Practical Identity-Based Encryption Without Random Oracles

Craig Gentry*

Stanford University
cgentry@cs.stanford.edu

Abstract. We present an Identity Based Encryption (IBE) system that is fully secure in the standard model and has several advantages over previous such systems – namely, computational efficiency, shorter public parameters, and a "tight" security reduction, albeit to a stronger assumption that depends on the number of private key generation queries made by the adversary. Our assumption is a variant of Boneh et al.'s decisional Bilinear Diffie-Hellman Exponent assumption, which has been used to construct efficient hierarchical IBE and broadcast encryption systems. The construction is remarkably simple. It also provides recipient anonymity automatically, providing a second (and more efficient) solution to the problem of achieving anonymous IBE without random oracles. Finally, our proof of CCA2 security, which has more in common with the security proof for the Cramer-Shoup encryption scheme than with security proofs for other IBE systems, may be of independent interest.

Keywords: Identity Based Encryption.

1 Introduction

An Identity Based Encryption (IBE) system [25, 8] is a public key encryption system in which a user's public key may be an arbitrary string, such as an email address or other identifier. The user's private key is generated by a trusted authority, called a Private Key Generator (PKG), which applies its master key to the user's identity after the user authenticates itself. Shamir [25] proposed the notion of IBE in 1984 as a way to simplify public key and certificate management. Rather than obtaining the disparate public keys of its intended recipients separately, a message sender who knows the identities of its recipients needs only to obtain the public parameters of the PKG; public key certificates are eliminated altogether.

Boneh and Franklin [8, 9] described the first secure and truly practical IBE system. Their system uses bilinear maps (or "pairings"), and they proved its

* Supported by the Herbert Kunzel Stanford Graduate Fellowship

security in the random oracle model. Canetti et al. [15] presented an IBE system whose security could proven without random oracles, but in a weaker "selective-ID" model, in which the adversary must declare at the beginning of its attack which identity it will target. Boneh and Boyen [4] provided more practical IBE systems in the selective-ID model. Shortly thereafter, Boneh and Boyen [5] presented a fully secure scheme – i.e., one in which the adversary may choose the target identity adaptively – without random oracles. Waters [27] simplified the scheme described in [5], substantially improving its efficiency.

Previous IBE Systems. Moni Naor observed that every IBE system secure against an adaptive-ID attack (as defined by Boneh and Franklin in [8]) implies a signature scheme secure against existential forgery under a chosen-message attack. The generic transformation is as follows: the PKG's parameters correspond to the public key of the signature scheme; private key generation queries to the PKG correspond to signature queries. If an adversary of the signature scheme can forge a signature on an unqueried message, it can generate a private key for an unqueried identity, thus breaking the IBE system. So, to design a secure IBE system, one begins (in some sense) by designing a secure signature scheme.

A common strategy for proving the security of a signature scheme in the random oracle model – e.g., for RSA with full-domain-hash – is as follows. The simulator responds to hash queries in such a way that it can generate a signature on most messages, but not all. The simulator aborts if the adversary requests a signature on a message that it cannot sign, or if the adversary's forgery is on a message that the simulator knows how to sign already. One can also use this strategy to design a secure private key generation procedure for an IBE system. Boneh and Franklin [8] did precisely that; the private key generation procedure in their system is essentially equivalent to the BLS signature scheme [12], which uses the proof strategy just described. (Though, inconveniently for our narrative, Boneh and Franklin's IBE system slightly pre-dates its associated signature scheme.)

When Boneh and Boyen [5] and later Waters [27] devised IBE systems fully secure without random oracles, their main innovation was in the private key generation procedures. Each of these procedures corresponds to a signature scheme that is fully secure (i.e., against a chosen-message attack) without random oracles. Interestingly, though, the (implicit) proof strategy for these standard-model signature schemes is still basically the same as above – i.e., the simulator constructs its public key in such a way that it can generate a signature on most messages, but not all. Since, intuitively speaking, the simulator follows the same strategy except for using its control of the public key (or public parameters, for an IBE system) to compensate for not controlling a random oracle, it should not be surprising that the public parameters for these IBE systems are quite large.

Another side effect of the above proof strategy is that the reduction is loose. If δ is the probability that the simulator can generate a private key for a random identity, then the probability that the simulator does not abort is at most $\delta^q(1-\delta)$, where q is the number of private key generation queries made by the adversary. Setting $\delta \approx 1 - 1/q$ maximizes this probability at $\mathcal{O}(1/q)$. Thus, the reduction loses a multiplicative factor of q. A lossy reduction is not merely a theoretical problem; if we take the lossiness seriously, we should augment the security parameter to compensate, making the system less efficient.

Almost all of the IBE systems since Boneh-Franklin follow the "common strategy" for proving security; consequently, they suffer from long parameters (when security is proven in the standard model) and lossy reductions (in the standard model or the random oracle model). However, we note a couple of exceptions. The IBE systems described in [4] have short parameters and achieve a tight reduction, but this is because they are proven secure only against selective-ID attacks. As noted in [4], one can generally transform a selective-ID scheme into a fully secure scheme by having the simulator guess which identity the adversary will ultimately select, but this transformation loosens the reduction by huge multiplicative factor – namely, by the total number of identities – that is super-polynomial and (much) larger than q. This transformation is also a very unsatisfying approach from a theoretical point of view. A second exception is the IBE system by Katz and Wang [23], which achieves a tight reduction in the random oracle model. In their system, the encryption of M under identity ID effectively consists of two ciphertexts under each of the derived identities $H(\mathsf{ID}, 0)$ and $H(\mathsf{ID}, 1)$ (for hash function H modeled as random oracle). Through its control of the random oracle, the simulator ensures that, for each ID, it knows the private key for exactly one of $H(\mathsf{ID}, 0)$ and $H(\mathsf{ID}, 1)$. It can thus answer any key generation query. The successful adversary partially decrypts the challenge ciphertext with the "wrong" private key with probability $1/2$, giving the simulator useful information. Though this system relies heavily on the random oracle model, it illustrates how a tight reduction for an IBE system can be achieved when the simulator can generate a private key for every identity. A recent paper [2] discusses the Katz-Wang system in detail.

Currently, there is no IBE system that is fully secure without random oracles, yet has short public parameters, or has a tight security reduction. Given this state of affairs, several papers [4, 5, 27] have encouraged work on the open problem of tight security; Waters posed [27] the open problem regarding compact public parameters.

Our Contributions. We present an IBE system that is fully secure without random oracles and has several advantages over previous such systems, including:

– Short public parameters (5 group elements for CCA2 security)
– A tight reduction, albeit based on a stronger assumption (see below)
– Recipient-anonymity

Our constructions are simple and efficient. For example, in the construction described in Section 4.1, which we prove secure against adaptive-ID and adaptive chosen-ciphertext attacks, a ciphertext consists of four group elements. Encryption and decryption require only a small constant number of group operations, while user private keys and the PKG's public parameters are compact. Compare, for example, the public parameters in our IBE system (five group elements and a hash function) to those in [27] ($n + 4$ group elements, where an identity is a bitstring of length n).

An IBE system is recipient-anonymous, roughly speaking, if it hard for an eavesdropper to distinguish which identity was used to generate a given ciphertext. Boneh et al. [7] discuss how anonymous IBE is useful in the context of searchable public key encryption; Abdalla et al. [1] propose the open problem of finding an anonymous IBE system secure without random oracles. Boyen and Waters recently presented the first such anonymous IBE system at the rump session of Crypto 2005 (see [14]). Our IBE system represents a second, but more efficient, solution to this problem; it gives recipient-anonymity basically "for free." The security proof for our scheme is also much simpler. However, we note that the Boyen-Waters approach offers *hierarchical* anonymous IBE.

Regarding the open problem of constructing an IBE system with a tight security reduction, our contribution is less clear. Our decision q-ABDHE assumption, discussed in Section 2.3, is related to the q-BDHE assumption, which has been used to construct efficient hierarchical IBE and broadcast encryption systems [6, 10], but it is stronger than the decision BDH assumption used in [5, 27]. We obtain a tight reduction based on q-ABDHE in the sense that the simulator's time complexity and success probability are identical to that of the adversary in breaking the system, except for *additive* factors depending on q. However, since our assumption is stronger, we cannot claim that a tighter reduction is necessarily an improvement. Moreover, it is not obvious what it means to have an asymptotically tight reduction based on the q-ABDHE assumption, since this assumption varies as q varies. However, we can analyze the concrete security of our system for specific values of q, as we do in Section 3.3. One conclusion of this analysis is that if we assume decision q-ABDHE is no easier than decision BDH (which may or may not be true), then our tighter reduction (for specific reasonable values of q) allows us to choose a smaller security parameter, adding to the efficiency advantages of our scheme. But perhaps this is not a very satisfying "solution" to the open problem; certainly, it would be preferable to obtain a tight reduction under a more natural assumption, such as decision BDH.

A final contribution of this paper is our proof technique, which differs substantially from the "common strategy" described above. Interestingly, our proof strategy draws inspiration from the Cramer-Shoup signature scheme [18] (and strong-RSA based signature schemes, generally) for our private key generation procedure, as well as from the Cramer-Shoup encryption scheme [17] for our approach to proving security against chosen-ciphertext attacks.

Strong-RSA based signatures typically achieve a tight reduction and have short public keys. Intuitively, this is related to the fact that, in the reduction, the simulator can produce a signature for any message. Similarly, unlike in previous IBE systems fully secure in the standard model, the simulator in our reduction can generate a private key for any identity. One can view our private key generation procedure as a strongly existentially unforgeable signature scheme that is "tightly" secure in the standard model under the q-strong DH assumption: that it is hard to compute a pair $(c, g^{1/(\alpha-c)})$ given $\{g^{\alpha^i} : i \in [0, q]\}$, where q corresponds to the anticipated number of queries. The savvy reader may notice that this signature scheme has direct analogue based on strong RSA. In the procedure, the

PKG (signer) publishes groups $\mathbb{G}$ and $\mathbb{G}_T$, and bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$, along with generators $g, g_1, h \in \mathbb{G}$, where $g_1 = g^\alpha$. A private key for identity $\mathsf{ID} \in \mathbb{Z}_p$ is a pair $(r_{\mathsf{ID}}, h_{\mathsf{ID}})$, where $r_{\mathsf{ID}} \in \mathbb{Z}_p$ and $h_{\mathsf{ID}} = (hg^{-r_{\mathsf{ID}}})^{1/(\alpha-\mathsf{ID})}$; if the private key for ID is requested more than once, the PKG uses the same value of r_{ID} each time. In the reduction, the simulator is given $g_i = g^{(\alpha^i)}$ for all $i \in [0, q]$, where q is (roughly) the anticipated number private key generation queries. Given $\{g_i\}$, the simulator computes h by generating a random q-degree polynomial $f(x) \in \mathbb{Z}_p[x]$, and setting $h = g^{f(\alpha)}$. To generate a private key for ID, it sets $r_{\mathsf{ID}} = f(\mathsf{ID})$ and $h_{\mathsf{ID}} = (hg^{-r_{\mathsf{ID}}})^{1/(\alpha-\mathsf{ID})} = g^{(f(\alpha)-f(\mathsf{ID}))/(\alpha-\mathsf{ID})}$; the simulator can compute the latter value from $\{g_i\}$, since $(f(x) - f(\mathsf{ID}))/(x - \mathsf{ID})$ is a $(q-1)$-degree polynomial in x. The values of r_{ID_i} in the simulation for $i \in [1, q]$ appear uniformly random, since $f(x)$ is a random polynomial of degree q. If the adversary can generate a private key $(r'_{\mathsf{ID}}, h'_{\mathsf{ID}})$ for ID for which $r'_{\mathsf{ID}} \neq r_{\mathsf{ID}}$, the simulator can efficiently compute $g^{1/(\alpha-\mathsf{ID})}$.

The fact that the simulator in our system can generate exactly one private key for any identity dovetails nicely with the proof strategy used in the Cramer-Shoup encryption scheme, where the simulator actually knows exactly one valid decryption key: its scalars (x_1, y_1, z_1), along with the dependent values (x_2, y_2, z_2). Roughly speaking, in their proof, Cramer and Shoup show that these scalars remain unconditionally hidden from the adversary (with overwhelming probability), and thus the adversary cannot (except with negligible probability) construct an invalid ciphertext that passes the simulator's validity test, or guess with advantage how the simulator would decrypt its own challenge ciphertext when that challenge ciphertext is incorrectly distributed. How do we adapt their technique to our (multi-user) IBE system? We augment the public parameters to include group elements $h_1, h_2, h_3 \in \mathbb{G}$ (rather than just h), where $h_i = g^{f_i(\alpha)}$ and $f_i(x) \in \mathbb{Z}_p[x]$ is a random and independent q-degree polynomial. The three scalars $r_{\mathsf{ID},i} = f_i(\mathsf{ID})$, which a user receives as part of its private key, play a role analogous to the scalars z_1, x_1, and y_1, respectively, in Cramer-Shoup; the values $r_{\mathsf{ID},2}$ and $r_{\mathsf{ID},3}$ are used in a projective-hash ciphertext validity test. The three scalars remain hidden from the adversary with overwhelming probability, even if the adversary obtains the scalars $r_{\mathsf{ID}',i} = f_i(\mathsf{ID}')$ for less than $q - 1$ identities $\mathsf{ID}' \neq \mathsf{ID}$, since $f_i(x)$ is random and has degree q. Interestingly, previous IBE systems fully secure without random oracles use an entirely different approach to proving chosen-ciphertext security. They employ results by Canetti et al. [16] (later improved by Boneh and Katz [11] and further by Boyen, Mie and Waters [13]) that a chosen-ciphertext-secure IBE system follows from a chosen-plaintext-secure 2-level hierarchical IBE system.

2 Preliminaries

Below, we review the definition of security for an IBE system. We also review the definition of a bilinear map and discuss the complexity assumption on which the security of our system is based.

2.1 Security Model for Identity-Based Encryption

An IBE system consists of four algorithms [25, 8]: *Setup*, *KeyGen*, *Encrypt*, and *Decrypt*. *Setup* establishes the PKG's parameters *params* and a master key *master-key*. *KeyGen* applies the *master-key* to an identity to generate the private key for that identity. *Encrypt* takes a message, an identity and *params* as input, and outputs a ciphertext. *Decrypt* decrypts a ciphertext for an identity using a private key for that identity.

Boneh and Franklin [8, 9] define chosen ciphertext security for IBE systems under a chosen identity attack via the following game.

Setup: The challenger runs *Setup*, and forwards *params* to the adversary.
Phase 1: Proceeding adaptively, the adversary issues queries $q_1, \ldots, q_m$ where q_i is one of the following:
– Key generation query $\langle \mathsf{ID}_i \rangle$: the challenger runs *KeyGen* on ID_i and forwards the resulting private key to the adversary.
– Decryption query $\langle \mathsf{ID}_i, C_i \rangle$. The challenger runs *KeyGen* on ID_i, decrypts C_i with the resulting private key, and sends the result to the adversary.
Challenge: The adversary submits two plaintexts $M_0, M_1 \in \mathcal{M}$ and an identity ID. ID must not have appeared in any key generation query in Phase 1. The challenger selects a random bit $b \in \{0, 1\}$, sets $C = \mathit{Encrypt}(\mathit{params}, \mathsf{ID}, M_b)$, and sends C to the adversary as its challenge ciphertext.
Phase 2: This is identical to Phase 1, except that the adversary may not request a private key for ID or the decryption of (ID, C).
Guess: The adversary submits a guess $b' \in \{0, 1\}$. The adversary wins if $b = b'$.

We call an adversary $\mathcal{A}$ in the above game a IND-ID-CCA adversary.

Definition 1. *An IBE system is* $(t, q_{\mathsf{ID}}, q_C, \epsilon)$ IND-ID-CCA *secure if all* t*-time* IND-ID-CCA *adversaries making at most* q_{ID} *private key queries and at most* q_C *chosen ciphertext queries have advantage at most* ϵ *in winning the above game.*

IND-ID-CPA security is defined similarly, but with the restriction that the adversary cannot make decryption queries.

Definition 2. *An IBE system is* $(t, q_{\mathsf{ID}}, \epsilon)$ IND-ID-CPA *secure if it is* $(t, q_{\mathsf{ID}}, 0, \epsilon)$ IND-ID-CCA *secure.*

Recipient-Anonymity. Informally, we say that an IBE system is anonymous if an adversary cannot distinguish the public key ID under which a ciphertext was generated. More formally, we can incorporate anonymity into our game above through the following simple modification. In the Challenge phase, the adversary outputs two identities ID_0 and ID_1 not queried in Phase 1 and two messages M_0 and M_1. The challenger picks two random bits $b, c \in \{0, 1\}$, uses ID_b to encrypt M_c, and sends the resulting ciphertext C to the adversary. Phase 2 is like Phase 1, except that the adversary cannot request a private key for ID_0 or ID_1, or the decryption of C under either identity. Finally, in the Guess phase, the adversary guesses two bits b', c' and wins if $b = b'$ and $c = c'$; we define the adversary's advantage in this game to be $|\Pr[b = b' \wedge c = c'] - \frac{1}{4}|$.

Definition 3. *We say that an IBE system* $\mathcal{E}$ *is* $(t, q_{\mathsf{ID}}, q_C, \epsilon)$ ANON-IND-ID-CCA *secure if all* t*-time* ANON-IND-ID-CCA *adversaries making at most* q_{ID} *private key queries and at most* q_C *chosen ciphertext queries have advantage at most* ϵ *in the modified game. We define* ANON-IND-ID-CPA *security similarly.*

2.2 Bilinear Maps

We review bilinear maps, using the following standard notation [8, 4, 27]:

1. $\mathbb{G}$ and $\mathbb{G}_T$ are two (multiplicative) cyclic groups of prime order p;
2. g is a generator of $\mathbb{G}$.
3. $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ is a bilinear map.

Let $\mathbb{G}$ and $\mathbb{G}_T$ be two groups as above. A bilinear map is a map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ with the following properties:

1. Bilinear: for all $u, v \in \mathbb{G}$ and $a, b \in \mathbb{Z}$, we have $e(u^a, v^b) = e(u, v)^{ab}$.
2. Non-degenerate: $e(g, g) \neq 1$.

We say that $\mathbb{G}$ is a bilinear group if the group action in $\mathbb{G}$ can be computed efficiently and there exists a group $\mathbb{G}_T$ and an efficiently computable bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ as above. Note that $e(,)$ is symmetric since $e(g^a, g^b) = e(g, g)^{ab} = e(g^b, g^a)$.

2.3 Complexity Assumptions

The security of our system is based on a complexity assumption that we call the decisional augmented bilinear Diffie-Hellman exponent assumption (decisional ABDHE). First, we recall the q-BDHE problem [6, 10], which is as follows: Given a vector of $2q + 1$ elements

$$\left(g', g, g^{\alpha}, g^{(\alpha^2)}, \ldots, g^{(\alpha^q)}, g^{(\alpha^{q+2})}, \ldots, g^{(\alpha^{2q})}\right) \in \mathbb{G}^{2q+1}$$

as input, output $e(g, g')^{(\alpha^{q+1})} \in \mathbb{G}_T$. Since the input vector is missing the term $g^{(\alpha^{q+1})}$, the bilinear map does not seem to help compute $e(g, g')^{(\alpha^{q+1})}$.

We define the q-ABDHE problem almost identically: Given a vector of $2q + 2$ elements

$$\left(g', g'^{(\alpha^{q+2})}, g, g^{\alpha}, g^{(\alpha^2)}, \ldots, g^{(\alpha^q)}, g^{(\alpha^{q+2})}, \ldots, g^{(\alpha^{2q})}\right) \in \mathbb{G}^{2q+2}$$

as input, output $e(g, g')^{(\alpha^{q+1})} \in \mathbb{G}_T$. Introducing the additional term $g'^{(\alpha^{q+2})}$ still does not appear to ease the computation of $e(g, g')^{(\alpha^{q+1})}$, since the input vector is missing the term $g^{(\alpha^{-1})}$.

The q-ABDHE problem is actually more than we need for our IBE system. Instead, we can use a truncated version of the q-ABDHE problem, in which the terms $(g^{(\alpha^{q+2})}, \ldots, g^{(\alpha^{2q})})$ are omitted from the input vector. Clearly, the truncated q-ABDHE problem is hard if the q-ABDHE problem is hard. An algorithm $\mathcal{A}$ has advantage ϵ in solving truncated q-ABDHE if

$$\Pr\left[\mathcal{A}\left(g', g'_{q+2}, g, g_1, \ldots, g_q\right) = e(g_{q+1}, g')\right] \geq \epsilon$$

where we use g_i and g'_i to denote $g^{(\alpha^i)}$ and $g'^{(\alpha^i)}$, and where the probability is over the random choice of generators g, g' in $\mathbb{G}$, the random choice of α in $\mathbb{Z}_p$, and the random bits used by $\mathcal{A}$.

The decisional version of truncated q-ABDHE is defined as one would expect. An algorithm $\mathcal{B}$ that outputs $b \in \{0, 1\}$ has advantage ϵ in solving truncated decision q-ABDHE if

$$\left| \Pr\left[\mathcal{B}\big(g', g'_{q+2}, g, g_1, \ldots, g_q, e(g_{q+1}, g')\big) = 0\right] - \Pr\left[\mathcal{B}\big(g', g'_{q+2}, g, g_1, \ldots, g_q, Z\big) = 0\right] \right| \geq \epsilon$$

where the probability is over the random choice of generators g, g' in $\mathbb{G}$, the random choice of α in $\mathbb{Z}_p$, the random choice of $Z \in \mathbb{G}_T$, and the random bits consumed by $\mathcal{B}$. We refer to the distribution on the left as $\mathcal{P}_{ABDHE}$ and the distribution on the right as $\mathcal{R}_{ABDHE}$.

Definition 4. *We say that the truncated (decision) (t, ϵ, q)-ABDHE assumption holds in $\mathbb{G}$ if no t-time algorithm has advantage at least ϵ in solving the truncated (decision) q-ABDHE problem in $\mathbb{G}$.*

As an aside, we note that the truncated q-ABDHE problem is also closely related to the q-bilinear Diffie-Hellman inversion (q-BDHI) problem, which has been used to construct an IBE system secure without random oracles under a selective-ID attack [4] and a verifiable random function [20]. Specifically, let us define the q-augmented BDHI (q-ABDHI) problem as follows: given a vector of $q + 2$ elements

$$\left(g^{(\alpha^{-q-2})}, g, g^{\alpha}, g^{(\alpha^2)}, \ldots, g^{(\alpha^q)}\right) \in \mathbb{G}^{q+1}$$

as input, output $e(g, g)^{1/\alpha} \in \mathbb{G}_T$. The q-ABDHI problem is identical to the q-BDHI problem, except that the former adds the term $g^{(\alpha^{-q-2})}$ to the input vector, which does not seem to help compute $e(g, g)^{1/\alpha}$. One can reduce (decision) q-ABDHI to truncated (decision) q-ABDHE simply by setting $(g', g'^{\alpha^{q+2}}) = ((g^{(\alpha^{-q-2})})^x, g^x)$ for random $x \in \mathbb{Z}_p^*$, and deriving $e(g, g)^{1/\alpha}$ as $e(g_{q+1}, g')^{1/x}$.

3 Construction I: Chosen-Plaintext Security

We now present an efficient IBE system that is ANON-IND-ID-CPA secure without random oracles under the truncated decision $(q_{\mathsf{ID}} + 1)$-ABDHE assumption.

Though this construction is substantially similar to the construction presented in Section 4.1, which is ANON-IND-ID-CCA secure, we present this construction separately because there are applications (such as searchable public key encryption [7, 1]) that only require chosen-plaintext security, and because we believe the reader may benefit from seeing this construction's (relatively) simple proof of security without being distracted by the additional machinery needed to prove chosen-ciphertext security.

3.1 Construction

Let $\mathbb{G}$ and $\mathbb{G}_T$ be groups of order p, and let $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ be the bilinear map. The IBE system works as follows.

***Setup*:** The PKG picks random generators $g, h \in \mathbb{G}$ and random $\alpha \in \mathbb{Z}_p$. It sets $g_1 = g^\alpha \in \mathbb{G}$. The public *params* and private *master-key* are given by

$$params = (g, g_1, h) \qquad master\text{-}key = \alpha \ .$$

***KeyGen*:** To generate a private key for identity $\mathsf{ID} \in \mathbb{Z}_p$, the PKG generates random $r_{\mathsf{ID}} \in \mathbb{Z}_p$, and outputs the private key

$$d_{\mathsf{ID}} = (r_{\mathsf{ID}},\ h_{\mathsf{ID}}), \quad \text{where} \quad h_{\mathsf{ID}} = (hg^{-r_{\mathsf{ID}}})^{1/(\alpha-\mathsf{ID})} \ .$$

If $\mathsf{ID} = \alpha$, the PKG aborts. We require that the PKG always use the same random value r_{ID} for ID. This can be accomplished, for example, using a PRF or an internal log to ensure consistency.

***Encrypt*:** To encrypt $m \in \mathbb{G}_T$ using identity $\mathsf{ID} \in \mathbb{Z}_p$, the sender generates random $s \in \mathbb{Z}_p$ and sends the ciphertext

$$C = (g_1^s g^{-s\cdot\mathsf{ID}},\ e(g,g)^s,\ m \cdot e(g,h)^{-s}) \ .$$

Notice that encryption does not require any pairing computations once $e(g,g)$ and $e(g,h)$ have been pre-computed. Alternatively, $e(g,g)$ and $e(g,h)$ can be included in the system parameters, in which case h can be dropped.

***Decrypt*:** To decrypt ciphertext $C = (u, v, w)$ with ID, the recipient outputs

$$m \ = \ w \cdot e(u, h_{\mathsf{ID}}) v^{r_{\mathsf{ID}}} \ .$$

Correctness: Assuming the ciphertext is well-formed for ID:

$$e(u, h_{\mathsf{ID}}) v^{r_{\mathsf{ID}}} \ = \ e(g^{s(\alpha-\mathsf{ID})}, h^{1/(\alpha-\mathsf{ID})} g^{-r_{\mathsf{ID}}/(\alpha-\mathsf{ID})}) e(g,g)^{s r_{\mathsf{ID}}} \ = \ e(g,h)^s \ ,$$

as required.

Intuitively, the recipient can decrypt because it possess a $(\alpha - \mathsf{ID})$-th root of h (after h is perturbed by $g^{-r_{\mathsf{ID}}}$). When this is paired with u, a $(\alpha - \mathsf{ID})$-th power of g^s, the recipient obtains the mask $e(g,h)^s$ after removing the perturbation.

3.2 Security

We now prove that the above IBE system is ANON-IND-ID-CPA secure under the truncated decision $(q_{\mathsf{ID}}+1)$-ABDHE assumption.

Theorem 1. *Let $q = q_{\mathsf{ID}}+1$. Assume the truncated decision (t, ϵ, q)-ABDHE assumption holds for $(\mathbb{G}, \mathbb{G}_T, e)$. Then, the above IBE system is $(t', \epsilon', q_{\mathsf{ID}})$ ANON-IND-ID-CPA secure for $t' = t - \mathcal{O}(t_{exp} \cdot q^2)$ and $\epsilon' = \epsilon + 2/p$, where t_{exp} is the time required to exponentiate in $\mathbb{G}$.*

Proof. Let $\mathcal{A}$ be an adversary that $(t', \epsilon', q_{\mathsf{ID}})$-breaks the ANON-IND-ID-CPA security of the IBE system described above. We construct an algorithm, $\mathcal{B}$, that solves the truncated decision q-ABDHE problem, as follows. $\mathcal{B}$ takes as input a random truncated decision q-ABDHE challenge $(g', g'_{q+2}, g, g_1, \ldots, g_q, Z)$, where Z is either $e(g_{q+1}, g')$ or a random element of $\mathbb{G}_T$ (recall that $g_i = g^{(\alpha^i)}$). Algorithm $\mathcal{B}$ proceeds as follows.

Setup: $\mathcal{B}$ generates a random polynomial $f(x) \in \mathbb{Z}_p[x]$ of degree q. It sets $h = g^{f(\alpha)}$, computing h from $(g, g_1, \ldots, g_q)$. It sends the public key (g, g_1, h) to $\mathcal{A}$. Since g, α, and $f(x)$ are chosen uniformly at random, h is uniformly random and this public key has a distribution identical to that in the actual construction.

Phase 1: $\mathcal{A}$ makes key generation queries. $\mathcal{B}$ responds to a query on $\mathsf{ID} \in \mathbb{Z}_p$ as follows. If $\mathsf{ID} = \alpha$, $\mathcal{B}$ uses α to solve truncated decision q-ABDHE immediately. Else, let $F_{\mathsf{ID}}(x)$ denote the $(q-1)$-degree polynomial $(f(x)-f(\mathsf{ID}))/(x-\mathsf{ID})$. $\mathcal{B}$ sets the private key $(r_{\mathsf{ID}}, h_{\mathsf{ID}})$ to be $(f(\mathsf{ID}), g^{F_{\mathsf{ID}}(\alpha)})$. This is a valid private key for ID, since $g^{F_{\mathsf{ID}}(\alpha)} = g^{(f(\alpha)-f(\mathsf{ID}))/(\alpha-\mathsf{ID})} = (hg^{-f(\mathsf{ID})})^{1/(\alpha-\mathsf{ID})}$, as required. We will describe why this private key appears to $\mathcal{A}$ to be correctly distributed below.

Challenge: $\mathcal{A}$ outputs identities $\mathsf{ID}_0, \mathsf{ID}_1$ and messages M_0, M_1. Again, if $\alpha \in \{\mathsf{ID}_0, \mathsf{ID}_1\}$, $\mathcal{B}$ uses α to solve truncated decision q-ABDHE immediately. Else, $\mathcal{B}$ generates bits $b, c \in \{0,1\}$, and computes a private key $(r_{\mathsf{ID}_b}, h_{\mathsf{ID}_b})$ for ID_b as in Phase 1. Let $f_2(x) = x^{q+2}$ and let $F_{2,\mathsf{ID}_b}(x) = (f_2(x) - f_2(\mathsf{ID}_b))/(x - \mathsf{ID}_b)$, which is a polynomial of degree $q+1$. $\mathcal{B}$ sets

$$u = g'^{f_2(\alpha)-f_2(\mathsf{ID}_b)}, \quad v = Z \cdot e(g', \prod_{i=0}^{q} g^{F_{2,\mathsf{ID}_b,i}\alpha^i}) \quad w = M_c/e(u, h_{\mathsf{ID}_b})v^{r_{\mathsf{ID}_b}} \ ,$$

where $F_{2,\mathsf{ID}_b,i}$ is the coefficient of x^i in $F_{2,\mathsf{ID}_b}(x)$. It sends (u, v, w) to $\mathcal{A}$ as the challenge ciphertext.

Let $s = (\log_g g')F_{2,\mathsf{ID}_b}(\alpha)$. If $Z = e(g_{q+1}, g')$, then $u = g^{s(\alpha-\mathsf{ID}_b)}$, $v = e(g,g)^s$, and $M_c/w = e(u, h_{\mathsf{ID}_b})v^{r_{\mathsf{ID}_b}} = e(g,h)^s$; thus (u, v, w) is a valid ciphertext for (ID_b, M_c) under randomness s. Since $\log_g g'$ is uniformly random, s is uniformly random, and so (u, v, w) is a valid, appropriately-distributed challenge to $\mathcal{A}$.

Phase 2: $\mathcal{A}$ makes key generation queries, and $\mathcal{B}$ responds as in Phase 1.

Guess: Finally, the adversary outputs guesses $b', c' \in \{0,1\}$. If $b = b'$ and $c = c'$, $\mathcal{B}$ outputs 0 (indicating that $Z = e(g_{q+1}, g')$); otherwise, it outputs 1.

Perfect Simulation: When $Z = e(g_{q+1}, g')$, the public key and challenge ciphertext issued by $\mathcal{B}$ comes from a distribution identical to that in the actual construction; however, we still must show that the private keys issued by $\mathcal{B}$ are appropriately distributed. Let $\mathcal{I}$ be a set consisting of α, ID_b, and the identities queried by $\mathcal{A}$; observe that $|\mathcal{I}| \leq q+1$. To show that the keys issued by $\mathcal{B}$ are appropriately distributed, it suffices to show that, from $\mathcal{A}$'s view, the values $\{f(a) : a \in \mathcal{I}\}$ are uniformly random and independent. But this follows from the fact that $f(x)$ is a uniformly random polynomial of degree q.

Probability Analysis: If $Z = e(g_{q+1}, g')$, then the simulation is perfect, and $\mathcal{A}$ will guess the bits (b, c) correctly with probability $1/4 + \epsilon'$. Else, Z is uniformly random, and thus (u, v) is a uniformly random and independent element of $\mathbb{G} \times \mathbb{G}_T$. In this case, the inequalities $v \neq e(u,g)^{1/(\alpha - \mathsf{ID}_0)}$ and $v \neq e(u,g)^{1/(\alpha - \mathsf{ID}_1)}$ both hold with probability $1 - 2/p$. When these inequalities hold, the value of $e(u, h_{\mathsf{ID}_b}) v^{r_{\mathsf{ID}_b}} = e(u, (hg^{-r_{\mathsf{ID}_b}})^{1/(\alpha - \mathsf{ID}_b)}) v^{r_{\mathsf{ID}_b}} = e(u,h)^{\alpha - \mathsf{ID}_b} (v/e(u,g)^{1/(\alpha - \mathsf{ID}_b)})^{r_{\mathsf{ID}_b}}$ is uniformly random and independent from $\mathcal{A}$'s view (except for the value w), since r_{ID_b} is uniformly random and independent from $\mathcal{A}$'s view (except for the value w). Thus, w is uniformly random and independent, and (u, v, w) can impart no information regarding the bits (b, c).

Assuming that no queried identity equals α (which would only increase $\mathcal{B}$'s success probability), we see that $|\Pr[\mathcal{B}(g', g'_{q+2}, g, g_1, \ldots, g_q, Z) = 0] - 1/4| \leq 2/p$ when $(g', g'_{q+2}, g, g_1, \ldots, g_q, Z)$ is sampled from $\mathcal{R}_{ABDHE}$. However, we have that $|\Pr[\mathcal{B}(g', g'_{q+2}, g, g_1, \ldots, g_q, Z) = 0] - 1/4| \geq \epsilon'$ when $(g', g'_{q+2}, g, g_1, \ldots, g_q, Z)$ is sampled from $\mathcal{P}_{ABDHE}$. Thus, for uniformly random g, g', α and Z, we have that

$$\Big| \Pr\left[\mathcal{B}\big(g', g'_{q+2}, g, g_1, \ldots, g_q, e(g_{q+1}, g')\big) = 0\right] - \Pr\left[\mathcal{B}\big(g', g'_{q+2}, g, g_1, \ldots, g_q, Z\big) = 0\right] \Big| \geq \epsilon' - 2/p \ .$$

Time-Complexity: In the simulation, $\mathcal{B}$'s overhead is dominated by computing $g^{F_{\mathsf{ID}}(\alpha)}$ in response to $\mathcal{A}$'s key generation query on ID, where $F_{\mathsf{ID}}(x)$ is a polynomial of degree $q-1$. Each such computation requires $\mathcal{O}(q)$ exponentiations in $\mathbb{G}$. Since $\mathcal{A}$ makes at most $q-1$ such queries, $t = t' + \mathcal{O}(t_{exp} \cdot q^2)$.

This concludes the proof of Theorem 1. □

3.3 Remarks on the Tightness of the Reduction

In the reduction, $\mathcal{B}$'s success probability and time complexity are the same as $\mathcal{A}$'s, except for *additive* factors depending on q. So, one could say that our IBE system has a tight security reduction in the standard model, addressing an open problem posed in [4, 5, 27]. However, it would be misleading to claim that a tight reduction from decision q-ABDHE is necessarily better than the loose reduction from decision BDH (for the IBE systems described by Boneh and Boyen [4] and Waters [27]), for a couple of reasons. First, decision q-ABDHE is a stronger assumption than decision BDH. Second, it is not even obvious what "a tight reduction from decision q-ABDHE" means, since the assumption is not fixed when q varies; it becomes stronger as the number of queries increases. Given these considerations, let's examine the significance (if any) of the "tight reduction" in closer detail.

Not much is known about the relative hardness of the decision q-ABDHE and decision BDH problems; they could be equally hard, or the former could be significantly easier. Decision q-ABDHE is a new problem, less natural and less well-studied than decision BDH, though it seems closely connected to the decision q-BDHE and decision q-BDHI problems that were used in [6, 10, 4, 20]. Interestingly, Boneh et al. [6] give some evidence that the decision q-ABDHE problem is easier to solve in the generic group model. In particular, Boneh et al. [6] show (roughly) that a generic attacker's advantage in deciding whether an element of $\mathbb{G}_T$ equals $g_1^{f(\alpha)}$ – when given oracle access to the group operation and the values $g \in G$, $g_1 \in \mathbb{G}_T$ and $g^{f_i(\alpha)} \in \mathbb{G}$ for polynomials $f_1, \ldots, f_s$ – is at most $(t + 2s + 2)^2 d/2p$, where p is the group order, t is the number of oracle queries, and $d = \max\{\deg(f), \deg(f_1), \ldots, \deg(f_s)\}$. Since $d = q$ for the decision q-ABDHE problem, Boneh et al.'s result suggests that a generic attacker's advantage in decision q-ABDHE may be about q times greater than in decision BDH (for fixed t and p, and assuming $q \ll t$). This factor of q seems to offset the factor of q that we eliminated by making our reduction tight. On the other hand, this generic-group result doesn't tell us much about relative hardness of the decision q-ABDHE and decision BDH problems in the real world, since the fastest algorithms for solving them are likely non-generic (and sub-exponential). Ultimately, it is unclear whether or not our tighter reduction under a stronger assumption improves security.

However, for the sake of argument, let's try to assess the impact of our tighter reduction under the assumption that the decision q-ABDHE and decision BDH problems are equally hard. Since it is not very useful simply to characterize our reduction as "tight" asymptotically, let's make such a statement more precise by fixing reasonable values of q and assessing the security and efficiency implications concretely. Suppose that we want to choose our security parameter such that, to succeed with probability at least ϵ', the time complexity of $\mathcal{A}$'s attack must be 2^{100}. Suppose also that it is infeasible for $\mathcal{A}$ make more than 2^{30} key generation queries, and that $t_{exp} = 2^{30}$. In this case, we should choose our security parameter such that $t = 2^{100} + \mathcal{O}(t_{exp} \cdot q^2)$. Since 2^{90} is much smaller than 2^{100}, it essentially suffices to choose the security parameter such that $t \approx 2^{100}$.

On the other hand, consider an IBE system whose reduction loses a *multiplicative* factor of q in time-complexity (without much loss in the success probability). In this setting, to ensure that $\mathcal{A}$'s time complexity is 2^{100}, we must choose our security parameter such that $t \approx 2^{130}$. The security parameter in this setting thus must be at least 30% greater (even more if sub-exponential attacks are possible against the system). Assuming, as a rough approximation, that exponentiation takes time proportional to the cube of the security parameter, the increase in the security parameter size more than doubles the time needed to exponentiate, which significantly impacts the computational efficiency of the system. So, our "tight reduction" significantly enhances the efficiency advantages of our system over previous IBE systems that have been proven fully secure in the standard model (under decision BDH), at least when we assume that decision q-ABDHE and decision BDH are equally hard.

Since the relative hardness of decision q-ABDHE and decision BDH is unknown, however, we stress that it remains an excellent open problem to construct an IBE system that has a tight reduction in the standard model under a more natural assumption, such as decision BDH.

4 Construction II: Chosen-Ciphertext Security

We now present an efficient IBE system that is ANON-IND-ID-CCA secure without random oracles under the truncated decision $(q_{\mathsf{ID}} + 2)$-ABDHE assumption.

4.1 Construction

Let $\mathbb{G}$ and $\mathbb{G}_T$ be groups of order p, and let $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ be the bilinear map. The IBE system works as follows.

***Setup*:** The PKG picks a random generators $g, h_1, h_2, h_3 \in \mathbb{G}$ and a random $\alpha \in \mathbb{Z}_p$. It sets $g_1 = g^\alpha \in \mathbb{G}$. It chooses a hash function H from a family of universal one-way hash functions. The public *params* and private *master-key* are given by

$$params = (g, g_1, h_1, h_2, h_3, H) \qquad master\text{-}key = \alpha \ .$$

***KeyGen*:** To generate a private key for identity $\mathsf{ID} \in \mathbb{Z}_p$, the PKG generates random $r_{\mathsf{ID},i} \in \mathbb{Z}_p$ for $i \in \{1, 2, 3\}$, and outputs the private key

$$d_{\mathsf{ID}} = \{(r_{\mathsf{ID},i},\ h_{\mathsf{ID},i}) : i \in \{1, 2, 3\}\}, \quad \text{where} \quad h_{\mathsf{ID},i} = (h_i g^{-r_{\mathsf{ID},i}})^{1/(\alpha - \mathsf{ID})} \ .$$

If $\mathsf{ID} = \alpha$, the PKG aborts. As before, we require that the PKG always use the same random values $\{r_{\mathsf{ID},i}\}$ for ID.

***Encrypt*:** To encrypt $m \in \mathbb{G}_T$ using identity $\mathsf{ID} \in \mathbb{Z}_p$, the sender generates random $s \in \mathbb{Z}_p$ and sends the ciphertext

$$C = (g_1^s g^{-s \cdot \mathsf{ID}},\ e(g, g)^s,\ m \cdot e(g, h_1)^{-s},\ e(g, h_2)^s e(g, h_3)^{s\beta}) \ .$$

Above, for $C = (u, v, w, y)$, we set $\beta = H(u, v, w)$. As before, encryption does not require any pairing computations once $e(g, g)$, and $\{e(g, h_i)\}$ have been pre-computed or alternatively included in *params*.

***Decrypt*:** To decrypt ciphertext $C = (u, v, w, y)$ with ID, the recipient sets $\beta = H(u, v, w)$ and tests whether

$$y = e(u, h_{\mathsf{ID},2} {h_{\mathsf{ID},3}}^{\beta}) v^{r_{\mathsf{ID},2} + r_{\mathsf{ID},3}\beta} .$$

If the check fails, the recipient outputs $\perp$. Otherwise, it outputs

$$m = w \cdot e(u, h_{\mathsf{ID},1}) v^{r_{\mathsf{ID},1}} .$$

Correctness: Assuming the ciphertext is well-formed for ID:

$$\begin{aligned} & e(u, h_{\mathsf{ID},2} {h_{\mathsf{ID},3}}^{\beta}) v^{r_{\mathsf{ID},2} + r_{\mathsf{ID},3}\beta} \\ & \quad = e(g^{s(\alpha - \mathsf{ID})}, (h_2 {h_3}^{\beta})^{1/(\alpha - \mathsf{ID})} g^{-(r_{\mathsf{ID},2} + r_{\mathsf{ID},3}\beta)/(\alpha - \mathsf{ID})}) e(g, g)^{s(r_{\mathsf{ID},2} + r_{\mathsf{ID},3}\beta)} \\ & \quad = e(g^{s(\alpha - \mathsf{ID})}, (h_2 {h_3}^{\beta})^{1/(\alpha - \mathsf{ID})}) = e(g, h_2)^s e(g, h_3)^{s\beta} . \end{aligned}$$

Thus, the check passes. Moreover, as in the $\mathsf{ANON\text{-}IND\text{-}ID\text{-}CPA}$ scheme,

$$e(u, h_{\mathsf{ID},1}) v^{r_{\mathsf{ID},1}} = e(g^{s(\alpha - \mathsf{ID})}, h_1^{1/(\alpha - \mathsf{ID})} g^{-r_{\mathsf{ID},1}/(\alpha - \mathsf{ID})}) e(g, g)^{s r_{\mathsf{ID},1}} = e(g, h_1)^s ,$$

as required.

4.2 Security

We now prove that the above construction is $\mathsf{ANON\text{-}IND\text{-}ID\text{-}CCA}$ secure under the truncated decision $(q_{\mathsf{ID}} + 2)$-ABDHE assumption. We will refer the reader to the proof of Theorem 1 for some portions of the present proof that would otherwise be duplicative.

Theorem 2. *Let $q = q_{\mathsf{ID}} + 2$. Assume the truncated decision (t, ϵ, q)-ABDHE assumption holds for $(\mathbb{G}, \mathbb{G}_T, e)$. Then, the above IBE system is $(t', \epsilon', q_{\mathsf{ID}}, q_C)$ $\mathsf{ANON\text{-}IND\text{-}ID\text{-}CCA}$ secure for $t' = t - \mathcal{O}(t_{exp} \cdot q^2)$ and $\epsilon' = \epsilon + 4q_C/p$, where t_{exp} is the time required to exponentiate in $\mathbb{G}$.*

Proof. Let $\mathcal{A}$ be an adversary that $(t', \epsilon', q_{\mathsf{ID}}, q_C)$-breaks the $\mathsf{ANON\text{-}IND\text{-}ID\text{-}CCA}$ security of the IBE system described above. We construct an algorithm, $\mathcal{B}$, that solves the truncated decision q-ABDHE problem, as follows. $\mathcal{B}$ takes as input a random truncated decision q-ABDHE challenge $(g', g'_{q+2}, g, g_1, \ldots, g_q, Z)$, where Z is either $e(g_{q+1}, g')$ or a random element of $\mathbb{G}_T$. Algorithm $\mathcal{B}$ proceeds as follows.

Setup: $\mathcal{B}$ generates random polynomials $f_i(x) \in \mathbb{Z}_p[x]$ of degree q for $i \in \{1, 2, 3\}$. It sets $h_i = g^{f_i(\alpha)}$. It sends the public key (g, g_1, h_1, h_2, h_3) to $\mathcal{A}$. Since g, α, and $f_i(x)$ for $i \in \{1, 2, 3\}$ are chosen uniformly at random, h_1, h_2, and h_3 are uniformly random and the public key has a distribution identical to that in the actual construction.

Phase 1: $\mathcal{A}$ makes key generation queries. $\mathcal{B}$ responds to a query on $\mathsf{ID} \in \mathbb{Z}_p$ as follows. If $\mathsf{ID} = \alpha$, $\mathcal{B}$ uses α to solve truncated decision q-ABDHE immediately. Else, to generate a pair $(r_{\mathsf{ID},1}, h_{\mathsf{ID},1})$ such that $h_{\mathsf{ID},1} = (h_1 g^{-r_{\mathsf{ID},1}})^{1/(\alpha+\mathsf{ID})}$, $\mathcal{B}$ sets $r_{\mathsf{ID},1} = f_1(\mathsf{ID})$ and computes $h_{\mathsf{ID},1}$ as before (in the proof of Theorem 1). It computes the remainder of the private key similarly. As before, the private key generated for ID in this fashion is valid.
$\mathcal{A}$ also makes decryption queries. To respond to a decryption query on (ID, C), $\mathcal{B}$ generates a private key for ID as above. It then decrypts C by performing the usual *Decrypt* algorithm with this private key.

Challenge: As before, $\mathcal{A}$ outputs identities $\mathsf{ID}_0, \mathsf{ID}_1$ and messages M_0, M_1. If $\alpha \in \{\mathsf{ID}_0, \mathsf{ID}_1\}$, $\mathcal{B}$ uses α to solve truncated decision q-ABDHE immediately. Else, as before, $\mathcal{B}$ generates bits $b, c \in \{0, 1\}$. After computing a private key $\{(r_{\mathsf{ID},i}, h_{\mathsf{ID},i}) : i \in \{1, 2, 3\}\}$ for ID_b, it also computes (u, v, w) as before, using the $(r_{\mathsf{ID}_b,1}, h_{\mathsf{ID}_b,1})$ portion of the key to compute w. After setting $\beta = H(u, v, w)$, $\mathcal{B}$ sets $y = e(u, h_{\mathsf{ID},2} h_{\mathsf{ID},3}{}^{\beta}) v^{r_{\mathsf{ID},2} + r_{\mathsf{ID},3}\beta}$. If $Z = e(g_{q+1}, g')$, then (u, v, w, y) is a valid, appropriately-distributed challenge to $\mathcal{A}$ for essentially the same reason as before.

Phase 2: $\mathcal{A}$ makes key generation and decryption queries, and $\mathcal{B}$ responds as in Phase 1.

Guess: As before.

Now, since the time-complexity analysis is as in the proof of Theorem 1, Theorem 2 follows from the following lemmata.

Lemma 1. *When $\mathcal{B}$'s input is sampled according to $\mathcal{P}_{ABDHE}$, the joint distribution of $\mathcal{A}$'s view and the bits (b, c) is indistinguishable from that in the actual construction, except with probability $2q_C/p$.*

Lemma 2. *When $\mathcal{B}$'s input is sampled according to $\mathcal{R}_{ABDHE}$, the distribution of the bits (b, c) is independent from the adversary's view, except with probability $2q_C/p$.*

Our approach to proving these claims closely follows the proof of security for the Cramer-Shoup encryption scheme [17], in that both proofs rely heavily on the notion of *linear independence*. More specifically, when one expresses the adversary's knowledge (from the public key, queries, etc.) as equations in the simulator's private key variables, one may ask whether a target equation that the adversary is trying to solve is linearly independent to the equations in its knowledge base; if so, then in certain circumstances, the adversary can be said to have an unconditionally negligible probability of finding a solution to the target equation. This will become clearer below.

Proof of Lemma 1: When $\mathcal{B}$'s input is sampled according to $\mathcal{P}_{ABDHE}$, $\mathcal{B}$'s simulation appears perfect to $\mathcal{A}$ if $\mathcal{A}$ makes only key generation queries, as in the proof of Theorem 1. $\mathcal{B}$'s simulation still appears perfect if $\mathcal{A}$ makes decryption queries only on identities for which it queries the private key, since $\mathcal{B}$'s responses give $\mathcal{A}$ no additional information. Furthermore, querying well-formed ciphertexts to the decryption oracle does not help $\mathcal{A}$ distinguish between the simulation and the actual construction, since, by the correctness of *Decrypt*, well-formed ciphertexts will be accepted in either case. Finally, querying a non-well-formed ciphertext (u', v', w', y') for ID for which $v' = e(u', g)^{1/(\alpha - \mathsf{ID})}$ does not help $\mathcal{A}$ distinguish, since this ciphertext will fail the *Decrypt* check under *every* valid private key for ID. Thus, the lemma follows from the following claim:

Claim: The decryption oracle, in the simulation and in the actual construction, rejects all invalid ciphertexts under identities not queried by $\mathcal{A}$, except with probability q_C/p.

We say a ciphertext (u', v', w', y') for ID is "invalid" if $v' \neq e(u', g)^{1/(\alpha - \mathsf{ID})}$.

Let (u', v', w', y') be an invalid ciphertext queried by $\mathcal{A}$ for ID, an identity not queried by $\mathcal{A}$. Let $\{(r_{\mathsf{ID},i}, h_{\mathsf{ID},i}) : i \in \{1, 2, 3\}\}$ be $\mathcal{B}$'s private key for ID. Let $a_{u'} = \log_g u'$, $a_{v'} = \log_{e(g,g)} v'$, and $a_{y'} = \log_{e(g,g)} y'$. For (u', v', w', y') to be accepted, we must have $y' = e(u', h_{\mathsf{ID},2} h_{\mathsf{ID},3}{}^{\beta'}) v'^{r_{\mathsf{ID},2} + r_{\mathsf{ID},3}\beta'}$ – i.e.,

$$a_{y'} = a_{u'}(\log_g h_{\mathsf{ID},2} + \beta' \log_g h_{\mathsf{ID},3}) + a_{v'}(r_{\mathsf{ID},2} + \beta' r_{\mathsf{ID},3}) \ , \tag{1}$$

for $\beta' = H(u', v', w')$. To compute the probability that $\mathcal{A}$ can generate such a y', we must consider the distribution of $\{(r_{\mathsf{ID},i}, h_{\mathsf{ID},i}) : i \in \{2, 3\}\}$ from $\mathcal{A}$'s view.

First, $\mathcal{A}$ knows that

$$\log_g h_1 = (\alpha - \mathsf{ID}) \log_g h_{\mathsf{ID},1} + r_{\mathsf{ID},1} \tag{2}$$

$$\log_g h_2 = (\alpha - \mathsf{ID}) \log_g h_{\mathsf{ID},2} + r_{\mathsf{ID},2} \tag{3}$$

$$\log_g h_3 = (\alpha - \mathsf{ID}) \log_g h_{\mathsf{ID},3} + r_{\mathsf{ID},3} \tag{4}$$

by the construction of the private key. In light of Equations 3 and 4, $\mathcal{A}$'s task may be re-phrased as finding a y' such that

$$a_{y'} = (a_{u'}/(\alpha - \mathsf{ID}))(\log_g h_2 + \beta' \log_g h_3) + (a_{v'} - a_{u'}/(\alpha - \mathsf{ID}))(r_{\mathsf{ID},2} + \beta' r_{\mathsf{ID},3}) \ . \tag{5}$$

Note that $a_{v'} - a_{u'}/(\alpha - \mathsf{ID}) \neq 0$, since the ciphertext is invalid. Let $z' = a_{v'} - a_{u'}/(\alpha - \mathsf{ID})$.

In the actual construction, the values of $r_{\mathsf{ID},i}$ for $i \in \{2, 3\}$ are chosen independently for different identities; however, this is not true in the simulation. Since $f_i(\mathsf{ID}) = r_{\mathsf{ID},i}$, $\mathcal{A}$ could conceivably gain information regarding $(r_{\mathsf{ID},2}, r_{\mathsf{ID},3})$ from its information regarding $(f_2(x), f_3(x))$, which includes the evaluations of $(f_2(x), f_3(x))$ at α (from the public key components (h_2, h_3)) and at $q - 2$ identities (from its key generation queries). We may represent the knowledge gained from these evaluations as a matrix product:

$$[f_{2,0}, f_{2,1}, \ldots, f_{2,q}, f_{3,0}, f_{3,1}, \ldots, f_{3,q}] \begin{bmatrix} 1 & 1 & \cdots & 1 & 0 & 0 & \cdots & 0 \\ x_1 & x_2 & \cdots & x_{q-1} & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_1^q & x_2^q & \cdots & x_{q-1}^q & 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & 1 & 1 & \cdots & 1 \\ 0 & 0 & \cdots & 0 & x_1 & x_2 & \cdots & x_{q-1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & x_1^q & x_2^q & \cdots & x_{q-1}^q \end{bmatrix} ,$$

where $f_{i,j}$ is the coefficient of x^j in $f_i(x)$, $x_k \in \mathbb{Z}_p$ is the k-th identity queried by $\mathcal{A}$ to the key generation oracle, and $x_{q-1} = \alpha$. Let $\boldsymbol{f}$ denote the vector on the left and let V denote the matrix on the right. Note that V contains two $(q+1) \times (q-1)$ Vandermonde matrices; its columns are linearly independent. From $\mathcal{A}$'s view, since V has four more rows than columns, the solution space for $\boldsymbol{f}$ is four-dimensional.

Let $\boldsymbol{\gamma}_{\mathsf{ID}}$ denote the vector $(1, \mathsf{ID}, \ldots, \mathsf{ID}^q)$. When we re-phrase Equation 5 in terms of the simulator's private key vector $\boldsymbol{f}$, we obtain:

$$a_{y'} \ = \ \text{“public” terms} \ + \ z'(\boldsymbol{f} \ \cdot \ \boldsymbol{\gamma}_{\mathsf{ID}} \| \beta' \boldsymbol{\gamma}_{\mathsf{ID}}) \ , \tag{6}$$

where "$\cdot$" denotes the dot product and $\boldsymbol{\gamma}_{\mathsf{ID}} \| \beta' \boldsymbol{\gamma}_{\mathsf{ID}}$ denotes the $2(q+1)$-dimensional vector formed by concatenating the coefficients of $\boldsymbol{\gamma}_{\mathsf{ID}}$ and $\beta' \boldsymbol{\gamma}_{\mathsf{ID}}$. If $\boldsymbol{\gamma}_{\mathsf{ID}} \| \beta' \boldsymbol{\gamma}_{\mathsf{ID}}$ were in the linear span of V, then potentially $\mathcal{A}$ could use knowledge gained from its key generation queries to compute a solution y' to Equation 6. However, one can easily see that $\boldsymbol{\gamma}_{\mathsf{ID}} \| \beta' \boldsymbol{\gamma}_{\mathsf{ID}}$ is linearly independent. Thus, as in the security proof of Cramer-Shoup, it follows that the decryption oracle will reject (u', v', w', y') for ID with probability $1 - 1/p$ if it is the first invalid ciphertext queried by $\mathcal{A}$, since there is only a $1/p$ chance that $\boldsymbol{f}$ is contained in the 3-dimensional solution space (with p^3 points) defined by Equation 6 and the columns of V, given that $\boldsymbol{f}$ is in the 4-dimensional solution space (with p^4 points) defined by the columns of V.

Each time the decryption oracle rejects an invalid ciphertext in the simulation, the solution space for $\boldsymbol{f}$ is "punctured" in a 3-dimensional space that $\mathcal{A}$ then concludes does not contain $\boldsymbol{f}$; consequently, the probability that $\mathcal{A}$'s i-th invalid ciphertext is accepted is at most $1/(p - i + 1)$. The probability that q_C invalid ciphertexts (on identities not queried to the key generation oracle) are all rejected is at least $1 - q_C/p$. This bound also holds for the actual construction (where $\mathcal{A}$'s attack is less effective). This concludes the proof of Lemma 1.

Proof of Lemma 2: The lemma follows from the following two claims.

Claim 1: If the decryption oracle rejects all invalid ciphertexts, then $\mathcal{A}$ has advantage at most q_C/p in guessing the bits (b, c).

Claim 2: The decryption oracle rejects all invalid ciphertexts, except with probability q_C/p.

Let $a_u = \log_g u$, $a_v = \log_{e(g,g)} v$ and $a_y = \log_{e(g,g)} y$ for challenge ciphertext (u, v, w, y) on (ID_b, M_c). Since (u, v, w, y) is generated by sampling from $\mathcal{R}_{ABDHE}$ in this case, (a_u, a_v) is a uniformly random element of $\mathbb{Z}_p \times \mathbb{Z}_p$ in $\mathcal{A}$'s view. From the challenge ciphertext and Equations 2-4, $\mathcal{A}$ obtains the equations

$$\log(M_c/w) = (a_u/(\alpha - \mathsf{ID}_b)) \log h_1 + (a_v - a_u/(\alpha - \mathsf{ID}_b)) r_{\mathsf{ID}_b,1} \tag{7}$$

$$a_y = (a_u/(\alpha - \mathsf{ID}_b))(\log_g h_2 + \beta \log_g h_3) + (a_v - a_u/(\alpha - \mathsf{ID}_b))(r_{\mathsf{ID}_b,2} + \beta r_{\mathsf{ID}_b,3}) \tag{8}$$

where $\beta = H(u, v, w)$.

Regarding Claim 1, if no invalid ciphertexts are accepted, then $\mathcal{B}$'s responses to decryption queries leak no information about $r_{\mathsf{ID}_b,1}$. Furthermore, $\mathcal{A}$'s key generation queries do not constrain $r_{\mathsf{ID}_b,1} = f_1(\mathsf{ID}_b)$, since f_1 is of degree q. Thus the distribution of M_c/w – conditioning on (b,c) and everything in $\mathcal{A}$'s view other than w – is uniform. As in Cramer-Shoup, M_c/w serves as a perfect one-time pad; w is uniformly random and independent, and c is independent of $\mathcal{A}$'s view.

The only part of the ciphertext that can reveal information about b is y, since $\mathcal{A}$ views (u,v,w) as a uniformly random and independent element of $\mathbb{G}\times\mathbb{G}_T\times\mathbb{G}_T$. The $2q-2$ equations corresponding to the columns of V intersect Equation 8 in at least a three-dimensional space in $\mathbb{Z}_p^{2(q+1)}$. $\mathcal{A}$ views $\boldsymbol{f}$ as being contained in one of two three-dimensional spaces, since b has two possible values. By an argument similar to above, each of $\mathcal{A}$'s invalid ciphertext queries punctures each of these three-dimensional spaces in a plane, removing each of the two planes from consideration as containing $\boldsymbol{f}$. Since no invalid ciphertext is accepted, each three-dimensional space is left with at least $p^3 - q_C p^2$ (out of p^3) candidates. Thus, $\mathcal{A}$ cannot distinguish b, except with advantage at most q_C/p.

Regarding Claim 2, suppose that $\mathcal{A}$ submits an invalid ciphertext (u',v',w',y') for unqueried identity ID, where $(u',v',w',y',\mathsf{ID}) \neq (u,v,w,y,\mathsf{ID}_b)$. Let $\beta' = H(u',v',w')$. There are three cases to consider:

1. $(u',v',w') = (u,v,w)$: In this case, the hashes are also equal. If $\mathsf{ID} = \mathsf{ID}_b$ but $y' \neq y$, the ciphertext will certainly be rejected. If $\mathsf{ID} \neq \mathsf{ID}_b$, $\mathcal{A}$ must generate a y' that satisfies Equation 6. However, we claim that the vector $\boldsymbol{\gamma}_{\mathsf{ID}}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}}$ (corresponding to Equation 6) is linearly independent in $\mathbb{Z}_p^{2(q+1)}$ to $\boldsymbol{\gamma}_{\mathsf{ID}_b}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}_b}$ (corresponding to the challenge ciphertext) and the columns of V, implying (via arguments analogous to those above) that $\mathcal{A}$ cannot generate such a y' except with probability $1/(p-i+1)$, where (u',v',w',y') is the i-th invalid ciphertext. Let $V_1,\ldots,V_{2q-2}$ be the columns of V. Suppose that there exist integers $(a_1,\ldots,a_{2q})$, not all zero, such that $a_1V_1 + \cdots + a_{2q-2}V_{2q-2} + a_{2q-1}(\boldsymbol{\gamma}_{\mathsf{ID}}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}}) + a_{2q}(\boldsymbol{\gamma}_{\mathsf{ID}_b}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}_b})$ is the zero vector in $\mathbb{Z}_p^{2(q+1)}$. Then, either $(a_1,\ldots,a_{q-1},a_{2q-1},a_{2q})$ or $(a_q,\ldots,a_{2q-2},a_{2q-1},a_{2q})$ is not all zeros; wlog, assume the former. The first $q+1$ coordinates of the vectors $(V_1,\ldots,V_{q-1},\boldsymbol{\gamma}_{\mathsf{ID}},\boldsymbol{\gamma}_{\mathsf{ID}_b})$ form a Vandermonde matrix (with nonzero determinant), but the first $q+1$ coordinates of $a_1V_1+\cdots+a_{q-1}V_{q-1} + a_{2q-1}(\boldsymbol{\gamma}_{\mathsf{ID}}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}}) + a_{2q}(\boldsymbol{\gamma}_{\mathsf{ID}_b}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}_b})$ is the zero vector in $\mathbb{Z}_p^{q+1}$ – a contradiction.
2. $(u',v',w') \neq (u,v,w)$ and $\beta' = \beta$: This violates the universal one-wayness of the hash function H, by an argument analogous to that in Cramer-Shoup.
3. $(u',v',w') \neq (u,v,w)$ and $\beta' \neq \beta$: In this case, $\mathcal{A}$ must generate, for some ID, a y' that satisfies Equation 6. For essentially the same reason as discussed in Item 1, $\mathcal{A}$ can do this with only negligible probability when $\mathsf{ID} \neq \mathsf{ID}_b$. If $\mathsf{ID} = \mathsf{ID}_b$, then $\boldsymbol{\gamma}_{\mathsf{ID}}\|\beta'\boldsymbol{\gamma}_{\mathsf{ID}}$ and $\boldsymbol{\gamma}_{\mathsf{ID}_b}\|\beta\boldsymbol{\gamma}_{\mathsf{ID}_b}$ generate $\boldsymbol{\gamma}_{\mathsf{ID}_b}\|0^{q+1}$ and $0^{q+1}\|\boldsymbol{\gamma}_{\mathsf{ID}_b}$ since $\beta \neq \beta'$. These vectors are clearly linearly independent to each other and the columns of V, and thus the standard analysis applies.

This completes the proof of Lemma 2.

5 Conclusions and Open Problems

We presented a fully secure IBE system that is quite practical, has very compact public parameters, and has a tight security reduction (though based on a stronger assumption that depends on the anticipated number of private key generation queries). The scheme is recipient-anonymous, and its proof extends Cramer-Shoup-type techniques to IBE systems.

Since a tight reduction based on decision q-ABDHE is not necessarily better than a loose reduction based on decision BDH (or some other natural assumption), it remains an outstanding open problem to construct a fully secure IBE system (without random oracles) that has a tight reduction based on a more natural assumption. Another interesting problem is to construct a hierarchical IBE system that has a reduction based on a reasonable assumption, either in the standard model or the random oracle model, that is polynomial in q and the number of levels.

Acknowledgments

We thank Dan Boneh, Brent Waters and the anonymous reviewers of Eurocrypt 2006 for insightful comments and helpful suggestions.

References

[1] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, and H. Shi. Searchable Encryption Revisited: Consistency Properties, Relation to Anonymous IBE, and Extensions. In *Advances in Cryptology – Crypto 2005*, volume 3621 of *LNCS*, pages 205–222. Springer-Verlag, 2005.

[2] N. Attrapadung, B. Chevallier-Mames, J. Furukawa, T. Gomi, G. Hanaoka, H. Imai, and R. Zhang. Efficient Identity Based Encryption with Tight Security Reduction. Cryptology ePrint Archive 2005/320.

[3] M. Bellare and P. Rogaway. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *Proc. of ACM CCS*, pages 62–73, 1993.

[4] D. Boneh and X. Boyen. Efficient Selective-ID Identity Based Encryption without Random Oracles. In *Advances in Cryptology – Eurocrypt 2004*, volume 3027 of *LNCS*, pages 223–238. Springer-Verlag, 2004.

[5] D. Boneh and X. Boyen. Secure Identity Based Encryption without Random Oracles. In *Advances in Cryptology – Crypto 2004*, volume 3152 of *LNCS*, pages 443–459. Springer-Verlag, 2004.

[6] D. Boneh, X. Boyen, and E.-J. Goh. Hierarchical Identity Based Encryption with Constant Size Ciphertext. In *Advances in Cryptology – Eurocrypt 2005*, volume 3494 of *LNCS*, pages 440–456. Springer-Verlag, 2005.

[7] D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano. Public Key Encryption with Keyword Search. In *Advances in Cryptology – Eurocrypt 2004*, volume 3027 of *LNCS*, pages 506–522. Springer-Verlag, 2004.

[8] D. Boneh and M. Franklin. Identity Based Encryption from the Weil pairing. In *Advances in Cryptology – Crypto 2001*, volume 2139 of *LNCS*, pages 213–229. Springer-Verlag, 2001.

[9] D. Boneh and M. Franklin. Identity Based Encryption from the Weil pairing. *SIAM Journal of Computing*, 32(3):586–615, 2003.

[10] D. Boneh, C. Gentry, and B. Waters. Collusion-Resistant Broadcast Encryption with Short Ciphertexts and Private Keys. In *Advances in Cryptology – Crypto 2005*, volume 3621 of *LNCS*, pages 258–275. Springer-Verlag, 2005.

[11] D. Boneh and J. Katz. Improved Efficiency for CCA-Secure Cryptosystems Built Using Identity Based Encryption. In *Proc. of CT-RSA*, volume 3376 of *LNCS*, pages 87–103. Springer-Verlag, 2005.

[12] D. Boneh, B. Lynn, and H. Shacham, *Short signatures from the Weil pairing*, Advances in Cryptology — Asiacrypt 2001, Lecture Notes in Computer Science 2248 (2001), Springer, 514–532.

[13] X. Boyen, Q. Mie and B. Waters, *Direct Chosen Ciphertext Security from Identity Based Techniques*, In *Proc. of ACM CCS*, pages 320–329, 2005.

[14] X. Boyen and B. Waters. Anonymous Hierarchical Identity-Based Encryption (without Random Oracles). Cryptology ePrint Archive 2006/085.

[15] R. Canetti, S. Halevi, and J. Katz. A Forward-Secure Public-Key Encryption Scheme. In *Advances in Cryptology – Eurocrypt 2003*, volume 2656 of *LNCS*, pages 255–271. Springer-Verlag, 2003.

[16] R. Canetti, S. Halevi, and J. Katz. Chosen-Ciphertext Security from Identity-Based Encryption. In *Advances in Cryptology – Eurocrypt 2004*, volume 3027 of *LNCS*, pages 207-222. Springer-Verlag, 2004.

[17] R. Cramer and V. Shoup. A Practical Public Key Cryptosystem Provably Secure Against Adaptive Chosen Ciphertext Attacks. In *Advances in Cryptology – Crypto 1998*, volume 1462 of *LNCS*, pages 13–25. Springer-Verlag, 1998.

[18] R. Cramer and V. Shoup. Signature Schemes Based on the Strong RSA Assumption. In *Proc. of ACM CCS*, pages 46–51, 1999.

[19] Y. Dodis. Efficient Construction of (Distributed) Verifiable Random Functions. In *Proc. of Public Key Cryptography*, volume 2567 of *LNCS*, pages 1–17. Springer-Verlag, 2002.

[20] Y. Dodis and A. Yampolskiy. A Verifiable Random Function with Short Proofs and Keys. In *Proc. of Public Key Cryptography*, volume 3386 of *LNCS*, pages 416–431. Springer-Verlag, 2005.

[21] C. Gentry and A. Silverberg. Hierarchical ID-Based Cryptography. In *Advances in Cryptology – Asiacrypt 2002*, volume 2501 of *LNCS*, pages 548–566. Springer-Verlag, 2002.

[22] J. Horwitz and B. Lynn. Toward Hierarchical Identity-Based Encryption. In *Advances in Cryptology – Eurocrypt 2002*, volume 2332 of *LNCS*, pages 466–481. Springer-Verlag, 2002.

[23] J. Katz and N. Wang. Efficiency Improvements for Signature Schemes with Tight Security Reductions. In *Proc. of ACM CCS*, pages 155–164, 2003.

[24] K. Kurosawa and Y. Desmedt. A New Paradigm of Hybrid Encryption Scheme. In *Advances in Cryptology – Crypto 2004*, volume 3152 of *LNCS*, pages 426–442. Springer-Verlag, 2004.

[25] A. Shamir. Identity-Based Cryptosystems and Signature Schemes. In *Advances in Cryptology – Crypto 1984*, volume 196 of *LNCS*, pages 47–53. Springer-Verlag, 1984.

[26] V. Shoup. Lower Bounds for Discrete Logarithms and Related Problems. In *Advances in Cryptology – Eurocrypt 1997*, volume 1233 of *LNCS*, pages 256–266. Springer-Verlag, 1997.

[27] B. Waters. Efficient Identity-Based Encryption without Random Oracles. In *Advances in Cryptology – Eurocrypt 2005*, volume 3494 of *LNCS*, pages 114–127. Springer-Verlag, 2005.