

On the Relations Between Non-Interactive Key Distribution, Identity-Based Encryption and Trapdoor Discrete Log Groups

Kenneth G. Paterson and Sriramkrishnan Srinivasan

Information Security Group,
Royal Holloway, University of London,
Egham, Surrey TW20 0EX, U.K.
{kenny.paterson,s.srinivasan}@rhul.ac.uk

Abstract. We investigate the relationships between identity-based non-interactive key distribution and identity-based encryption. We provide constructions for these schemes that make use of general trapdoor discrete log groups. We then investigate the schemes that result in two concrete settings, obtaining new, provably secure, near-practical identity-based encryption schemes.

Keywords: Identity-based encryption; identity-based non-interactive key distribution; trapdoor discrete logs.

1 Introduction

The concept of identity-based cryptography (IBC) was first introduced by Shamir [33]. In IBC, arbitrary identifying strings such as e-mail addresses or IP addresses can be used to form public keys for users, with the corresponding private keys being created by a Trusted Authority (TA) who is in possession of a system-wide master secret. Then a party Alice who wishes, for example, to encrypt to a party Bob needs only know Bob's identifier and the system-wide public parameters. This approach eliminates certificates and the associated processing and management overheads from public key cryptography.

Shamir [33] was able to construct an identity-based signature (IBS) scheme, but the first fully secure and efficient identity-based encryption (IBE) schemes were not forthcoming until much later [6, 32]. Since then there has been an explosion of interest in IBC and the related field of cryptography based on pairings. A major open problem is to construct an efficient IBE scheme whose security is *not* based on hardness assumptions in pairing-friendly groups (c.f. [30]). A candidate solution is the Cocks IBE scheme [11], whose security is based on the hardness of the quadratic residuosity problem modulo an RSA composite N . But this scheme suffers from ciphertext expansion. More recently, Boneh *et al.* gave an IBE scheme with security also based on the quadratic residuosity problem [7]. This scheme has short ciphertexts but rather large private keys, and both the encryption and decryption algorithms require non-trivial computational effort.

In a separate strand of work that pre-dates the pairing-based approach to IBE, Maurer and Yacobi [23–25] explored the use of special RSA moduli to construct groups in which knowledge of trapdoor information makes the discrete log problem (DLP) in the group easier to solve than when the trapdoor is not available. They used these groups to construct identity-based non-interactive key distribution (ID-NIKD) schemes. Such a scheme allows two entities who have each received a private key from a TA to compute a shared key without exchanging any messages; this key can then be used to derive a key for symmetric encryption, for example. According to [27], around the same time, Murakami and Kasahara adopted a similar approach as Maurer and Yacobi. Unfortunately,

the schemes of Maurer and Yacobi have a somewhat troubled history, with a series of attacks and fixes having been presented for them [24, 22, 25, 19, 27]. Put simply, none of the schemes in [23–25] are secure in any reasonable security model for ID-NIKD. (However, as we shall see, they can be repaired by appropriate use of hashing.) Moreover, these schemes are not actually IBE schemes. This is because there is a subtle but important distinction between IBE and ID-NIKD: in the former, any party can send Bob an encrypted message knowing only Bob’s identifier and the system parameters, while in the latter, both parties need to be enrolled in the system and in possession of their private keys in order to establish a shared key for encryption.

1.1 Our Contributions

Our first contribution is to provide a new security model for ID-NIKD that accurately captures the security desired from this primitive. Our new model is stronger than the existing model of Dupont and Enge [13] and is inspired by earlier models for (interactive) authenticated key exchange [3, 5, 2, 9]. Since ID-NIKD is itself a very useful cryptographic primitive with a wide range of applications in cryptography (see for example [35, 8, 1]), our new model may be of independent interest. We show that the well-known ID-NIKD scheme of Sakai *et al.* [31] is secure in our model.

Our second contribution is to shed light on the relationship between IBE and ID-NIKD schemes. We give a generic conversion that takes any ID-NIKD scheme that is secure in our model and that satisfies some mild technical conditions, and produces from it an IBE scheme that is secure in the IND-ID-CPA security model of [6]. The conversion itself works in the standard model. Chosen-ciphertext security in the random oracle model can easily be obtained by applying a secondary conversion, for example the conversion of [36]. Our conversion provides a framework within which existing ID-NIKD and IBE schemes can be related. For example, we are able to show how to build an IBE scheme from (a repaired version of) the Maurer-Yacobi ID-NIKD scheme, and how the Boneh-Franklin IBE scheme [6] arises from the ID-NIKD scheme of Sakai *et al.* [31]. Thus these quite different IBE schemes can be seen to arise in a uniform way from a single construction applied to well-known ID-NIKD schemes.

Our third contribution is to show how to construct a secure ID-NIKD scheme from *any* trapdoor discrete log (TDL) group G , under the assumptions that the CDH problem is hard in G and that there exists an efficient hash function $H_1 : \{0, 1\}^* \rightarrow G$. This hash function is modelled as a random oracle in our security analysis. Our construction generalises the approach taken by Maurer and Yacobi [23–25] to arbitrary TDL groups and provably fixes the flaws in these earlier schemes by appropriate use of hashing. Applying our generic ID-NIKD-to-IBE conversion, we also obtain a fully secure IBE scheme from any TDL group, under the same assumptions.

Our final contribution is to investigate the ID-NIKD and IBE schemes that result when these constructions are instantiated in two different candidate TDL groups. The first instantiation yields a provably secure IBE scheme for the RSA-based setting that was first introduced by Maurer and Yacobi. The second uses a TDL group construction based on combining Weil descent and disguising elliptic curves via isogenies that was first sketched in [14] and fully developed in [34]. At an 80-bit security level, the two IBE schemes are on the verge of practicality, the only issue being that private key generation is rather expensive for the TA. For example, for the IBE scheme built from elliptic curves and isogenies, the TA faces a one-time setup cost of about 2^{48} bit operations, after which individual private key extractions are straightforward. This IBE scheme operates on regular elliptic curves defined over $\mathbb{F}_{2^{161}}$, enjoys compact public parameters and private keys, and has fast

encryption and decryption. Thus we obtain a near-practical IBE scheme whose security is not based on hardness assumptions in pairing-friendly groups.

As well as answering (at least in part) a major open problem in IBC, our work gives further motivation, should it be needed, for the construction of practical trapdoor discrete log groups.

1.2 Related Work

Murakami and Kasahara [27] recently revisited the security issues apparent in the Maurer-Yacobi approach to constructing ID-NIKD schemes in the RSA trapdoor discrete log setting. They gave a careful exposition of the options for constructing schemes and emphasized the need for appropriate hashing to avoid the square root attacks that plague the Maurer-Yacobi schemes. However, [27] contains no provable security analysis and does not consider IBE schemes. In fact, we will establish the security of one of the ID-NIKD schemes from [27] in this paper. In [20], Kunihiro *et al.* also revisited the Maurer-Yacobi approach and considered its practicality. However, they only study flawed versions of the Maurer-Yacobi schemes.

We are not aware of previous work pointing out the relationship between ID-NIKD and IBE. Indeed, these two different primitives are sometimes confused in the literature. In the particular case of the ID-NIKD scheme of Sakai *et al.* [31] and the Boneh-Franklin IBE scheme [6], it has been noted by many authors that the same algebraic setting and keying method is used. But, as far as we are aware, their exact relationship has not been clarified until now. The fact that ID-NIKD schemes are known to be useful in constructing other cryptographic primitives is evident from [35, 8, 1], for example.

Several papers [16, 18, 28] sketch how TDL groups can be used to build ID-NIKD or IBE schemes. But to our knowledge, none of the previous work properly formalises this idea, nor provides any security analysis. In this sense some of our constructions and security results can be regarded as formalising folklore using the techniques of provable security. More generally, the utility of TDL groups in cryptography is widely acknowledged – see for example [29, 12]. Teske [34] suggests the use of TDL groups for escrowed encryption applications: each user selects his own public key and provides the relevant trapdoor information to the escrow authority. A virtue is then made of the fact that solving the DLP is *not* easy for the authority – this prevents the agency from decrypting individual user’s communications too easily. In our application of the ideas of [14, 34] to IBE, we use a single TDL group G as part of the public parameters, and then exploit the fact that many discrete logs in G can be extracted relatively cheaply after a significant pre-computation. We note that no mention of ID-based applications is made in [14, 34].

Rivest [30] outlines a construction for IBE based on the existence of trapdoor groups with infeasible inverses (TGIIs) that is attributed to Vaikuntanathan, and asks whether this is the only known sufficient condition for IBE outside of bilinear maps (i.e. pairing-friendly groups). Our work, in particular our construction of IBE from any TDL group in which CDH is hard and for which hashing onto the group is possible, answers this question positively. Moreover, in both of our concrete constructions, taking inverses in the relevant group is easy, showing that the notion of a TDL group is distinct from that of a TGII. Indeed, there do not currently appear to be any good candidates for TGIIs.

Several authors [17, 21] have considered the construction of IBE schemes secure against collusion involving a limited number of participants. The benefit of their approach is that any group in which DDH is hard can be used, but the drawback is that the schemes are not secure in the

accepted security model for IBE [6]. Moreover, the complexity of the encryption and decryption algorithms of the schemes is linear in k , the threshold of cheating participants that can be tolerated.

2 Background and Definitions

In this section, we provide basic definitions needed for the remainder of the paper.

Definition 1 (Trapdoor Discrete Log Generator). A *Trapdoor Discrete Log group generator* (TDL group generator) is defined by a pair algorithms *TDLGen* and *SolveDL*:

- *TDLGen*: An algorithm that takes a security parameter 1^k as input and outputs (G, r, g, T) where G is a (description of a) cyclic group of some order r with generator g and T denotes trapdoor information.
- *SolveDL*: An algorithm which takes as input (G, r, g, T) and a group element h and outputs $a \in \mathbb{Z}_r$ such that $h = g^a$.

Note that we do not insist that the order r of the group G in the above definition be prime. This allows us to handle both RSA and elliptic curve settings. In the RSA setting, r must be kept secret by the party running the *TDLGen* algorithm, and, for technical reasons, we will need to assume that a suitable bound R on the group order is available as part of the description of G . We will see that a large multiple of the modulus N can play the role of R in the RSA setting. When disclosure of r does not impact on security we may include r directly as part of the group description. Note also that we do not insist that *SolveDL* run in time polynomial in k – this allows us to handle situations where private key generation may require time that is super-polynomial or even exponential in k , but that is still small enough for practical application for specific values of k .

Definition 2. We define the advantage of an algorithm \mathcal{A} in solving the Computational Diffie-Hellman (CDH) problem in the TDL group G with generator g to be

$$Adv_{\mathcal{A}}^{CDH}(k) = \Pr(\mathcal{A}(G, g, g^a, g^b) = g^{ab})$$

where $a, b \xleftarrow{\$} \mathbb{Z}_r$.

Definition 3. We define the advantage of an algorithm \mathcal{A} in solving the Decisional Diffie-Hellman (DDH) problem in the TDL group G with generator g to be

$$Adv_{\mathcal{A}}^{DDH}(k) = \frac{1}{2} \left| \Pr(\mathcal{A}(G, g, g^a, g^b, g^{ab}) = 1) - \Pr(\mathcal{A}(G, g, g^a, g^b, Z) = 1) \right|$$

where $a, b \xleftarrow{\$} \mathbb{Z}_r$ and $Z \xleftarrow{\$} G$.

We also want to work with pairing-friendly groups, and introduce a second generator for this setting:

Definition 4. A *pairing-friendly group generator* *PairingGen* is a polynomial time algorithm with input 1^k and output a tuple (G, G_T, e, q, P) . Here G, G_T are groups of prime order q , P generates G , and $e : G \times G \rightarrow G_T$ is a bilinear, non-degenerate and efficiently computable map. By convention, G is an additive group and G_T multiplicative.

For ease of presentation, we work exclusively in the setting where e is symmetric; our definitions and results can be generalised to the asymmetric setting where $e : G_1 \times G_2 \rightarrow G_T$, with G_1 and G_2 being different groups. Further details concerning the basic choices that are available when using pairings in cryptography can be found in [15].

Definition 5. We define the advantage of an algorithm \mathcal{A} in solving the Bilinear Diffie-Hellman (BDH) problem in (G, G_T) to be

$$\text{Adv}_{\mathcal{A}}^{\text{BDH}}(k) = \Pr(\mathcal{A}(aP, bP, cP) = e(P, P)^{abc})$$

where $a, b, c \xleftarrow{\$} \mathbb{Z}_q$. Here, we implicitly assume that parameters (G, G_T, e, q, P) are given to \mathcal{A} as additional inputs.

Definition 6. We define the advantage of an algorithm \mathcal{A} in solving the Decisional Bilinear Diffie-Hellman (DBDH) problem in (G, G_T) to be:

$$\text{Adv}_{\mathcal{A}}^{\text{DBDH}}(k) = \frac{1}{2} \left| \Pr(\mathcal{A}(aP, bP, cP, e(P, P)^{abc}) = 1) - \Pr(\mathcal{A}(aP, bP, cP, Z) = 1) \right|$$

where $a, b, c \xleftarrow{\$} \mathbb{Z}_q$ and $Z \xleftarrow{\$} G_T$. Again, we assume that \mathcal{A} has the expected additional inputs.

Definition 7. A function $\epsilon(k)$ is said to be negligible if, for every c , there exists k_c such that $\epsilon(k) \leq k^{-c}$ for every $k \geq k_c$.

Definition 8. Consider a family of functions $f = (f_k)_{k \in \mathbb{N}}$ where $f_k : \{0, 1\}^* \rightarrow \mathcal{T}_k$ for some set \mathcal{T}_k . We define the advantage of an algorithm \mathcal{A} in breaking the one-wayness of function f_k to be:

$$\text{Adv}_{\mathcal{A}}^{\text{OW}}(k) = \Pr(f_k(x) = y : x' \xleftarrow{\$} \{0, 1\}^*, y = f_k(x'), x \leftarrow \mathcal{A}(y)).$$

Here, \mathcal{A} is given a description of f_k as part of its input. The family $(f_k)_{k \in \mathbb{N}}$ is said to be one-way if $\text{Adv}_{\mathcal{A}}^{\text{OW}}(k)$ is a negligible function for all algorithms \mathcal{A} running in time polynomial in k .

3 ID-based Non-interactive Key Distribution

We formally define an ID-based non-interactive key distribution (ID-NIKD) scheme by three distinct algorithms: **Setup**, **Extract** and **SharedKey**. Algorithms **Setup** and **Extract** are executed by the Trusted Authority (TA), while **SharedKey** can be executed by any entity in possession of its private key and the identifier of any other entity with which it wishes to generate a shared key.

- **Setup:** On input 1^k , outputs a master public key (or system parameters) mpk and master secret key msk .
- **Extract:** On input mpk , msk and identifier $\text{ID} \in \{0, 1\}^*$, returns a private key S_{ID} from some space of private keys \mathcal{SK} .
- **SharedKey:** On input mpk , a private key S_{ID_A} and an identifier $\text{ID}_B \in \{0, 1\}^*$, where $\text{ID}_B \neq \text{ID}_A$, this algorithm returns a key $K_{A,B}$ from some space of shared keys \mathcal{SHK} specified in mpk .

We require that, for any pair of identities ID_A , ID_B , and corresponding private keys S_{ID_A} , S_{ID_B} , **SharedKey** satisfies the constraint:

$$\text{SharedKey}(mpk, S_{\text{ID}_A}, \text{ID}_B) = \text{SharedKey}(mpk, S_{\text{ID}_B}, \text{ID}_A).$$

This ensures that entities A and B can indeed generate a shared key without any interaction. We will normally assume that \mathcal{SHK} , the space of shared keys, is $\{0, 1\}^{n(k)}$ for some function $n(k)$. In practice, this can be arranged by hashing a “raw” key.

3.1 Definition of Security for ID-based Non-interactive Key Distribution

Dupont and Enge [13] introduced a basic security notion for ID-NIKD. Their model allows an adversary to extract an arbitrary number of private keys for different entities, thus capturing collusion attacks, but does not give the adversary any direct oracle access to shared keys. This is somewhat analogous to the weak “No Reveals” adversary that is sometimes considered in security models for key exchange. This restriction means that the model of [13] does not capture the very natural requirement that the adversary, even after capturing keys shared between some pairs of entities, should still not be able to compute further shared keys. An adversary can of course compute any given shared key after extracting the private key of one of the relevant entities. But these different key types may be afforded different levels of protection in practice, so differentiating between the different types of compromise that are possible gives us a more refined model that may better represent real applications. The model of [13] also requires the adversary to compute the shared key held between two entities in order to be judged successful, whereas an indistinguishability-based definition is stronger, and more closely aligned with existing models for key exchange. In what follows, we strengthen the ID-NIKD security definition of [13] to include adversarial access to a shared key oracle and to use an indistinguishability-based success criterion.

Our definition is stated in terms of a game between an adversary \mathcal{A} and a challenger \mathcal{C} . \mathcal{C} takes as input the security parameter 1^k , runs algorithm **Setup** of the ID-NIKD scheme and gives \mathcal{A} mpk . It keeps msk to itself. \mathcal{A} then makes queries of the following three types:

- **Extract**(ID): \mathcal{C} responds by running algorithm **Extract** of the ID-based NIKD scheme with input (mpk, msk, ID) to generate a private key S_{ID} . \mathcal{C} gives S_{ID} to \mathcal{A} .
- **Reveal**(ID_A, ID_B): \mathcal{C} responds by running **Extract**(mpk, msk, ID_A) to obtain a private key S_A and then **SharedKey**(mpk, S_A, ID_B) to generate a shared key $K_{A,B}$. \mathcal{C} gives $K_{A,B}$ to \mathcal{A} .
- **Test**(ID_A, ID_B): \mathcal{C} responds by calculating $K_{A,B}$ as above. \mathcal{C} then selects $b \xleftarrow{\$} \{0, 1\}$. If $b = 0$ then \mathcal{C} gives $K_{A,B}$ to \mathcal{A} ; if $b = 1$, then \mathcal{C} gives \mathcal{A} a random element from \mathcal{SHK} .

\mathcal{A} 's queries may be made adaptively and are arbitrary in number, except that \mathcal{A} is allowed to make only one **Test** query (as is usual in security models for key exchange). Moreover, no query to the **Reveal** oracle is allowed on the pair of identities selected for the **Test** query (in either order), and no **Extract** query is allowed on either of the identities involved in the **Test** query.

Finally, \mathcal{A} outputs a bit b' , and wins the game if $b' = b$. \mathcal{A} 's advantage in this IND-SK (indistinguishability of shared key) security game is defined to be

$$\text{Adv}_{\mathcal{A}}^{\text{IND-SK}}(k) = |\Pr[b = b'] - 1/2|.$$

We say that an ID-NIKD scheme is IND-SK secure if for any polynomial time adversary \mathcal{A} , the function $\text{Adv}_{\mathcal{A}}^{\text{IND-SK}}(k)$ is negligible.

Note that a weaker computational version of this model is easily obtained by removing the **Test** oracle and changing the win condition for the game to require that \mathcal{A} output the actual shared key $K_{A,B}$ for two identities ID_A, ID_B neither of which is the subject of an **Extract** query, and that are not together the subject of a **Reveal** query. This model is still stronger than the model of [13], and is useful in the sequel. We refer to COMP-SK security in this case.

3.2 ID-based Non-interactive Key Distribution from Trapdoor Discrete Log Groups

We specify how to obtain an ID-NIKD scheme from any TDL group generator. We need to assume the existence of efficient hash functions $H_1 : \{0, 1\}^* \rightarrow G$ and $H_2 : G \rightarrow \{0, 1\}^n$ in our construction;

these are modelled as random oracles in our security proof. The component algorithms of our TDL-based ID-NIKD scheme are defined as follows:

- **Setup:** On input 1^k , this algorithm runs **TDLGen** of the TDL generator and obtains a tuple (G, r, g, T) . It outputs $mpk = (G, g, H_1, H_2, n)$ where $H_1 : \{0, 1\}^* \rightarrow G$ and $H_2 : G \rightarrow \{0, 1\}^n$ are hash functions, and $\mathcal{SHK} = \{0, 1\}^n$. It also outputs $msk = (G, g, H_1, H_2, n, r, T)$. Here $n = n(k)$ will be the bit-length of shared keys in the ID-NIKD scheme.
- **Extract:** On input mpk, msk and identifier $ID \in \{0, 1\}^*$, this algorithm runs algorithm **SolveDL** on input $H_1(ID)$ to obtain a value $S_{ID} \in \mathbb{Z}_r$ such that $g^{S_{ID}} = H_1(ID)$. The algorithm then outputs S_{ID} .
- **SharedKey:** On input mpk , a private key S_{ID_A} and an identifier $ID_B \in \{0, 1\}^*$, where $ID_B \neq ID_A$, this algorithm outputs the value $H_2(H_1(ID_B)^{S_{ID_A}}) = H_2(g^{S_{ID_A} S_{ID_B}}) \in \{0, 1\}^n$.

It is clear that **SharedKey** defined in this way satisfies the requirement that entities A and B are able to compute a common key. Note that g need not be included in the public parameters. However, including it slightly simplifies our later presentation.

Theorem 1. *The TDL-based ID-NIKD scheme is secure assuming the hardness of the CDH problem in groups G produced by the TDL group generator. In more detail, for any IND-SK adversary \mathcal{A} against the TDL-based ID-NIKD scheme that makes q_i queries to hash function H_i for $i = 1, 2$, there is an algorithm \mathcal{B} that solves the CDH problem in groups G produced by the TDL group generator with*

$$Adv_{\mathcal{B}}^{CDH}(k) \geq Adv_{\mathcal{A}}^{IND-SK}(k)/q_1^2 q_2.$$

Moreover, \mathcal{B} runs in time $O(\text{time}(\mathcal{A}))$. The proof is in the random oracle model, i.e. H_1, H_2 are modelled as random oracles.

The proof of this result is given in the appendix.

A variant scheme whose security is based on the hardness of the DDH problem in the TDL group is obtained simply by omitting the hash function H_2 , so that the shared key is defined to be the “raw” value $g^{S_A S_B} \in G$. We can obtain COMP-SK security based on the hardness of the CDH problem in the TDL group for this variant scheme. The proofs of these results use similar techniques as the proof of Theorem 1.

Our TDL-based ID-NIKDS presented here is a generalisation of the schemes given in [23] from specific TDL groups in the RSA setting to the setting of general TDL groups. As we noted in the introduction, the schemes of [23] and their successors [24, 25] lack formal security analysis and all suffer from attacks of various kinds [22, 19, 27]. Our analysis above shows that what is needed to prevent the various attacks on the Maurer-Yacobi ID-NIKD schemes is the introduction of appropriate hash functions. We return to the issue of how to select parameters and construct appropriate hash functions in Section 5, where we examine RSA-based instantiations of our TDL-based schemes. A scheme of this type in the setting of anomalous elliptic curves over \mathbb{Z}_n was also sketched in [28], but the specific scheme was soon found to be insecure by the authors.

3.3 Security of the ID-NIKD Scheme of Sakai, Ohgishi and Kasahara

In this section, we prove the security of the ID-NIKD scheme of Sakai, Ohgishi and Kasahara [31] in our extended security model. This strengthens the main result of [13].

The Sakai-Ohgishi-Kasahara (SOK) ID-NIKD scheme makes use of a pairing-friendly group generator **PairingGen** and has the following algorithms:

- **Setup**: On input 1^k , this algorithm runs `PairingGen` to obtain a tuple (G, G_T, e, q, P) with the usual properties. It selects $s \xleftarrow{\$} \mathbb{Z}_r$ and outputs $mpk = (G, G_T, e, q, P, P_0 = sP, H_1, H_2, n)$ where $H_1 : \{0, 1\}^* \rightarrow G$ and $H_2 : G_T \rightarrow \{0, 1\}^n$ are hash functions, and $\mathcal{SHK} = \{0, 1\}^n$. It also outputs $msk = s$.
- **Extract**: On input mpk , msk and identifier $ID \in \{0, 1\}^*$, this algorithm outputs $S_{ID} = sH_1(ID)$.
- **SharedKey**: On input mpk , a private key S_{ID_A} and an identifier $ID_B \in \{0, 1\}^*$, where $ID_B \neq ID_A$, this algorithm outputs the value $H_2(e(S_{ID_A}, H_1(ID_B))) \in \{0, 1\}^n$.

It is clear from bilinearity of the map e that `SharedKey` defined in this way satisfies the requirement that entities A and B are able to compute a common key. It also should be pointed out that Dupont and Enge [13] analyse a slight generalisation of the SOK ID-NIKD scheme that operates in the more general setting of asymmetric pairings. Our analysis can also be transferred to this setting. Note too that, strictly speaking, there is no need to include the value $P_0 = sP$ in the public parameters of the scheme. However, including it simplifies our later presentation.

Theorem 2. *The SOK ID-NIKD scheme is secure assuming the hardness of the BDH problem in groups (G, G_T) produced by the pairing-friendly group generator `PairingGen`. In more detail, for any IND-SK adversary \mathcal{A} against the SOK ID-NIKD scheme that makes q_i queries to hash function H_i for $i = 1, 2$, there is an algorithm \mathcal{B} that solves the BDH problem in groups (G, G_T) produced by `PairingGen` with*

$$Adv_{\mathcal{B}}^{BDH}(k) \geq Adv_{\mathcal{A}}^{IND-SK}(k)/q_1^2 q_2.$$

Moreover, \mathcal{B} runs in time $O(\text{time}(\mathcal{A}))$. The proof is in the random oracle model, i.e. H_1, H_2 are modelled as random oracles.

The proof of this result can be found in the appendix.

A variant scheme whose security is based on the hardness of the DBDH problem is obtained simply by omitting the hash function H_2 , so that the shared key is defined to be the “raw” value $e(S_{ID_A}, H_1(ID_B)) \in G_T$. We can also obtain COMP-SK security based on the hardness of the BDH problem for this variant scheme.

4 From ID-based Non-interactive Key Distribution to Identity-based Encryption

We show how to construct an IBE scheme from any ID-NIKD scheme that meets two additional technical conditions. We will then show that if the ID-NIKD scheme is IND-SK secure, then the IBE scheme that results from our conversion is IND-ID-CPA secure in the sense of [6]. We note here that we think that a generic construction in the reverse direction is unlikely to be possible, since it seems difficult to obtain a non-interactive primitive from one that is intrinsically interactive. A KEM formulation of our construction is also possible. In some ways this would be a more natural approach to take, since both ID-NIKD schemes and KEMs are concerned with keys, and one can generically obtain IBE from suitable KEMs using results of [4]. However, we are interested in exploring the relationship between existing ID-NIKD and IBE schemes, and so have focussed here on an IBE formulation instead.

We begin by recalling the formal definitions of IBE and its security, then give our conversion.

An IBE scheme is defined in terms of four algorithms:

- **Setup**: On input 1^k , outputs a master public key (or system parameters) mpk and master secret key msk .
- **Extract**: On input mpk , msk and identifier $ID \in \{0, 1\}^*$, returns a private key S_{ID} .
- **Encrypt**: On input mpk , identifier $ID \in \{0, 1\}^*$ and message M (from some message space), returns a ciphertext C .
- **Decrypt**: On input mpk , a private key S_{ID} and a ciphertext C , returns either a message M or a failure symbol \perp .

We have the obvious correctness requirement that decryption “undoes” encryption.

Security models for IBE were first established in [6]. Chosen-ciphertext security is defined in terms of the following IND-ID-CCA game between an adversary \mathcal{A} and a challenger \mathcal{C} . \mathcal{C} takes as input the security parameter 1^k , runs algorithm **Setup** of the IBE scheme and gives \mathcal{A} mpk . It keeps msk to itself. \mathcal{A} then runs in two phases:

- **Phase 1**: \mathcal{A} issues a series of adaptively selected **Extract** and **Decrypt** queries on identities ID and identifier/ciphertext combinations (ID, C) of its choice. These are replied to by \mathcal{C} by using algorithms **Extract** and **Decrypt** and knowledge of msk .
- **Challenge**: After \mathcal{A} decides to end Phase 1, it outputs two equal length messages M_0, M_1 and a challenge identifier ID^* . \mathcal{C} selects $b \xleftarrow{\$} \{0, 1\}$, sets $C^* = \text{Encrypt}(mpk, ID^*, M_b)$ and gives C^* to \mathcal{A} . We require that ID^* not be the subject of any **Extract** query in Phase 1.
- **Phase 2**: This phase proceeds as in Phase 1, with the constraints that ID^* not be the subject of any **Extract** query and that the pair (ID^*, C^*) not be the subject of any **Decrypt** query.
- **Guess**: Finally \mathcal{A} outputs a bit $b' \in \{0, 1\}$.

The advantage of \mathcal{A} against the IBE scheme in the above IND-ID-CCA security game is defined to be:

$$Adv_{\mathcal{A}}^{\text{IND-ID-CCA}}(k) = \left| \Pr[b = b'] - \frac{1}{2} \right|$$

where the probability is measured over the random choices of coins of \mathcal{A} and \mathcal{C} . An IBE scheme is said to be IND-ID-CCA secure if the function $Adv_{\mathcal{A}}^{\text{IND-ID-CCA}}(k)$ is negligible for all polynomial time adversaries \mathcal{A} .

A weaker notion of IND-ID-CPA security for IBE is obtained by removing the adversary’s access to the **Decrypt** oracle. The advantage of \mathcal{A} against the IBE scheme in the resulting IND-ID-CPA security game is defined to be:

$$Adv_{\mathcal{A}}^{\text{IND-ID-CPA}}(k) = \left| \Pr[b = b'] - \frac{1}{2} \right|$$

and an IBE scheme is said to be IND-ID-CPA secure if the function $Adv_{\mathcal{A}}^{\text{IND-ID-CPA}}(k)$ is negligible for all polynomial time adversaries \mathcal{A} .

4.1 The Generic Conversion

We are now ready to present our generic conversion from ID-NIKD to IBE. We require the ID-NIKD scheme to satisfy the following requirements:

1. The **Extract** and **SharedKey** algorithms of the ID-NIKD scheme should, as a first step, hash the input identifier ID using a member $h = h_k$ of a one-way hash function family $(h_k)_{k \in \mathbb{N}}$ to produce an element U_{ID} in some set \mathcal{PK} , with all further computations in the **Extract** and **SharedKey** algorithms depending only on U_{ID} and not on ID . We assume that h is described in the public parameters of the scheme. We may think of U_{ID} as being the public key corresponding to the string ID . If this condition is satisfied, then, from the algorithm **SharedKey** with inputs mpk, S_{ID_A}, ID_B we can construct a new algorithm **SharedKey'** with inputs $mpk, S_{ID_A}, U_{ID_B} = h(ID_B)$ that has the same output as **SharedKey**.
2. There should exist a randomized algorithm **Sample** that on input mpk , outputs pairs $(S, U) \in \{\mathcal{SK}\} \times \{\mathcal{PK}\}$ with S being a private key corresponding to public key U and $U \xleftarrow{\$} \{\mathcal{PK}\}$. Note that an identifier corresponding to U will not be obtainable from U when it is generated in this way without inverting the one-way hash function H .

These conditions are certainly satisfied for the TDL-based and the SOK ID-NIKD schemes considered in Section 3. In both schemes, a hash function H_1 is first used to convert identifiers into group elements before any further processing. Modelling H_1 as a random oracle in the security analysis is a stronger assumption than H_1 being one-way. For the second condition for the TDL-based scheme, we can define **Sample** as setting $S \xleftarrow{\$} \mathbb{Z}_r$ and $U = g^S$. If r is not part of mpk , we may set $S \xleftarrow{\$} \mathbb{Z}_R$ instead. For the SOK scheme, we can define **Sample** as setting $b \xleftarrow{\$} \mathbb{Z}_q$, $S = bP_0$ and $U = bP$. Here we see the need for including g and P_0 in the public parameters of the ID-NIKD schemes.

Assuming these two conditions are met for an ID-NIKD scheme \mathcal{S} , we construct an IBE scheme $\mathcal{IBE}(\mathcal{S})$ as follows:

- **Setup**: On input 1^k , this algorithm runs the **Setup** of \mathcal{S} (with the same input) to obtain mpk, msk . The master public key of $\mathcal{IBE}(\mathcal{S})$ is set to mpk and the master secret key is set to msk . We assume that mpk contains a description of the shared key space \mathcal{SHK} , and that $\mathcal{SHK} = \{0, 1\}^n$; this will also define the message space of $\mathcal{IBE}(\mathcal{S})$.
- **Extract**: On input (mpk, msk, ID) , this algorithm runs the **Extract** algorithm of the ID-NIKD scheme \mathcal{S} with the same input to obtain a private key S_{ID} . The output is S_{ID} .
- **Encrypt**: Let the input be $mpk, ID, M \in \{0, 1\}^n$. This algorithm runs **Sample** to obtain a pair (S, U) . It then runs **SharedKey** of scheme \mathcal{S} on input (mpk, S, ID) to obtain a key $K \in \{0, 1\}^n$. The output is $C = (U, V)$ where $V = M \oplus K$.
- **Decrypt**: Let the input be (mpk, S_{ID}, C) with $C = (U, V)$. This algorithm runs **SharedKey'** (derived from **SharedKey** of \mathcal{S}) on input (mpk, S_{ID}, U) to obtain a key K . The output is $M = V \oplus K$.

Note that the **Setup** and **Extract** algorithms of the IBE scheme $\mathcal{IBE}(\mathcal{S})$ are almost identical to those of the ID-NIKD scheme \mathcal{S} . The main idea of the construction is that the encrypting party A can generate an “on-the-fly” key-pair (S, U) for each encryption to B ; running **SharedKey** (mpk, S, ID_B) allows A to generate a shared key with B having identifier ID_B , and sending the public key U to B as part of the ciphertext allows B to compute the same shared key by running **SharedKey'** (mpk, S_{ID_B}, U) . This shared key is then used to protect the message.

The scheme can be generalised to handle an ID-NIKD scheme whose shared key space \mathcal{SHK} is any set equipped with a group operation; messages are then constrained to also lie in \mathcal{SHK} . Alternatively, the key K produced during **Encrypt** can be used to derive a key for a symmetric encryption scheme. The following security result is proved in the appendix.

Theorem 3. *Suppose the ID-NIKD scheme \mathcal{S} is IND-SK secure and satisfies the two conditions above. Suppose also that $(h_k)_{k \in \mathbb{N}}$ used in the construction of \mathcal{S} is a one-way function family. Then the IBE scheme $\mathcal{IBE}(\mathcal{S})$ is IND-ID-CPA secure.*

The above construction provides only IND-ID-CPA security for the scheme $\mathcal{IBE}(\mathcal{S})$. In the random oracle model, we may apply the generic conversion of [36] to obtain an IBE scheme with IND-ID-CCA security, provided $\mathcal{IBE}(\mathcal{S})$ satisfies a mild technical condition (γ -uniformity). The resulting scheme (using either conversion) will be only a little less efficient than $\mathcal{IBE}(\mathcal{S})$. In turn, $\mathcal{IBE}(\mathcal{S})$ has roughly the same performance characteristics as the ID-NIKD scheme that it is built from. In fact, the generic conversion of [36] only requires a one-way security notion for the starting IBE scheme in order to obtain IND-ID-CCA security. We can achieve this notion of security for our ID-NIKD-to-IBE conversion under the weaker requirement that the ID-NIKD scheme be COMP-SK secure. This can allow a slightly more efficient overall construction for an IND-ID-CCA secure IBE scheme, since we can typically obtain COMP-SK security using one less hash function than is needed for IND-SK security, and with a slightly tighter reduction. Moreover, it can be seen from the proof of Theorem 3 that no access to the **Reveal** oracle is needed during the simulation. This means that security in the weaker model for ID-NIKD originally proposed in [13] is actually sufficient for our application to IBE.

In the next sections, we will examine some specific IBE schemes that result from our conversion.

4.2 Applying the Conversion

We have already explained how the TDL-based ID-NIKD scheme and the SOK ID-NIKD scheme meet the requirements for our conversion to be applicable.

In the TDL case, we obtain an IBE scheme with the following algorithms:

- **Setup:** On input 1^k , this algorithm runs **TDLGen** of the TDL generator and obtains a tuple (G, r, g, T) . It outputs $mpk = (G, g, H_1, H_2, n)$ where $H_1 : \{0, 1\}^* \rightarrow G$ and $H_2 : G \rightarrow \{0, 1\}^n$ are hash functions. It also outputs $msk = (G, g, H_1, H_2, n, r, T)$. Here n is the size of plaintext messages.
- **Extract:** On input mpk , msk and identifier $ID \in \{0, 1\}^*$, this algorithm runs algorithm **SolveDL** on input $H_1(ID)$ to obtain a value $S_{ID} \in \mathbb{Z}_r$ such that $g^{S_{ID}} = H_1(ID)$. The algorithm then outputs S_{ID} .
- **Encrypt:** On input mpk , identifier $ID \in \{0, 1\}^*$ and message M , this algorithm returns a ciphertext $C = (U, V)$ where $S \xleftarrow{\$} \mathbb{Z}_r$, $U = g^S$ and $V = M \oplus H_2(H_1(ID)^S)$.
- **Decrypt:** On input mpk , a private key S_{ID} and a ciphertext $C = (U, V)$, this algorithm outputs $M = V \oplus H_2(U^{S_{ID}})$.

The IND-ID-CPA security of this IBE scheme is guaranteed by Theorems 1 and 3, assuming the hardness of the CDH problem in groups G produced by **TDLGen**. The scheme is an ID-based analogue of the ElGamal encryption scheme; indeed it can be constructed and its security analysed directly without going via our ID-NIKD-to-IBE conversion. However, we are not aware of this IBE scheme having been explicitly presented in the literature before. A variant of it without explicit hash functions and without any security analysis was sketched in [18]. We will explore specific instantiations of this TDL-based scheme in more detail in the next section.

In the SOK case, it is easy to see that the IND-ID-CPA secure IBE scheme that results from applying our generic conversion is nothing other than the BasicIdent IBE scheme of Boneh and

Franklin [6]: sampling produces a pair $S = bP_0$ and $U = bP$, and the ciphertext has the form $C = (U, V)$ where $V = M \oplus H_2(e(S, H_1(\text{ID})))$. Here,

$$e(S, H_1(\text{ID})) = e(bP_0, H_1(\text{ID})) = e(bP, sH_1(\text{ID})) = e(U, sH_1(\text{ID})) = e(U, S_{\text{ID}}),$$

showing how the recipient may perform decryption.

Thus our approach provides a “new” proof of security for this scheme (and since the IND-ID-CCA secure scheme FullIdent of [6] is effectively obtained via the later generic conversion of [36], a proof for FullIdent too). More importantly, our conversion demonstrates how seemingly quite different IBE schemes can be seen as arising in a uniform way from a common underlying primitive, namely ID-NIKD. Specifically, it explains in a new way the relationship between the SOK ID-NIKD scheme and the IBE schemes of Boneh and Franklin [6].

5 Instantiating the TDL-based Schemes

5.1 The RSA Setting

We present and evaluate an instantiation of a TDL group generator for the RSA setting. Our presentation borrows in part from [27] (but note that [27] contains only a heuristic security analysis). We assume the reader is familiar with basic number theory and algorithms for factorisation and discrete logs. See [26] for further background.

Let $N = pq$ where $p = 3 \pmod 4$, $q = 1 \pmod 4$ and $\gcd(p-1, q-1) = 2$. Then the maximal order of any element in \mathbb{Z}_N^* is $(p-1)(q-1)/2$. For $x \in \mathbb{Z}_N$, let $(\frac{x}{N})$ denote the Jacobi symbol. Let g be an element of \mathbb{Z}_N^* such that $g_p = g \pmod p$ is primitive in \mathbb{Z}_p and $g_q = g \pmod q$ is primitive in \mathbb{Z}_q . Then g has maximal order in \mathbb{Z}_N^* and $(\frac{g}{N}) = 1$. Let G denote the subgroup of \mathbb{Z}_N^* generated by g ; G will be our trapdoor discrete log group. We next explain how to hash onto G , and then discuss how to select p and q in order to build a trapdoor for extracting discrete logs in G .

It is easy to verify that G is exactly the set of all elements of \mathbb{Z}_N with Jacobi symbol 1. Moreover, from the choice of p and q , we have $(\frac{-1}{N}) = -1$. Let H be a hash function from $\{0, 1\}^*$ onto \mathbb{Z}_N , and define

$$H_1(\text{ID}) = \left(\frac{H(\text{ID})}{N} \right) \cdot H(\text{ID}).$$

Since -1 has Jacobi symbol -1 in \mathbb{Z}_N^* , we see that outputs of H_1 have Jacobi symbol 1, and so lie in G (unless a hash output lying in $\mathbb{Z}_N \setminus \mathbb{Z}_N^*$ can be found, in which case N can be factored). If $H : \{0, 1\}^* \rightarrow \mathbb{Z}_N$ is modelled as a random oracle, then H_1 is a random oracle with output in G . In practice, it is easy to instantiate H with output in \mathbb{Z}_N using a hash function such as SHA-256.

We next discuss in more detail how p and q should be chosen in order to obtain a suitable trapdoor for the DLP. With the information p and q , discrete logs can be extracted in G by finding discrete logs in \mathbb{Z}_p and \mathbb{Z}_q and combining the results using Chinese remaindering. The idea, as originally sketched in [23], is to choose p and q so that $p-1$ and $q-1$ are both B -smooth, where B is some bound to be determined. Then the Pohlig-Hellman algorithm can be combined with Pollard’s rho algorithm to find discrete logs in \mathbb{Z}_p and \mathbb{Z}_q in time roughly $O(\ell B^{1/2})$ where ℓ is the number of prime factors in $p-1$ and $q-1$, using essentially no storage. This algorithm can be parallelised. Now consider the case when the trapdoor information is not known. It is shown in [25, Lemma 2] that extracting discrete logs to the base g in this situation is at least as difficult as factoring N . On the other hand when N is large enough, the best known algorithm for factoring

N will either be the Number Field Sieve (with running time $L_N(1/3, c)$ for some constant c) or Pollard’s $p - 1$ factoring algorithm. The latter algorithm requires running time $O(B \log N / \log B)$ if $p - 1$ and $q - 1$ are both B -smooth and if both do have a prime factor of size comparable to B . It does not seem to be capable of parallelisation.

Thus we see that, provided N is chosen so that the Number Field Sieve is not effective, solving the DLP in G with the trapdoor information takes time $O(\ell B^{1/2})$ while, without the trapdoor information, it takes time $O(B \log N / \log B)$. Roughly speaking, then, with the trapdoor, finding discrete logs takes time 2^k , and without it, time 2^{2k} , for some value k that can be selected at will. This gives us the basis of a trapdoor for the DLP in our group G , albeit where knowledge of the trapdoor only gives a square-root speed-up rather than a polynomial-time algorithm for discrete logs. For example, suppose $N = pq$ with p and q as above has 1024 bits and $B = 2^{80}$. Then the NFS has complexity roughly 2^{80} operations, as does Pollard’s $p - 1$ algorithm. So without the trapdoor, the DLP in G takes effort roughly 2^{80} , while with the trapdoor, it takes effort roughly 2^{40} .

The above discussion can be formalised in the form of a trapdoor discrete log generator. Applying our constructions, we obtain concrete ID-NIKD and IBE schemes that are provably secure in the random oracle model, assuming the hardness of the CDH problem in the group G consisting of elements in \mathbb{Z}_N with Jacobi symbol 1. It is known that the CDH problem in such groups is at least as hard as factoring (see [26, Fact 3.80]). This security guarantee is, to the best of our knowledge, in contrast to all previously presented versions of these schemes. The schemes are fairly efficient, with encryption requiring two exponentiations modulo N and decryption requiring one. Public parameters and private keys are compact. However, private key extraction is somewhat inefficient, requiring about 2^{40} effort for each private key at the 80-bit security level. This is because, with the Pohlig-Hellman and Pollard rho algorithms, it does not seem possible to re-use any of the effort expended in extracting one log to find another. On the other hand, this effort can be spread across multiple processors and is storage-free. Re-use of effort would be possible if one instead used an index-calculus style algorithm for discrete logs in \mathbb{Z}_p and \mathbb{Z}_q . For parameters of practical interest, this would effectively trade a large pre-computation of the order of $L_p(1/3, c)$ for rapid extraction of individual discrete logs thereafter. In our example where p has roughly 512 bits, this pre-computation is significant but feasible. The factorisation of RSA-512 [10] took 8000 mips years using the Number Field Sieve and one would expect a roughly similar pre-computation cost for finding discrete logs in \mathbb{Z}_p with p having 512 bits.

Finally, we note that an alternative approach to constructing N using more than 2 prime factors appears to give a greater asymmetry for the DLP with and without the trapdoor, at least for parameters of practical interest. Analysis in [23] and more recently [20] proceeds on this basis. However, we do not know how to securely and efficiently hash onto a large cyclic subgroup of \mathbb{Z}_N^* in that case, and as the attacks on earlier ID-NIKD instantiations [23–25] and our security analysis show, proper hashing is vital for security. Thus this alternative approach does not seem to be compatible with security.

5.2 Trapdoor Discrete Logs from Isogenies and Weil Descent

Galbraith *et al.* [14] proposed a method by which elliptic curves on which the DLP can be solved using Weil descent can be disguised using isogenies to look like random curves. Teske [34] developed these ideas, giving a more detailed complexity and security analysis. In essence, the construction is as follows.

We start with a special elliptic curve E defined over $\mathbb{F}_{2^{161}}$ for which it is possible to construct an explicit group homomorphism Φ from E into $J_C(\mathbb{F}_{2^{23}})$, the Jacobian of a hyperelliptic curve of genus 7 or 8 defined over $\mathbb{F}_{2^{23}}$. Let $\langle P \rangle$ of order r be a large cyclic subgroup of the group of points on E , and assume $P \notin \ker(\Phi)$. Then the DLP in $G = \langle P \rangle$ may be transferred using Φ to $J_C(\mathbb{F}_{2^{23}})$, where sub-exponential algorithms using index-calculus techniques are available. It is estimated in [34, Section 2.2] that the DLP in $J_C(\mathbb{F}_{2^{23}})$ can be solved using around 2^{48} bit operations, involving a factor base of size 2^{22} . Most of this work can be parallelised. In fact, individual discrete log problems in $J_C(\mathbb{F}_{2^{23}})$ can be solved relatively easily once a pre-computation of this order of magnitude has been carried out: each new problem requires the generation of an extra relation for the index-calculus algorithm and a small amount of linear algebra. So, given that finding 2^{22} relations takes about 2^{48} bit operations, finding each additional relation should take about a further 2^{26} bit operations.

Now from E we create a second elliptic curve E' having the same order as E by applying a “random walk of isogenies” to E . Explicit algorithms for calculating the defining equations of E' are given in [34]. The new curve E' is expected to be vulnerable to a Weil descent attack with very low probability, and indeed the best algorithm for solving the DLP in E' seems most likely to be the Pollard rho method, with running time $O(2^{80})$ for this size of curve. It is possible to compute an explicit representation of the chain of the isogenies which maps E' back to E . Let Ψ denote the composition of the isogenies in this chain, and let $P' \in E'$ denote the pre-image of P under this map. Then the DLP in $G' = \langle P' \rangle$ of order r can be efficiently transferred into G using Ψ , and thence to a feasible DLP in $J_C(\mathbb{F}_{2^{23}})$ using Φ .

The construction of a trapdoor discrete log group in this setting is now straightforward: the trapdoor group is G' with generator P' . Without the trapdoor, this appears to be a subgroup of order r on a random curve over $\mathbb{F}_{2^{161}}$ for which the most efficient algorithm for solving the DLP is the Pollard rho method. Detailed analysis supporting this statement can be found in [34, Section 5]. The trapdoor information is the pair of maps (Ψ, Φ) . These can be used to efficiently map the DLP in G' into $J_C(\mathbb{F}_{2^{23}})$, where an index-calculus algorithm can be used to solve the DLP.

Applying our ID-NIKD and IBE constructions in this setting gives us schemes that are extremely efficient in every respect except for private key extraction, and which carry proofs of security based on the hardness of the CDH problem in G' , giving roughly 80 bits of security. To illustrate this efficiency claim further, consider the TDL-based IBE scheme of Section 4.2. Hashing onto elements $(x, y) \in G'$ is straightforward. We may use an ordinary hash function to map bitstrings ID onto elements $x \in \mathbb{F}_{2^{161}}$. We then use an iterated approach to obtain a suitable point on E' : we repeatedly extend ID with prefixes and hash to produce an x -coordinate, until we can solve for $y \in \mathbb{F}_{2^{161}}$ such that $(x, y) \in E'$. The success probability is roughly $1/2$ on each iteration. Then we need to multiply by the co-factor $|E'|/r$, which can be arranged to be small, say 2 or 4 [34], to obtain an element of G' . Each private key is an element of \mathbb{Z}_r and can be represented using just 162 bits; public parameters are compact, containing a description of E' , r and the generator P' . Encryption (resp. decryption) requires just 2 (resp. 1) point multiplications in E' , together with some hashing. Ciphertexts (in the IND-ID-CPA secure scheme) have lengths the same as plaintexts plus one group element, and so have an overhead of only 162 bits.

As we have indicated above, private key extraction, being based on an index-calculus style algorithm, requires a pre-computation step during which many relations are gathered and a large linear algebra computation to obtain a reduced matrix is carried out. After this step, extracting individual private keys is fairly trivial because each key needs only one additional relation to be produced and then a small amount of linear algebra. For the specific parameters in [34], the pre-

computation cost can be estimated at about 2^{48} bit operations, and the cost of each private key extraction at about 2^{26} bit operations. The pre-computation task can be parallelised. Notice that this deployment scenario is rather different from the one envisaged in [34] – there every individual generates his own curves E, E' and gives the trapdoor data (Ψ, Φ) to an escrow agency, with the aim being to force the agency to expend considerable effort to break each key.

6 Conclusion and Open Problems

We have explored the relationships between ID-based non-interactive key distribution and identity-based encryption, showing how the IBE scheme of Boneh and Franklin [6] can be seen as arising from the ID-NIKD scheme of Sakai *et al.* [31]. We have given constructions for secure ID-NIKD and IBE schemes in trapdoor discrete log groups, and studied the instantiations of these schemes in two settings, the RSA-based setting of Maurer and Yacobi, and the disguised elliptic curve setting of Galbraith *et al.* [14] and Teske [34]. We obtained two provably secure IBE schemes with attractive performance in every respect except private key generation. These IBE schemes may be reasonably said to be near-practical.

Several open problems are raised by our work. From a theoretical perspective, it would be interesting to extend our constructions based on TDL groups to the hierarchical ID-based setting and to find constructions offering security in the standard model. It would also be interesting to explore if further known IBE schemes can be related to (presumably as yet unknown) ID-NIKD schemes. From a practical perspective, it would be nice to find an efficient and secure way of hashing onto \mathbb{Z}_N^* when N has more than two prime factors, since this would allow more flexibility in parameter choices in the RSA setting. In particular, this may allow us to make better use of index calculus algorithms in \mathbb{Z}_p for each factor p , shifting most of the cost of private key extraction to a pre-computation. One considerable drawback of the disguised elliptic curve version of our TDL-based ID-NIKD and IBE schemes is that the scheme parameters do not seem to scale very well to higher security levels. A closer examination of the techniques in [34] may reveal variations on this approach with more attractive scaling properties. It may be worthwhile to investigate more closely the performance characteristics of schemes arising from the TDL groups proposed in [18]. However, Dupont and Enge [13] indicate that the gap between the hardness of the DLP with and without the trapdoor is not that great. Of course, the most important question of all is to find examples of TDL groups for which solving the DLP with the trapdoor and hashing onto the group are both easy. These would have many applications in cryptography.

Acknowledgements

This research was supported in part by the European Commission under contract IST-2002- 507932 (ECRYPT).

This research was also sponsored in part by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

The second author is supported by a Dorothy Hodgkin Postgraduate Award, funded by EPSRC and Vodafone and administered by Royal Holloway, University of London.

The authors wish to thank Alex Dent, Andreas Enge, Steven Galbraith, Ueli Maurer, Juanma Gonzalez Nieto and Mike Scott for insightful comments during the development of this paper.

References

1. D. Balfanz, G. Durfee, N. Shankar, D. Smetters, J. Staddon and H.-C. Wong. Secret handshakes from pairing-based key agreements. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pp. 180-196. IEEE Computer Society Press, 2003.
2. M. Bellare, R. Canetti and H. Krawczyk. A modular approach to the design and analysis of authentication and key exchange protocols. In *30th STOC*, pp. 419-428. ACM Press, 1998.
3. M. Bellare and P. Rogaway. Entity authentication and key distribution. In *D.R. Stinson (ed.): CRYPTO'93*, LNCS 773, pp. 232-249. Springer-Verlag, 1994.
4. K. Bentahar, P. Farshim, J. Malone-Lee and N.P. Smart. Generic constructions of identity-based and certificate-less KEMs. Available from <http://eprint.iacr.org/2005/058>, 2005.
5. S. Blake-Wilson, D. Johnson and A. Menezes. Key agreement protocols and their security analysis. In *M. Darnell (ed.): Cryptography and Coding, 6th IMA International Conference*, LNCS 1355, pp. 30-45. Springer-Verlag, 1997.
6. D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In *J. Kilian (ed.): CRYPTO 2001*, LNCS 2139, pp. 213-229. Springer-Verlag, 2001.
7. D. Boneh, C. Gentry and M. Hamburg. Space-efficient identity based encryption without pairings. Cryptology ePrint Archive, Report 2007/177, <http://eprint.iacr.org/>, 2007.
8. C. Boyd, W. Mao, and K.G. Paterson. Key agreement using statically keyed authenticators. In *M. Jakobsson et al. (eds.): ACNS 2004*, LNCS 3089, pp. 248-262, Springer-Verlag, 2004.
9. R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In *B. Pfitzmann (ed.): EUROCRYPT 2001*, LNCS 2045, pp. 453-474. Springer-Verlag, 2001.
10. S. Cavallar, B. Dodson, A.K. Lenstra, W.M. Lioen, P.L. Montgomery, B. Murphy, H. te Riele, K. Aardal, J. Gilchrist, G. Guillerm, P.C. Leyland, J. Marchand, F. Morain, A. Muffett, C. Putnam, C. Putnam and P. Zimmermann. Factorization of a 512-Bit RSA Modulus. In *B. Preneel (ed.): EUROCRYPT 2000*, LNCS 1807, pp. 1-18. Springer-Verlag, 2000.
11. C. Cocks. An identity based encryption scheme based on quadratic residues. In *B. Honary (ed.) Cryptography and Coding, 8th IMA International Conference* LNCS 2260, pp. 360-363. Springer-Verlag, 2001.
12. A.W. Dent and S.D. Galbraith. Hidden pairings and trapdoor DDH groups. In *F. Hess, S. Pauli and M. Pohst (eds.): Algorithmic Number Theory: 7th International Symposium (ANTS VII)*, LNCS 4076, pp. 436-451. Springer-Verlag, 2006.
13. R. Dupont and A. Enge. Provably secure non-interactive key distribution based on pairings. *Discrete Applied Mathematics*, 154(2):270-276, 2006. See also Cryptology ePrint Archive, Report 2002/136, 2002, <http://eprint.iacr.org/>.
14. S. Galbraith, F. Hess, N.P. Smart. Extending the GHS Weil descent attack. In *L. Knudsen (ed.): EUROCRYPT 2002*, LNCS 2332, pp. 29-44, Springer-Verlag 2002.
15. S.D. Galbraith, K.G. Paterson and N.P. Smart. Pairings for cryptographers. Cryptology ePrint Archive: Report 2006/165, 2006. <http://eprint.iacr.org/>.
16. D.M. Gordon. Designing and detecting trapdoors for discrete log cryptosystems. In *E.F. Brickell (ed.): CRYPTO'92*, LNCS 740, pp. 66-75, Springer-Verlag, 1993.
17. S.-H. Heng and K. Kurosawa. k -resilient identity-based encryption in the standard model. In *T. Okamoto (ed.): CT-RSA 2004*, LNCS 2964, pp. 67-80. Springer-Verlag, 2004.
18. D. Hühnlein, M.J. Jacobson, Jr. and D. Weber. Towards practical non-interactive public-key cryptosystems using non-maximal imaginary quadratic orders. *Designs, Codes and Cryptography*, 39(3):281-299, 2003.
19. D. Kügler and M. Maurer. A note on the weakness of the Maurer-Yacobi squaring method. Technical report TI-15/99, Department of Computer Science, Technical University of Darmstadt, Darmstadt, Germany, 1999.
20. N. Kunihiro, W. Abe and K. Ohta. Maurer-Yacobi ID-based key distribution revisited. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 89(5):1421-1424, 2006.
21. W.-B. Lee and K.-C. Liao. Constructing identity-based cryptosystems for discrete logarithm based cryptosystems. *J. Network and Computer Applications*, 27:191-199, 2004.

22. C.H. Lim and P.J. Lee. Modified Maurer-Yacobi's scheme and its applications. In *J. Seberry and Y. Zheng (eds.): ASIACRYPT'92*, LNCS 718, pp. 308-323. Springer-Verlag, 1992.
23. U. Maurer and Y. Yacobi. Non-interactive public-key cryptography. In *D.W. Davies (ed.): EUROCRYPT91*, LNCS 547, pp. 498-507. Springer-Verlag, 1991.
24. U. Maurer and Y. Yacobi. A remark on a non-interactive public-key distribution system. In *R.A. Rueppel (ed.): EUROCRYPT92*, LNCS 658, pp. 458-460. Springer-Verlag, 1993.
25. U.M. Maurer and Y. Yacobi. A Non-interactive public-key distribution System. *Designs, Codes and Cryptography*, 9(3):305-316, 1996.
26. A.J. Menezes, P.C. van Oorschot and S.A. Vanstone. *Handbook of applied cryptography*. CRC Press, 1997.
27. Y. Murakami and M. Kasahara. Murakami-Kasahara ID-based key sharing scheme revisited – in comparison with Maurer-Yacobi schemes. Cryptology ePrint Archive, Report 2005/306, 2005. <http://eprint.iacr.org/>.
28. T. Okamoto and S. Uchiyama. Security of an identity-based cryptosystem and the related reductions. In *K. Nyberg (ed.): EUROCRYPT98*, LNCS 1403, pp. 546-560. Springer-Verlag, 1998.
29. P. Paillier. Public-key cryptosystems based on composite-degree residuosity. In *J. Stern (ed.): EUROCRYPT99*, LNCS 1592, pp. 223-238. Springer-Verlag, 1999.
30. R. Rivest. Controlled algebras and GII's. Talk given at IPAM Workshop on "Securing Cyberspace: Applications and Foundations of Cryptography and Computer Security", October 2006. Available from <https://www.ipam.ucla.edu/>.
31. R. Sakai, K. Ohgishi and M. Kasahara. Cryptosystems based on pairing. *The 2000 Symposium on Cryptography and Information Security, Okinawa, Japan, January*, pp. 26-28, 2000.
32. R. Sakai, K. Ohgishi and M. Kasahara. Cryptosystems based on pairing over elliptic curve. In *2001 Symposium on Cryptography and Information Security (SCIS2000)*, January 2001.
33. A. Shamir. Identity-based cryptosystems and signature schemes. In *G.R. Blakley and D. Chaum (eds.): CRYPTO'84*, LNCS 196, pp. 47-53. Springer-Verlag, 1985.
34. E. Teske. An elliptic curve trapdoor system. *J. Cryptology*, 19(1):115-133, 2006.
35. Y.-M. Tseng and J.-K. Jan. ID-based cryptographic schemes using a non-interactive public-key distribution system. In *ACSAC 1998*, pp. 237-243. IEEE Computer Society, 1998.
36. P. Yang, T. Kitagawa, G. Hanaoka, R. Zhang, K. Matsuura and H. Imai. Applying Fujisaki-Okamoto to identity-based encryption. In *M. Fossorier et al. (eds.): AAEC 2006*, LNCS 3857, pp. 183-192. Springer-Verlag, 2006.

Appendix

Proof of Theorem 1

Suppose there is an adversary \mathcal{A} against the TDL-based ID-NIKD scheme with advantage ϵ and running time t . We show how to construct an algorithm \mathcal{B} that uses \mathcal{A} to solve the CDH problem in groups G produced by the TDL generator.

\mathcal{B} 's input is (G, g, g^a, g^b) where G is a cyclic group produced by TDLGen, g is a generator of G and (g^a, g^b) is an instance of the CDH problem in G . \mathcal{B} 's task is to compute g^{ab} , and it does this by acting as a challenger for \mathcal{A} . We assume here that r , the group order, is part of the description of G . After the proof, we sketch the modifications necessary to handle the case where only a bound R on r is available.

\mathcal{B} gives \mathcal{A} $mpk = (G, g, H_1, H_2, n)$ where H_1 and H_2 are random oracles controlled by \mathcal{B} . Let q_1 be a bound on the number of queries made to H_1 by \mathcal{A} in the course of its attack; similarly let q_2 be a bound on the number of queries made to H_2 . \mathcal{B} chooses two distinct indices I and J uniformly at random from $\{1, 2, \dots, q_1\}$ and a third index L uniformly at random from $\{1, 2, \dots, q_2\}$. \mathcal{A} makes a series of queries which \mathcal{B} answers as follows:

- **H_1 queries:** \mathcal{B} maintains a table to handle \mathcal{A} 's H_1 queries. If ID already appears in an entry of the form (ID, c, h) in the table, then \mathcal{B} returns h in response to \mathcal{A} 's query. Otherwise, if \mathcal{A} 's i -th distinct query to H_1 is on ID_i , then \mathcal{B} proceeds as follows:

1. If $i = I$, then \mathcal{B} adds $(\text{ID}_I, \perp, g^a)$ to the H_1 table and returns g^a .
2. If $i = J$, then \mathcal{B} adds $(\text{ID}_J, \perp, g^b)$ to the H_1 table and returns g^b .
3. Otherwise, \mathcal{B} chooses c_i uniformly at random from \mathbb{Z}_r , adds $(\text{ID}_i, c_i, g^{c_i})$ to the H_1 table, and returns g^{c_i} .

\mathcal{B} 's responses to H_1 queries are uniformly and independently generated. Notice how knowledge of the group order r is used to ensure this.

- **H_2 queries:** \mathcal{B} maintains a table to handle \mathcal{A} 's H_2 queries. If the query s already appears in an entry of the form (s, K) in the table, then \mathcal{B} returns K in response to \mathcal{A} 's query. Otherwise, if \mathcal{A} 's i -th distinct query to H_2 is on s_i , then \mathcal{B} selects $K_i \xleftarrow{\$} \{0, 1\}^n$, adds (s_i, K_i) to the table, and returns K_i to \mathcal{A} . Again, \mathcal{B} 's responses to H_2 queries are uniformly and independently generated.
- **Extract queries:** If \mathcal{A} makes an **Extract** query on an identifier ID , \mathcal{B} first makes an H_1 query on ID if this has not already been done. If $\text{ID} \in \{\text{ID}_I, \text{ID}_J\}$ then \mathcal{B} aborts the simulation. Otherwise, \mathcal{B} finds an entry (ID, c, h) in the H_1 table and outputs c .
- **Reveal queries:** When \mathcal{A} makes a **Reveal** query on a pair of identities $\{\text{ID}_i, \text{ID}_j\}$, \mathcal{B} first makes H_1 queries on ID_i and ID_j if this has not already been done. If $\{\text{ID}_i, \text{ID}_j\} = \{\text{ID}_I, \text{ID}_J\}$ then \mathcal{B} aborts the simulation. Otherwise, suppose $|\{\text{ID}_i, \text{ID}_j\} \cap \{\text{ID}_I, \text{ID}_J\}| \leq 1$. Then \mathcal{B} can obtain from the H_1 table two entries (ID_i, c_i, h_i) and (ID_j, c_j, h_j) , where either $c_i \neq \perp$ or $c_j \neq \perp$. If $c_i \neq \perp$, then \mathcal{B} responds with the value $H_2(h_j^{c_i})$, first making the H_2 query if necessary. If $c_j \neq \perp$, then \mathcal{B} responds with the value $H_2(h_i^{c_j})$.
- **Test query:** At some point during the simulation \mathcal{A} makes a single **Test** query on a pair of identities. If \mathcal{A} does not choose ID_I and ID_J as the identities in this query, then \mathcal{B} aborts its interaction with \mathcal{A} and fails. Otherwise, \mathcal{B} outputs a randomly generated value $K \in \{0, 1\}^n$. Notice that the “correct” key that would be computed by \mathcal{B} in responding to this query is equal to $H_2(g^{ab})$.

This completes our description of \mathcal{B} 's simulation. When \mathcal{A} terminates by outputting a bit b' (or if \mathcal{A} exceeds its normal running time), then \mathcal{B} outputs the value s_L held in the L -th entry of the H_2 list.

We now assess \mathcal{B} 's success probability. Let \mathcal{F} denote the event that \mathcal{B} is not forced to abort during its simulation and let \mathcal{G} denote the event that a query to H_2 on input g^{ab} is made at some point during \mathcal{B} 's simulation. It is easy to see that $\Pr(\mathcal{F}) \geq 1/q_1^2$, and that, conditioned on event \mathcal{F} occurring, then up to the point where \mathcal{G} occurs, \mathcal{B} 's simulation is indistinguishable from that received when \mathcal{A} interacts with a true challenger. Moreover, if \mathcal{F} occurs, but \mathcal{G} does not occur, then it is easy to see that \mathcal{A} 's advantage is zero. For then the shared key that should be held between the two identities involved in the **Test** query (namely ID_I and ID_J) is equal to $H_2(g^{ab})$, and this is independent of \mathcal{A} 's view unless event \mathcal{G} occurs. Hence, assuming event \mathcal{F} occurs, we have

$$\begin{aligned}
\text{Adv}_{\mathcal{A}}^{\text{IND-SK}}(k) &= \left| \Pr[b = b'] - \frac{1}{2} \right| \\
&= \left| \Pr[b = b' | \mathcal{G}] \Pr[\mathcal{G}] + \Pr[b = b' | \neg \mathcal{G}] \Pr[\neg \mathcal{G}] - \frac{1}{2} \right| \\
&\leq |\Pr[\mathcal{G}]| + \left| \Pr[b = b' | \neg \mathcal{G}] - \frac{1}{2} \right| \\
&= \Pr[\mathcal{G}]
\end{aligned}$$

On the other hand, when events \mathcal{F} and \mathcal{G} do occur, \mathcal{B} is successful in outputting g^{ab} with probability at least $1/q_2$, since in this case we know that g^{ab} is on the H_2 list in some position, and \mathcal{B} selects its output from the list in a random position. Combining these facts, we see that $\text{Adv}_{\mathcal{B}}^{\text{CDH}}(k) \geq \epsilon/q_1^2 q_2$. This completes the proof.

The above proof uses \mathcal{B} 's knowledge of r to produce uniformly sampled elements from G with known discrete logs. When only an upper bound R on r is available (because the group order is hidden), \mathcal{B} cannot so easily achieve this. However, if we place \mathbb{Z}_r by \mathbb{Z}_R wherever it occurs in the proof, then \mathcal{B} 's simulation of H_1 deviates from being perfectly uniform by a fraction of at most r/R , and this can be made negligibly small simply by increasing the size of R . For example, in the RSA setting, we can take $R = 2^k N$ where 1^k is the security parameter and N the modulus. Although the private keys c returned by \mathcal{B} to \mathcal{A} as the result of **Extract** queries are now no longer in \mathbb{Z}_r , they are still valid keys (satisfying $g^c = H_1(\text{ID})$ in G). Strictly, the definition of the scheme should then be altered slightly to ensure that the outputs of the **Extract** algorithm are also from \mathbb{Z}_R .

Proof of Theorem 2

Suppose there is an adversary \mathcal{A} against the SOK ID-NIKD scheme with advantage ϵ and running time t . We show how to construct an algorithm \mathcal{B} that uses \mathcal{A} to solve the BDH problem in groups (G, G_T) produced by **PairingGen**.

\mathcal{B} 's input is $(G, G_T, e, q, P, aP, bP, cP)$ where G, G_T are cyclic groups of prime order q , P generates G , $e : G \times G \rightarrow G_T$ is a bilinear map, and (aP, bP, cP) is an instance of the BDH problem in G . \mathcal{B} 's task is to compute $e(P, P)^{abc}$, and it does this by acting as a challenger for \mathcal{A} .

\mathcal{B} gives \mathcal{A} $mpk = (G, G_T, e, q, P, P_0 = cP, H_1, H_2, n)$ where H_1 and H_2 are random oracles controlled by \mathcal{B} . Let q_1 be a bound on the number of queries made to H_1 by \mathcal{A} in the course of its attack; similarly let q_2 be a bound on the number of queries made to H_2 . \mathcal{B} chooses two distinct indices I and J uniformly at random from $\{1, 2, \dots, q_1\}$ and a third index L uniformly at random from $\{1, 2, \dots, q_2\}$. \mathcal{A} makes a series of queries which \mathcal{B} answers as follows:

- **H_1 queries:** \mathcal{B} maintains a table to handle \mathcal{A} 's H_1 queries. If ID already appears in an entry of the form (ID, d, h) in the table, then \mathcal{B} returns h in response to \mathcal{A} 's query. Otherwise, if \mathcal{A} 's i -th distinct query to H_1 is on ID_i , then \mathcal{B} proceeds as follows:
 1. If $i = I$, then \mathcal{B} adds (ID_I, \perp, aP) to the H_1 table and returns aP .
 2. If $i = J$, then \mathcal{B} adds (ID_J, \perp, bP) to the H_1 table and returns bP .
 3. Otherwise, \mathcal{B} chooses d_i uniformly at random from \mathbb{Z}_q , adds $(\text{ID}_i, d_i, d_i P)$ to the H_1 table, and returns $d_i P$.
- \mathcal{B} 's responses to H_1 queries are uniformly and independently generated.
- **H_2 queries:** \mathcal{B} maintains a table to handle \mathcal{A} 's H_2 queries. If the query s already appears in an entry of the form (s, K) in the table, then \mathcal{B} returns K in response to \mathcal{A} 's query. Otherwise, if \mathcal{A} 's i -th distinct query to H_2 is on s_i , then \mathcal{B} selects a random element K_i from G , adds (s_i, K_i) to the table, and returns K_i to \mathcal{A} . Again, \mathcal{B} 's responses to H_2 queries are uniformly and independently generated.
- **Extract queries:** If \mathcal{A} makes an **Extract** query on an identifier ID, \mathcal{B} first makes an H_1 query on ID if this has not already been done. If $\text{ID} \in \{\text{ID}_I, \text{ID}_J\}$ then \mathcal{B} aborts the simulation. Otherwise, \mathcal{B} finds an entry (ID, d, h) in the H_1 table and outputs $d(cP)$.

- **Reveal** queries: When \mathcal{A} makes a **Reveal** query on a pair of identities $\{\text{ID}_i, \text{ID}_j\}$, \mathcal{B} first makes H_1 queries on ID_i and ID_j if this has not already been done. If $\{\text{ID}_i, \text{ID}_j\} = \{\text{ID}_I, \text{ID}_J\}$ then \mathcal{B} aborts the simulation. Otherwise, suppose $|\{\text{ID}_i, \text{ID}_j\} \cap \{\text{ID}_I, \text{ID}_J\}| \leq 1$. Then \mathcal{B} can obtain from the H_1 table two entries (ID_i, d_i, h_i) and (ID_j, d_j, h_j) , where either $d_i \neq \perp$ or $d_j \neq \perp$. If $d_i \neq \perp$, then \mathcal{B} responds with the value $H_2(e(h_j, d_i(cP)))$, first making the H_2 query if necessary. If $d_j \neq \perp$, then \mathcal{B} responds with the value $H_2(e(h_i, d_j(cP)))$.
- **Test** query: At some point during the simulation \mathcal{A} makes a single **Test** query on a pair of identities. If \mathcal{A} does not choose ID_I and ID_J as the identities in this query, then \mathcal{B} aborts its interaction with \mathcal{A} and fails. Otherwise, \mathcal{B} outputs a randomly generated value $K \in \{0, 1\}^n$. Notice that because of the way in which the simulation is set up, the “correct” key that would be computed by \mathcal{B} in responding to this query is equal to $H_2(e(P, P)^{abc})$.

This completes our description of \mathcal{B} 's simulation. When \mathcal{A} terminates by outputting a bit b' (or if \mathcal{A} exceeds its normal running time), then \mathcal{B} outputs the value s_L held in the L -th entry of the H_2 list. Then the assessment of \mathcal{B} 's success probability is almost identical to that in the proof of Theorem 1, and we omit the details here.

Proof of Theorem 3

Let \mathcal{A} be an adversary against $\mathcal{IBE}(\mathcal{S})$ with advantage $\text{Adv}_{\mathcal{A}}^{\text{IND-ID-CPA}}(k)$. We construct from \mathcal{A} two distinct algorithms $\mathcal{B}_1, \mathcal{B}_2$. Algorithm \mathcal{B}_1 attempts to break the one-wayness of hash function h , while algorithm \mathcal{B}_2 attempts to break the ID-NIKD scheme. The subsequent joint analysis of these two algorithms will then give us our result.

\mathcal{B}_1 receives from its challenger \mathcal{C}_1 a value $U \in \mathcal{PK}$ and uses \mathcal{A} in an attempt to find a string $\text{ID}' \in \{0, 1\}^*$ such that $h(\text{ID}') = U$. Clearly, if \mathcal{B}_1 is successful, then it breaks the one-wayness of h . \mathcal{B}_1 runs **Setup** of the scheme $\mathcal{IBE}(\mathcal{S})$, obtaining msk, mpk . Note that h is assumed to be described in mpk . Now \mathcal{B}_1 gives mpk to \mathcal{A} . \mathcal{B}_1 handles \mathcal{A} 's **Extract** queries using its knowledge of msk to run algorithm **Extract** of $\mathcal{IBE}(\mathcal{S})$. When \mathcal{A} submits challenge messages M_0, M_1 and a challenge identifier ID^* , \mathcal{B}_1 sets $C^* = (U, V)$ where $V = M_b \oplus K$ and $K = \text{SharedKey}'(mpk, S_{\text{ID}^*}, U)$. Here, S_{ID^*} is obtained by running **Extract** and $b \stackrel{\$}{\leftarrow} \{0, 1\}$. Eventually \mathcal{A} outputs its bit b' . If at any point in the game, \mathcal{A} made a query to its **Extract** oracle on an identifier ID satisfying $h(\text{ID}) = U$, then \mathcal{B}_1 now outputs ID . If ID^* satisfies $h(\text{ID}^*) = U$, then \mathcal{B}_1 outputs ID^* . Otherwise, \mathcal{B}_1 fails when \mathcal{A} outputs its bit. It is clear that the attack environment provided by \mathcal{B}_1 to \mathcal{A} is indistinguishable from that provided by a real challenger. Moreover, if \mathcal{A} does at any point make a query (either an **Extract** query or during the challenge phase) involving an identifier $\text{ID}' \in \{0, 1\}^*$ such that $h(\text{ID}') = U$, then \mathcal{B}_1 breaks the one-wayness of h .

\mathcal{B}_2 receives from its challenger \mathcal{C}_2 the public parameters mpk of the ID-NIKD scheme \mathcal{S} and uses \mathcal{A} in an attempt to win against \mathcal{C}_2 in the IND-SK security game for \mathcal{S} . Let d denote the hidden bit used by \mathcal{C}_2 in responding to \mathcal{B}_2 's **Test** query. \mathcal{B}_2 begins by selecting $\text{ID}' \stackrel{\$}{\leftarrow} \{0, 1\}^*$ and computing $U = h(\text{ID}')$. \mathcal{B}_2 then passes mpk to \mathcal{A} . \mathcal{A} 's **Extract** queries on identifiers ID are handled by \mathcal{B}_2 by passing them to \mathcal{C}_2 as **Extract** queries in the IND-SK game. However, if $h(\text{ID}) = U$ for any of these queries, then \mathcal{B}_2 aborts. When \mathcal{A} submits challenge messages M_0, M_1 and a challenge identifier ID^* , \mathcal{B}_2 makes its **Test** query to \mathcal{C}_2 on input $(\text{ID}', \text{ID}^*)$, receiving a value K in return. \mathcal{B}_2 then selects $b \stackrel{\$}{\leftarrow} \{0, 1\}$ and sets $C^* = (U, V)$ where $V = M_b \oplus K$. However, if $h(\text{ID}^*) = U$, then \mathcal{B}_2 aborts. Eventually \mathcal{A} outputs its bit b' . If $b' = b$ then \mathcal{B}_2 outputs bit $d' = 0$; otherwise \mathcal{B}_2 outputs

$d' = 1$. The attack environment provided by \mathcal{B}_2 to \mathcal{A} is indistinguishable from that provided by a real challenger, provided that \mathcal{B}_2 does not abort. Moreover, when \mathcal{B}_2 does not abort, \mathcal{B}_2 makes only legal queries to its challenger \mathcal{C}_2 – we can be sure that ID' is distinct from ID^* and that ID' is not involved in any **Extract** query. Then \mathcal{B}_2 's advantage can, via a standard argument, be expressed as:

$$\text{Adv}_{\mathcal{B}_2}^{\text{IND-SK}}(k) = \frac{1}{2} |\Pr[d' = 0|d = 0] - \Pr[d' = 0|d = 1]|.$$

When $d = 1$, the key K returned by \mathcal{C}_2 to \mathcal{B}_2 as a result of the **Test** query is random in $\mathcal{SK} = \{0, 1\}^n$, and then the ciphertext C^* received by \mathcal{A} is independent of the bit b . Hence in this case, \mathcal{A} has zero advantage and $\Pr[b = b'] = 1/2$. So $\Pr[d' = 0|d = 1] = 1/2$. On the other hand, when $d = 0$, the key K returned by \mathcal{C}_2 to \mathcal{B}_2 is equal to the shared key for identifiers ID' , ID^* , and hence C^* is a proper encryption of M_b . Then $|\Pr[d' = 0|d = 0] - 1/2| = |\Pr[b = b'] - 1/2| = \text{Adv}_{\mathcal{A}}^{\text{IND-ID-CPA}}(k)$. Combining this information, we obtain

$$\text{Adv}_{\mathcal{B}_2}^{\text{IND-SK}}(k) = \frac{1}{2} \text{Adv}_{\mathcal{A}}^{\text{IND-ID-CPA}}(k),$$

assuming \mathcal{B}_2 does not abort.

Notice that \mathcal{A} 's view is identical when playing against either \mathcal{B}_1 or \mathcal{B}_2 , unless \mathcal{A} makes a query involving an identifier ID' satisfying $h(\text{ID}') = U$, where U is the first component of the challenge ciphertext C^* (in this situation, \mathcal{B}_2 aborts, but \mathcal{B}_1 does not). Let \mathcal{F} denote this event. So we see that the occurrence of \mathcal{F} is independent of which algorithm \mathcal{A} plays against. But if \mathcal{F} has non-negligible probability, then \mathcal{B}_1 is able to break the one-wayness of $(h_k)_{k \in \mathbb{N}}$. This contradicts our assumption about this hash family. Therefore we may assume that \mathcal{F} has negligible probability. Then, since \mathcal{B}_2 only aborts when \mathcal{F} occurs, we see that \mathcal{B}_2 's advantage against the IND-SK security of the scheme \mathcal{S} is negligibly different from $\frac{1}{2} \text{Adv}_{\mathcal{A}}^{\text{IND-ID-CPA}}(k)$. If \mathcal{A} 's advantage $\text{Adv}_{\mathcal{A}}^{\text{IND-ID-CPA}}(k)$ against $\mathcal{IBE}(\mathcal{S})$ is non-negligible, we have a contradiction to the IND-SK security of \mathcal{S} . This completes the proof.

A slightly more refined analysis using game hopping can be used to show that the advantage of any attacker against the IBE scheme is tightly related to the maximum advantage of any attacker against the ID-NIKD scheme and to the maximum advantage of any algorithm to break the one-wayness of the hash function family. We omit the details.