davida, G.I., Frankel, Y., Rees, O. (eds.) InfraSec 2002. LNCS, vol. 2437, pp. 260–275. Springer, Heidelberg (2002)
12. Chen, Y., Chen, L.: Twin bilinear Diffie-Hellman inversion problem and its application. To appear in the Proceedings of the 13th Annual International Conference on Information Security and Cryptology, ICISC 2010 (2010)
13. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
14. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM Journal on Computing 33, 167–226 (2001)
15. Damgård, I., Jurik, M.: A length-flexible threshold cryptosystem with applications. In: Safavi-Naini, R., Seberry, J. (eds.) ACISP 2003. LNCS, vol. 2727, pp. 350–364. Springer, Heidelberg (2003)
16. Delerablée, C., Paillier, P., Pointcheval, D.: Fully Collusion Secure Dynamic Broadcast Encryption with Constant-Size Ciphertexts or Decryption Keys. In: Takagi, T., Okamoto, T., Okamoto, E., Okamoto, T. (eds.) Pairing 2007. LNCS, vol. 4575, pp. 39–59. Springer, Heidelberg (2007)
17. Diffie, W., Hellman, M.E.: New directions in cryptograpgy. IEEE Transactions on Infomation Theory 22(6), 644–654 (1976)
18. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg (1999)
19. Galbraith, S., Paterson, K., Smart, N.P.: Pairings for cryptographers. Discrete Applied Mathematics 156(16), 3113–3121 (2008)
20. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: ACM Conference on Computer and Communications Security, ACM CCS 2006, pp. 89–98. ACM, New York (2006)
21. Libert, B., Quisquater, J.-J.: Identity Based Encryption Without Redundancy. In: Ioannidis, J., Keromytis, A.D., Yung, M. (eds.) ACNS 2005. LNCS, vol. 3531, pp. 285–300. Springer, Heidelberg (2005)
22. Rackoff, C., Simon, D.R.: Non-interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (1992)
23. Sakai, R., Kasahara, M.: ID based cryptosystems with pairing on elliptic curve, Cryptology ePrint Archive, Report 2003/054 (2003)
24. Shamir, A.: How to share a secret. Commun. ACM 22(11), 612–613 (1979)