Universally Anonymous IBE based on the Quadratic Residuosity Assumption*

Giuseppe Ateniese
Johns Hopkins University
ateniese@cs.jhu.edu

Paolo Gasti
University of Genoa
gasti@disi.unige.it

Abstract

We introduce the first universally anonymous, thus key-private, IBE whose security is based on the standard quadratic residuosity assumption. Our scheme is a variant of Cocks IBE (which is not anonymous) and is efficient and highly parallelizable.

1 Introduction

Identity-based encryption was introduced by Shamir in 1984 [24]. He asked whether it was possible to encrypt a message by just using the identity of the intended recipient. Several partial and inefficient solutions were proposed after Shamir's initial challenge but it was only in 2000 that Sakai et al. [22], Boneh and Franklin [9], and Cocks [12] came up with very practical solutions.

The Boneh-Franklin work has been the most influential of all: it did not just introduce the first practical IBE scheme but, more importantly, it provided appropriate assumptions and definitions and showed how to pick the right curves, how to encode and map elements into points, etc.

Cocks' scheme is *per se* revolutionary: it is the first IBE that does not use pairings but rather it works in standard RSA groups and its security relies on the standard quadratic residuosity assumption (within the random oracle model). Cocks IBE, however, encrypts the message bit by bit and thus it is considered very bandwidth consuming. On the other end, Cocks [12] observes that his scheme can be used in practice to encrypt short session keys in which case the scheme becomes very attractive. We may add that the importance of relying on such a standard assumption should not be underestimated. In fact, this is what motivated the recent work of Boneh, Gentry, and Hamburg [10] where a new space-efficient IBE scheme is introduced whose security is also based on the quadratic residuosity assumption. Unfortunately, as the authors point out [10], their scheme is not efficient, it is more expensive than Cocks IBE and in fact it is more expensive than all standard IBE and public-key encryption schemes since its complexity is quartic in the security parameter (in particular, the encryption algorithm may take several seconds to complete even on a fast machine).

However, the scheme of Boneh et al. [10] has an important advantage over the scheme of Cocks: it provides anonymity, i.e., nobody can tell who the intended recipient is by just looking at the ciphertext. Anonymity, or key-privacy, is a very important property that was first studied by Bellare et al. [5]. Recipient anonymity can be used, for example, to thwart traffic analysis, to

*This is the extended version of the paper that appears in CT-RSA, The Cryptographers' Track at the RSA Conference 2009 [3] enable searching on encrypted data [8], or to anonymously broadcast messages [1]. Several IBE schemes provide anonymity, for instance the Boneh-Franklin scheme is anonymous. Other schemes that do not originally provide anonymity can be either properly modified [11] or adapted to work in the XDH setting [23, 7, 4, 2].

At this point, it is natural to ask whether it is possible to enhance Cocks IBE and come up with a variant that provides anonymity and that, unlike Boneh et al.'s scheme [10], is as efficient as the original scheme of Cocks.

The first attempt in this direction has been proposed recently by Di Crescenzo and Saraswat [15]. They provide the first public-key encryption with keyword search (PEKS) that is not based on pairings. Although their scheme is suitable for PEKS, we note that when used as an IBE it becomes quite impractical: it uses four times the amount of bandwidth required by Cocks and it requires each user to store and use a very large number of secret keys (four keys per each bit of the plaintext). In addition, the security of their scheme is based on a new assumption they introduce, but we can show that their assumption is equivalent to the standard quadratic residuosity one.

Universal anonymity is a new and exciting notion introduced at Asiacrypt 2005 by Hayashi and Tanaka [19]. An encryption scheme is universally anonymous if ciphertexts can be made anonymous by anyone and not just by whoever created the ciphertexts. Specifically, a universally anonymizable public-key encryption scheme consists of a standard public-key encryption scheme and two additional algorithms: one is used to anonymize ciphertexts, which takes as input only the public key of the recipient, and the other is used by the recipient to decrypt anonymized ciphertexts. As described in [19], a company may mandate that internal emails should be encrypted but not anonymous (for efficiency and legal requirements) when routed internally but be anonymized whenever they are sent outside the domain of the company. Using a universally anonymous encryption, a gateway server can be easily setup to anonymize any ciphertext sent outside the domain while internal clients can encrypt messages efficiently.

The following observations are obvious but worth emphasizing: (1) A universally anonymous scheme is also key-private in the sense of Bellare et al. [5]. What makes universally anonymity interesting and unique is that anyone can anonymize ciphertexts using just the public key of the recipient. (2) Key-private schemes can be more expensive than their non-private counterparts. For instance, RSA-OAEP can be made key-private as shown in [5] but the new anonymous variant is more expensive. (3) The concept of universal anonymity makes sense also for schemes that are already key-private. For instance, ElGamal is key-private only by assuming that all keys are generated in the same group and participants share the same public parameters. But in many scenarios this is not the case. In PGP, for instance, parameters for each user are selected in distinct groups. Evidently, ElGamal applied in different algebraic groups is not anonymous anymore as one can test whether a given ciphertext is in a group or not.

Our contributions are:

(1) We enhance Cocks IBE and make it universally anonymous, and thus key-private in the sense of Bellare et al. [5]. Our variant of Cocks IBE can be seen as the most efficient anonymous IBE whose security is based on the quadratic residuosity assumption. The efficiency of our scheme is comparable to that of Cocks IBE. In fact, it is substantially more efficient than the recent scheme of Boneh et al. [10] and the IBE that derives from the PEKS construction by Di Crescenzo et al. [15]. In addition, the ciphertext expansion of our scheme is comparable to that of Cocks IBE.

(2) We implemented our variant and measured its performance. We show that in practice the efficiency of Cocks IBE and the variant we propose in this paper compare favorably even with that

of the Boneh-Franklin scheme.
(3) Incidentally, our solutions and techniques can be used to simplify the PEKS construction in [15] and our Lemma 2.2 in Section 2.3 can be used to show that the new security assumption introduced in [15] is actually equivalent to the standard quadratic residuosity assumption, thus making the elegant Di Crescenzo-Saraswat PEKS scheme the first one whose security is based solely on such a standard assumption (which was left as an open problem in the area). However, we will not elaborate on this point any further in this paper.

Hybrid Encryption and CCA-security. It is well-known that in order to encrypt long messages, asymmetric encryption can be used in combination with symmetric encryption for improved efficiency. This simple and well-known paradigm has been formalized only recently by Cramer and Shoup [14, 13] and Shoup [26]. It is introduced as the KEM-DEM construction which consists of two parts: the key encapsulation mechanism (KEM), used to encrypt a symmetric key, and the data encapsulation mechanism (DEM) that is used to encrypt the plaintext via a symmetric cipher.

Cramer and Shoup [14] showed that a hybrid encryption scheme is CCA secure (i.e., secure against the adaptive chosen ciphertext attack) in the standard model if the KEM component is CCA secure and the DEM component is a CCA-secure one-time symmetric encryption. Later, Kurosawa and Desmedt [20] showed that the KEM component does not have to be CCA-secure as long as the CCA-secure one-time symmetric encryption satisfies an extra condition (which is satisfied by the DEM scheme proposed by Shoup [14, 26, 27]). The construction of CCA-secure one-time symmetric encryption is standard and it is usually accomplished by coupling a message authentication code (MAC) with a symmetric encryption. In particular, it can be shown that applying a one-time MAC on the output of a CPA-secure symmetric encryption results in a CCA-secure symmetric encryption scheme (see e.g. Cramer and Shoup [14]).

The focus of this paper is on variants of Cocks IBE which can be proven secure only in the random oracle model. Thus, it makes sense to consider KEM-DEM constructions that are CCA-secure in such a model. In this case, the most relevant work is the one from Fujisaki and Okamoto [17] that shows how to build CCA-secure hybrid encryption schemes and how to convert any CPA-secure asymmetric scheme into a CCA-secure one in the random oracle model. Even more relevant is the work of Bentahar et al. [6] that formalizes the concept of id-based KEM-DEM and provides a generic transformation from any IBE scheme to CCA-secure ID-based KEM in the random oracle model. In particular, it is possible to show (see, e.g., Bentahar et al. [6]) that if a KEM returns $(\text{Encrypt}_{\mathsf{IBE}}(K), F(K))$, where $\text{Encrypt}_{\mathsf{IBE}}(K)$ is a one-way encryption for an identity and F is a hash function modeled as a random oracle, then the combination of this KEM with a CCA-secure DEM results in a CCA-secure hybrid encryption. (Note that, unless the encryption is a permutation, an additional random oracle is needed within the encryption to compute the randomness from the plaintext, see Bentahar et al. [6] for details.)

Since our scheme $\mathsf{UAnonIBE}$ and its efficient variants are CPA-secure in the random oracle model, the resulting hybrid encryption that follows from the paradigm above is a CCA-secure encryption in the random oracle model. (Note that one-way encryption is implied by CPA-security.)

2 Preliminaries

In this section, we recall first the IBE scheme proposed by Cocks [12]. Then we show that Cocks IBE is not anonymous due to a test proposed by Galbraith, as reported in [8]. Finally, we show that

Galbraith's test is the "best test" possible against the anonymity of Cocks IBE. We assume that N is a large-enough RSA-type modulus. Hence, throughout the paper, we will omit to consider cases where randomly picked elements are in $\mathbb{Z}_N$ but not in $\mathbb{Z}_N^*$ or, analogously, have Jacobi symbol over N equal to 0 since these cases occur only with negligible probability[1]. Therefore, for consistency, we always assume to work in $\mathbb{Z}_N^*$ rather than in $\mathbb{Z}_N$ even though $\mathbb{Z}_N^*$ is not closed under modular addition.

We will denote with $\mathbb{Z}_N^*[+1]$ ($\mathbb{Z}_N^*[-1]$) the set of elements in $\mathbb{Z}_N^*$ with Jacobi symbol $+1$ (-1, resp.) and with $\mathbb{QR}(N)$ the set of quadratic residues (or squares) in $\mathbb{Z}_N^*$. The security of Cocks IBE (and our variants) relies on the standard quadratic residuosity assumption which simply states that the two distributions $DQR(n) = \{(c, N) : (N, p, q) \leftarrow Gen(1^n),\ c \leftarrow \mathbb{QR}(N)\}$ and $DQRN(n) = \{(c, N) : (N, p, q) \leftarrow Gen(1^n),\ c \leftarrow \mathbb{Z}_N^*[+1] \setminus \mathbb{QR}(N)\}$ are computationally indistinguishable, where n is a security parameter and $Gen(\cdot)$ generates a RSA-type n-bit Blum modulus and its two prime factors.

2.1 Cocks' IBE Scheme

Let $N = pq$ be a Blum integer, i.e., where p and q are primes each congruent to 3 modulo 4. In addition, we consider $H : \{0,1\}^* \rightarrow \mathbb{Z}_N^*[+1]$ a full-domain hash which will be modeled as a random oracle in the security analysis.

Master Key: The secret key of the trusted authority is (p, q) while its public key is $N = pq$.

Key Generation: Given the identity ID, the authority generates $a = H(ID)$ (thus the Jacobi symbol $\left(\frac{a}{N}\right)$ is $+1$). The secret key for the identity ID is a value r randomly chosen in $\mathbb{Z}_N^*$ such that $r^2 \equiv a \bmod N$ or $r^2 \equiv -a \bmod N$. This value r is stored and returned systematically.

Encryption: To encrypt a bit $b \in \{-1, +1\}$ for identity ID, choose uniformly at random two independent values $t, v \in \mathbb{Z}_N^*$, such that $\left(\frac{t}{N}\right) = \left(\frac{v}{N}\right) = b$, and compute:

$$(c, d) = \left(t + \frac{a}{t} \bmod N, v - \frac{a}{v} \bmod N\right)$$

Decryption: Given a ciphertext (c, d), first set $s = c$ if $r^2 \equiv a \bmod N$ or $s = d$ otherwise. Then, decrypt by computing:

$$\left(\frac{s + 2r}{N}\right) = b$$

Notice that $s + 2r \equiv w(1 + r/w)^2 \bmod N$, thus the Jacobi symbol of $s + 2r$ is equal to that of w, where w is either t or v.

2.2 Galbraith's Test (GT)

As mentioned in the paper by Boneh et al. [8], Galbraith showed that Cocks' scheme is not anonymous. Indeed, let $a \in \mathbb{Z}_N^*[+1]$ be a public key and consider the following set:

[1] The Jacobi symbol of $a \in \mathbb{Z}_N$ is denoted as $\left(\frac{a}{N}\right)$ and is either -1, 0, or $+1$. However, $\left(\frac{a}{N}\right) = 0$ if and only if $\gcd(a, N) \neq 1$, thus this case happens only with negligible probability since the value $\gcd(a, N)$ would be a non-trivial factor of N.

$$S_a[N] = \left\{ t + \frac{a}{t} \bmod N \mid t \in \mathbb{Z}_N^* \right\} \cap \mathbb{Z}_N^*$$

Given two random public keys $a, b \in \mathbb{Z}_N^*[+1]$, Galbraith's test (which we will denote with "$GT(\cdot)$") allows us to distinguish the uniform distribution on the set $S_a[N]$ from the uniform distribution on the set $S_b[N]$. Given $c \in \mathbb{Z}_N^*$, the test over the public key a is defined as the Jacobi symbol of $c^2 - 4a$ over N, that is:

$$GT(a, c, N) = \left(\frac{c^2 - 4a}{N} \right)$$

Notice that when c is sampled from $S_a[N]$, the test $GT(a, c, N)$ will return $+1$ with overwhelming probability given that $c^2 - 4a = (t - (a/t))^2$ is a square. However, if c is sampled from $S_b[N]$ the test is expected to return $+1$ with probability negligibly close to $1/2$ since, in this case, the distribution of the Jacobi symbol of the element $c^2 - 4a$ in $\mathbb{Z}_N^*$ follows the uniform distribution on $\{-1, +1\}$.

It is mentioned in [8] that since Cocks ciphertext is composed of several values sampled from either $S_a[N]$ (and $S_{-a}[N]$) or $S_b[N]$ (and $S_{-b}[N]$, respectively), then an adversary can repeatedly apply Galbraith's test to determine with overwhelming probability whether a given ciphertext is intended for a or b. However, one must first prove some meaningful results about the distribution of Jacobi symbols of elements of the form $c^2 - 4b$ in $\mathbb{Z}_N^*$, for *fixed* random elements $a, b \in \mathbb{Z}_N^*[+1]$ and for $c \in S_a[N]$. These results are reported in the next section.

2.3 Relevant Lemmata and Remarks

Damgård in [16] studied the distribution of Jacobi symbols of elements in $\mathbb{Z}_N^*$ in order to build pseudo-random number generators. In his paper, Damgård reports of a study performed in the 50s by Perron in which it is proven that for a prime p and for any a, the set $a + \mathbb{QR}(p)$ contains as many squares as non squares in $\mathbb{Z}_p^*$ when $p \equiv 1 \mod 4$, or the difference is just 1 when $p \equiv 3 \mod 4$. It is possible to generalize Perron's result to study the properties of the set $a + \mathbb{QR}(N)$ in $\mathbb{Z}_N^*$ but we also point out that the security of Cocks IBE implicitly depends on the following Lemma:

Lemma 2.1. *Let (a, N) be a pair such that $(N, p, q) \leftarrow Gen(1^n)$ and $a \leftarrow \mathbb{Z}_N^*[+1]$. The distribution $\left\{ \left(\frac{t^2+a}{N} \right) : t \leftarrow \mathbb{Z}_N^* \right\}$ is computationally indistinguishable from the uniform distribution on $\{-1, +1\}$ under the quadratic residuosity assumption.*

To prove the Lemma above it is enough to observe that if we compute the Jacobi symbol of a value $c \in S_a[N]$ we obtain:

$$\left(\frac{c}{N} \right) = \left(\frac{(t^2 + a)/t}{N} \right) = \left(\frac{t^2 + a}{N} \right) \left(\frac{t}{N} \right)$$

However, the Jacobi symbol of t over N is the plaintext in Cocks IBE and thus Lemma 2.1 must follow otherwise the CPA-security of Cocks IBE would not hold.

Remark. Let's pick c uniformly at random from $\mathbb{Z}_N^*$. If $GT(a, c, N) = -1$, we can clearly conclude that $c \notin S_a[N]$. However, if $GT(a, c, N) = +1$, what is the probability that $c \in S_a[N]$? The answer is $1/2$ since a t exists such that $c = t + a/t$ whenever $c^2 - 4a$ is a square and this happens only half of the times. Clearly $GT(a, c, N)$ is equal to 0 with negligible probability hence we do not consider this case. To summarize:

$$GT(a,c,N) = \begin{cases} +1 \Longrightarrow c \in S_a[N] \text{ with prob. } 1/2 \\ -1 \Longrightarrow c \notin S_a[N] \end{cases}$$

We will argue that there is no *better* test against anonymity over an encrypted bit. That is, we show that a test that returns +1 to imply that $c \in S_a[N]$ with probability $1/2 + \delta$ (for a non-negligible $\delta > 0$) cannot exist under the quadratic residuosity assumption. We first notice that $c \in S_a[N]$ if and only if $\Delta = c^2 - 4a$ is a square in $\mathbb{Z}_N$. Indeed, if $c = t + a/t$ then $\Delta = (t - a/t)^2$. If Δ is a square[2] then the quadratic equation $c = t + a/t$ has solutions for t in $\mathbb{Z}_N^*$. Thus $S_a[N]$ can alternatively be defined as the set of all $c \in \mathbb{Z}_N^*$ such that $c^2 - 4a$ is a square in $\mathbb{Z}_N$.

Intuitively, we can see Galbraith's test as an algorithm that checks whether the discriminant Δ has Jacobi symbol +1 or −1, and this is clearly *the best it can do* since the factors of the modulus N are unknown. (Recall that we do not consider cases where the Jacobi symbol is 0 since they occur with negligible probability.) Indeed, if $\pm x$ and $\pm y$ are the four distinct square roots modulo N of Δ, then $t^2 - ct + a$ is congruent to 0 modulo N whenever t is congruent modulo N to any of the following four distinct values:

$$\frac{c \pm x}{2} \text{ and } \frac{c \pm y}{2}$$

We denote with $GT_a^N[+1]$ the set $\{c \in \mathbb{Z}_N^* \mid GT(a,c,N) = +1\}$. Analogously, we define $GT_a^N[-1]$ as the set $\{c \in \mathbb{Z}_N^* \mid GT(a,c,N) = -1\}$. We prove the following Lemma:

Lemma 2.2. [VQR–Variable Quadratic Residuosity] *The two distributions* $D_0(n) = \{(a,c,N) : (N,p,q) \leftarrow Gen(1^n),\ a \leftarrow \mathbb{Z}_N^*[+1],\ c \leftarrow S_a[N]\}$ *and* $D_1(n) = \{(a,c,N) : (N,p,q) \leftarrow Gen(1^n),\ a \leftarrow \mathbb{Z}_N^*[+1],\ c \leftarrow GT_a^N[+1] \setminus S_a[N]\}$ *are computationally indistinguishable under the quadratic residuosity assumption.*

Proof. We assume there is a PPT adversary $\mathcal{A}$ that can distinguish between $D_0(n)$ and $D_1(n)$ with non-negligible advantage and we use it to solve a random instance of the quadratic residuosity problem. We make no assumptions on how $\mathcal{A}$ operates and we evaluate it via oracle access.

The simulator is given a random tuple (N, x) where $(N,p,q) \leftarrow Gen(1^n)$ and $x \leftarrow \mathbb{Z}_N^*[+1]$. The simulation proceeds as follows:

1. Find a random $r \in \mathbb{Z}_N^*$ such that $a = (r^2 - x)/4$ has Jacobi symbol +1 (see Lemma 2.1). $\mathcal{A}$ receives as input (a, r, N), where $a \in \mathbb{Z}_N^*[+1]$ and $r \in GT_a^N[+1]$. Notice that releasing the public key a effectively provides $\mathcal{A}$ with the ability to generate several values in $S_a[N]$ (and $S_{-a}[N]$);

2. If $\mathcal{A}$ responds that $r \in S_a[N]$ then output *"x is a square"* otherwise output *"x is not a square"*.

The value a is distributed properly. Indeed, if $x = r^2 - 4a$ then about half of the values that x assumes for all $r \in \mathbb{Z}_N^*$ have Jacobi symbol +1 (thanks to Lemma 2.1). Let $\mathsf{Pairs}_a = \{(r^2, x) \mid x \in \mathbb{Z}_N^*[+1], r \in \mathbb{Z}_N^*, a = (r^2 - x)/4\}$. For any random $a' \leftarrow \mathbb{Z}_N^*[+1]$, with $a' \neq a$, the sets Pairs_a and $\mathsf{Pairs}_{a'}$ have about the same cardinality and their intersection is empty. Thus, the pair (r^2, x) provided by the simulator represents a random pair from the union over a of all sets Pairs_a, and therefore the value a is uniformly distributed to a PPT adversary.

[2] If $\Delta \in \mathbb{Z}_N \setminus \mathbb{Z}_N^*$ then Δ has one or two square roots but this happens with negligible probability.

It has already been established that $r \in S_a[N]$ if and only if $r^2 - 4a$ is a square. But $r^2 - 4a = x$, therefore $\mathcal{A}$ cannot have non-negligible advantage under the quadratic residuosity assumption.

□

The next Lemma easily follows from Lemma 2.1 since $c^2 - 4a$ can be written as $c^2 + h$ for a fixed $h \in \mathbb{Z}_N^*[+1]$.

Lemma 2.3. *Let (a, N) be a pair such that $(N, p, q) \leftarrow Gen(1^n)$ and $a \leftarrow \mathbb{Z}_N^*[+1]$. The distribution $\{GT(a, c, N) : c \leftarrow \mathbb{Z}_N^*\}$ is computationally indistinguishable from the uniform distribution on $\{-1, +1\}$.*

3 Our Basic Construction and Its Efficient Variants

We extend Cocks' scheme to support anonymity. Unlike previous proposals, our scheme UAnonIBE has efficiency, storage, and bandwidth requirements similar to those of the original scheme by Cocks (which is not anonymous). Our scheme is also the first universally anonymous IBE, according to the definition in [19] (although we do not include the extra algorithms as in [19] to keep the presentation simple).

3.1 The Basic Scheme

Let $H : \{0,1\}^* \rightarrow \mathbb{Z}_N^*[+1]$ be a full-domain hash modeled as a random oracle. Let n and m be two security parameters (we assume w.l.o.g. that m depends on n). The algorithms which form UAnonIBE are defined as follows (all operations are performed modulo N):

Master Key: The public key of the trusted authority is the n-bit Blum integer $N = pq$, where p and q are $n/2$-bit primes each congruent to 3 modulo 4.

Key Generation: Given the identity ID, the authority generates $a = H(ID)$ (thus the Jacobi symbol $\left(\frac{a}{N}\right)$ is $+1$). The secret key for the identity ID is a value r randomly chosen in $\mathbb{Z}_N^*$ such that $r^2 \equiv a \bmod N$ or $r^2 \equiv -a \bmod N$. This value r is stored and returned systematically.

Encryption: To encrypt a bit $b \in \{-1, +1\}$ for identity ID, choose uniformly at random two values $t, v \in \mathbb{Z}_N^*$, such that $\left(\frac{t}{N}\right) = \left(\frac{v}{N}\right) = b$, and compute $(c, d) = \left(t + \frac{a}{t}, v - \frac{a}{v}\right)$.
Then, compute the *mask* to anonymize the ciphertext (c, d) as follows:

1. Pick two indices k_1 and k_2 independently from the geometric distribution[3] D with probability parameter 1/2;
2. Choose $T, V \leftarrow \mathbb{Z}_N^*$ independently and uniformly at random and set $Z_1 = c+T$ and $Z_2 = d+V$;
3. For $1 \leq i < k_1$, choose $T_i \leftarrow \mathbb{Z}_N^*$ independently and uniformly at random such that $GT(a, Z_1 - T_i, N) = -1$;
4. For $1 \leq i < k_2$, choose $V_i \leftarrow \mathbb{Z}_N^*$ independently and uniformly at random such that $GT(-a, Z_2 - V_i, N) = -1$;

[3] The geometric distribution is a discrete memoryless random distribution for $k = 1, 2, 3, \ldots$ having probability function $\Pr[k] = p(1-p)^{k-1}$ where $0 < p < 1$. Therefore, for $p = 1/2$ the probability that $k_1 = k$ is 2^{-k}. For more details see, e.g., [28].

5. Set $T_{k_1} = T$ and $V_{k_2} = V$;

6. For $k_1 < i \leq m$, choose $T_i \leftarrow \mathbb{Z}_N^*$ independently and uniformly at random;

7. For $k_2 < i \leq m$, choose $V_i \leftarrow \mathbb{Z}_N^*$ independently and uniformly at random;

Finally, output $(Z_1, T_1, \ldots, T_m) \in (\mathbb{Z}_N^*)^{m+1}$ and $(Z_2, V_1, \ldots, V_m) \in (\mathbb{Z}_N^*)^{m+1}$.[4]

Decryption: Given a ciphertext $(Z_1, T_1, \ldots, T_m)$ and $(Z_2, V_1, \ldots, V_m)$, first discard one of the two tuples based on whether a or $-a$ is a square. Let's assume we keep the tuple $(Z_1, T_1, \ldots, T_m)$ and we discard the other. In order to decrypt, find the smallest index $1 \leq i \leq m$ s.t. $GT(a, Z_1 - T_i, N) = +1$ and output:

$$\left(\frac{Z_1 - T_i + 2r}{N}\right) = b$$

We run the same procedure above if the second tuple is actually selected and the first is discarded. It is enough to replace a with $-a$, Z_1 with Z_2, and T_i with V_i.

3.2 Security Analysis

We need to show that our scheme, UAnonIBE, is ANON-IND-ID-CPA-secure [1, 10], that is, the ciphertext does not reveal any information about the plaintext and an adversary cannot determine the identity under which an encryption is computed, even thought the adversary selects the identities and the plaintext.

In [18], Halevi provides a sufficient condition for a CPA public-key encryption scheme to meet the notion of key-privacy, or anonymity, as defined by Bellare et al. in [5]. In [1], Abdalla et al. extend Halevi's condition to identity-based encryption. In addition, their notion is defined within the random oracle model and Halevi's statistical requirement is weakened to a computational one. Informally, it was observed that if an IBE scheme is already IND-ID-CPA-secure then the challenger does not have to encrypt the message chosen by the adversary but can encrypt a random message of the same length. The game where the challenger replies with an encryption on a random message is called ANON-RE-CPA. In [1], it was shown that if a scheme is IND-ID-CPA-secure and ANON-RE-CPA-secure then it is also ANON-IND-ID-CPA-secure.

ANON-RE-CPA game. We briefly describe the security game introduced by Abdalla et al. in [1]. MPK represents the set of public parameters of the trusted authority. The adversary $\mathcal{A}$ has access to a random oracle H and to an oracle KeyDer that given an identity ID returns the private key for ID according to the IBE scheme.

```
Experiment $\mathsf{Exp}^{\mathsf{anon-re-cpa}}_{\mathsf{IBE},\mathcal{A}}(n)$ :
    pick random oracle $H$;
    $(ID_0, ID_1, msg, state) \leftarrow \mathcal{A}^{\mathsf{KeyDer}(\cdot),H}(\mathtt{find}, MPK)$;
    $b \leftarrow \{0, 1\}$;
    $W \leftarrow \{0, 1\}^{|msg|}; c \leftarrow \mathsf{Enc}^H(MPK, ID_b, W)$;
```

[4] Note that if we build a sequence $c_1, \ldots, c_m$ by selecting a k from D and setting $c_i = -1$ for $1 \leq i < k$, $c_k = 1$, and random $c_i \leftarrow \{-1, 1\}$ for $k < i \leq m$, we have that $\Pr[c_i = 1] = 2^{-i} + \sum_{j=2}^{i} 2^{-j} = 1/2$.

$b' \leftarrow \mathcal{A}^{\mathsf{KeyDer}(\cdot),H}(\mathsf{guess}, c, state)$;

return 1 if $b' = b$, 0 otherwise.

The adversary cannot request the private key for ID_0 or ID_1 and the message msg must be in the message space associated with the scheme. A scheme is said to be ANON-RE-CPA-secure if for all polynomial time adversaries $\mathcal{A}$ there exists a negligible function $\mathsf{negl}(\cdot)$ such that:

$$\Pr[\mathsf{Exp}^{\mathsf{anon-re-cpa}}_{\mathsf{IBE},\mathcal{A}}(n) = 1] \leq 1/2 + \mathsf{negl}(n)$$

We define the function $\mathsf{Adv}^{\mathsf{anon-re-cpa}}_{\mathsf{UAnonIBE},\mathcal{A}}(\cdot)$ as the advantage of the adversary $\mathcal{A}$ against $\mathsf{UAnonibe}$, that is:

$$\mathsf{Adv}^{\mathsf{anon-re-cpa}}_{\mathsf{UAnonIBE},\mathcal{A}}(n) \overset{\mathsf{def}}{=} \Pr[\mathsf{Exp}^{\mathsf{anon-re-cpa}}_{\mathsf{UAnonIBE},\mathcal{A}}(n) = 1] - 1/2$$

Note that the algorithm Enc has access to the random oracle H even though H is not used to construct the challenge c. In our basic scheme, H is used only to generate public keys from identity strings.

Theorem 3.1. $\mathsf{UAnonIBE}$ *is* ANON-IND-ID-CPA-*secure in the random oracle model under the quadratic residuosity assumption.*

Proof. It must be clear that $\mathsf{UAnonIBE}$ is IND-ID-CPA-secure since Cocks IBE is IND-ID-CPA-secure in the random oracle model under the quadratic residuosity assumption and the mask is computed without knowing the plaintext or the secret key of the intended recipient. Thus, we only need to show that $\mathsf{UAnonIBE}$ is ANON-RE-CPA-secure.

In order to simplify the proof of theorem 3.1, we present it as a sequence of experiments which we call *games* following the standard literature. Let win_i denote the event that the adversary $\mathcal{A}$ wins in game i.

Game 0: This game is identical to the ANON-RE-CPA experiment defined above, where the adversary $\mathcal{A}$ chooses two identities ID_0 and ID_1 and receives two sequences $(Z_1, T_1, \ldots, T_m) \in (\mathbb{Z}_N^*)^{m+1}$ and $(Z_2, V_1, \ldots, V_m) \in (\mathbb{Z}_N^*)^{m+1}$. Hence, by definition, we have that:

$$\Pr[\mathsf{win}_0] - 1/2 = \mathsf{Adv}^{\mathsf{anon-re-cpa}}_{\mathsf{UAnonIBE},\mathcal{A}}(n)$$

In the rest of the proof, we consider w.l.o.g. only the first sequence of the challenge ciphertext, that is the sequence $(Z_1, T_1, \ldots, T_m)$. To prove that this is indeed w.l.o.g. notice that the second sequence $(Z_2, V_1, \ldots, V_m)$ is completely independent from the first one. In particular, the parameters and elements used to build the second sequence are generated independently from those used to build the first one, except that both sequences encrypt the same message. However, since the encryption is CPA-secure, this affects the advantage of the adversary only negligibly.

Game 1: In this game, we modify the way the challenger constructs the ciphertext. In particular, first set $a_0 = H(ID_0)$ and $a_1 = H(ID_1)$. Choose a random bit $b \leftarrow \{0, 1\}$. Then, a sequence $(v_1, \ldots, v_m) \in (\mathbb{Z}_N^*)^m$ is built as follows: (1) Select an index k according to the geometric distribution D with parameter 1/2. (2) Choose an element c_b uniformly at random from $S_{a_b}[N]$ and set $v_k = c_b$. (3) Choose elements $v_1, \ldots, v_{k-1}$ independently and uniformly at random from $GT^N_{a_b}[-1]$ and elements $v_{k+1}, \ldots, v_m$ from $\mathbb{Z}_N^*$.

The challenger chooses Z_1 uniformly at random from $\mathbb{Z}_N^*$ and sets $T_i = Z_1 - v_i \bmod N$, for $1 \leq i \leq m$. The challenge ciphertext is $(Z_1, T_1, \ldots, T_m)$.

The adversary $\mathcal{A}$ cannot distinguish Game 1 from Game 0 since the index k is chosen from the same distribution and the values Z_1 and T_i, for $1 \leq i \leq m$, in Game 1 are distributed as in Game 0 with overwhelming probability. Therefore, there exists a negligible function $\mathsf{negl_0}(\cdot)$ such that:

$$|\Pr[\mathsf{win}_0] - \Pr[\mathsf{win}_1]| \leq \mathsf{negl_0}(n)$$

Game 2: In this game, we make only a minor change to Game 1. In particular, the challenger chooses c_b uniformly at random from $GT_{a_b}^N[+1]$ (and not from $S_{a_b}[N]$ as in Game 1). The rest of the construction is identical to the one in Game 1.

Thanks to Lemma 2.2 (VQR), the adversary $\mathcal{A}$ cannot distinguish Game 2 from Game 1. Namely, there exists a negligible function $\mathsf{negl_1}(\cdot)$ such that:

$$|\Pr[\mathsf{win}_1] - \Pr[\mathsf{win}_2]| \leq \mathsf{negl_1}(n)$$

Game 3: This game proceeds exactly as Game 2 except that the sequence $(v_1, \ldots, v_m)$ is now chosen uniformly at random from $(\mathbb{Z}_N^*)^m$ (i.e., steps (1) through (3) described in Game 1 and in Game 2 are not performed and each v_i is independently and uniformly distributed in $\mathbb{Z}_N^*$).

It is easy to show that Game 3 is indistinguishable from Game 2. Indeed, the only difference between the two games is in the way the sequence $(v_1, \ldots, v_m) \in (\mathbb{Z}_N^*)^m$ is built. In Game 2, an index k is picked according to the geometric distribution with parameter 1/2, v_k is uniformly distributed in $GT_{a_b}^N[+1]$, and the elements before and after v_k in the sequence are independently and uniformly distributed respectively in $GT_{a_b}^N[-1]$ and $\mathbb{Z}_N^*$. In Game 3, each v_i is chosen independently and uniformly at random from $\mathbb{Z}_N^*$. The adversary $\mathcal{A}$ will not detect any difference between these two distributions thanks to Lemma 2.3. Namely, we know that $\Pr[GT(a_b, v_i, N) = +1]$ is negligibly close to 1/2 whether v_i is from the distribution in Game 2 or from the one in Game 3. Therefore, there exists a negligible function $\mathsf{negl_2}(\cdot)$ such that:

$$|\Pr[\mathsf{win}_2] - \Pr[\mathsf{win}_3]| \leq \mathsf{negl_2}(n)$$

Combining the probabilities in all games, we have that there exists a negligible function $\mathsf{negl}(\cdot)$ such that:

$$\mathsf{Adv}^{\mathsf{anon-re-cpa}}_{\mathsf{UAnonIBE},\mathcal{A}}(n) = \Pr[\mathsf{win}_0] - 1/2 \leq \Pr[\mathsf{win}_3] - 1/2 + \mathsf{negl}(n)$$

To conclude the proof we must show that $\Pr[\mathsf{win}_3]$ is exactly 1/2. This follows simply because the challenge ciphertext $(Z_1, T_1, \ldots, T_m)$ in Game 3 is independent of the bit b. Indeed, $Z_1 \leftarrow \mathbb{Z}_N^*$ and $T_i = Z_1 - v_i \in \mathbb{Z}_N^*$ where the v_i's are independently and uniformly distributed in $\mathbb{Z}_N^*$ (recall that, as in the rest of the paper, we explicitly use $\mathbb{Z}_N^*$ rather than $\mathbb{Z}_N$ since elements uniformly distributed in $\mathbb{Z}_N$ are in $\mathbb{Z}_N^*$ with overwhelming probability.)

□

3.3 A First Efficient Variant: Reducing Ciphertext Expansion

The obvious drawback of the basic scheme is its ciphertext expansion. Indeed, for each bit of the plaintext $2 \cdot (m+1)$ values in $\mathbb{Z}_N^*$ must be sent while in Cocks IBE each bit of the plaintext requires two values in $\mathbb{Z}_N^*$. Therefore, we need a total of $2 \cdot (m+1) \cdot n$ bits for a single bit in the plaintext, where n and m are the security parameters (e.g., $n = 1024$ and $m = 128$). This issue, however, is easy to fix. Intuitively, since our scheme requires the random oracle model for its security, we could use another random oracle that expands a short seed into a value selected uniformly and independently in $\mathbb{Z}_N^*$. Specifically, a function $G : \{0,1\}^* \to \mathbb{Z}_N^*$ is used, which we model as a random oracle, that maps a e-bit string α to a random value in $\mathbb{Z}_N^*$. The parameter e must be large enough, e.g., $e = 160$.

It is tempting to use the oracle G and a single short seed α plus a counter to generate all values $T_1, \ldots, T_m$ and $V_1, \ldots, V_m$. This first solution would provide minimal ciphertext expansion, since only the seed α must be sent, however it may turn out to be computationally expensive. To see this, consider that an α must be found such that $GT(a, Z_1 - T_i, N) = -1$ for $1 \leq i < k_1$. Now, if k_1 happens to be large, say $k_1 = 20$, then clearly finding a suitable α could be computationally intensive. Nevertheless, we prove that this scheme is secure as long as the basic UAnonIBE scheme is secure. More importantly, we emphasize that the proof of security of all other schemes proposed after this first one can easily be derived from the proof of the following theorem.

Theorem 3.2. *The first efficient variant of* UAnonIBE *is* ANON-IND-ID-CPA-*secure in the random oracle model under the quadratic residuosity assumption.*

Proof. We let the simulator $\mathcal{S}$ play the role of a man-in-the-middle attacker between two ANON-RE-CPA games: the first game is against the basic UAnonIBE and the second game is against an adversary $\mathcal{A}$ that has non-negligible advantage in breaking the first variant of UAnonIBE. We show that $\mathcal{S}$ can use $\mathcal{A}$ to win in the first ANON-RE-CPA game, thus violating the quadratic residuosity assumption. The simulation is straightforward: $\mathcal{S}$ forwards the H-queries and KeyDer-queries to the respective oracles. When $\mathcal{A}$ challenges for identities ID_0 and ID_1, $\mathcal{S}$ challenges on the same identities in the first ANON-RE-CPA game. Then $\mathcal{S}$ receives the ciphertext $(Z_1, T_1, \ldots, T_m)$, $(Z_2, V_1, \ldots, V_m)$. $\mathcal{S}$ sends $(Z_1, \alpha), (Z_2, \beta)$ to $\mathcal{A}$, where α and β are chosen uniformly at random in $\{0,1\}^e$. At this point, the simulator responds to the G-queries as follows:

$$G(\alpha \mid\mid i) = T_i \text{ and } G(\beta \mid\mid i) = V_i \text{ , for } 1 \leq i \leq m,$$

and with random values in $\mathbb{Z}_N^*$ in any other cases. The adversary $\mathcal{A}$ eventually returns its guess which $\mathcal{S}$ uses in the first game in order to win with non-negligible advantage. □

The obvious next-best solution is to use a single seed per value. Thus, rather than sending the ciphertext as per our basic scheme, that is $(Z_1, T_1, \ldots, T_m)$ and $(Z_2, V_1, \ldots, V_m)$, the following values could be sent:

$$(Z_1, \alpha_1, \ldots, \alpha_m) \text{ and } (Z_2, \beta_1, \ldots, \beta_m),$$

where α_i, β_i are chosen uniformly at random in $\{0,1\}^e$ until the conditions in steps 3. and 4. of the encryption algorithm of the basic scheme are satisfied. The recipient would then derive the intended ciphertext by computing $T_i = G(\alpha_i)$ and $V_i = G(\beta_i)$, for $1 \leq i \leq m$. If we set e to be large enough, say $e = 160$, then clearly the security of this variant derives from the one of the basic scheme in the random oracle model and a single bit of the plaintext would require $2 \cdot (m \cdot e + n)$ bits rather than $2 \cdot (m \cdot n + n)$, where $e < n$. Hence, for $n = 1024$, $m = 128$ and $e = 160$, we need to send $2 \cdot (160 \cdot 128 + 1024)$ bits while Cocks' scheme requires only $2 \cdot (1024)$ bits.

On a closer look, however, it is easy to see that since G is a random oracle we just need to ensure that its inputs are repeated only with negligible probability. Let $X = x^{(1)}x^{(2)}\dots x^{(t)}$ be the plaintext of t bits. For each plaintext X, the sender selects a random message identifier $MID_X \in \{0,1\}^{e_1}$ which is sent along with the ciphertext. For bit $x^{(j)}$, the sender computes:

$$(Z_1^{(j)}, \alpha_1^{(j)}, \dots, \alpha_m^{(j)}) \text{ and } (Z_2^{(j)}, \beta_1^{(j)}, \dots, \beta_m^{(j)}),$$

where the coefficients $\alpha_i^{(j)}$, $\beta_i^{(j)}$ are chosen uniformly at random in $\{0,1\}^e$ until the conditions in steps 3. and 4. of the encryption algorithm of the basic scheme are satisfied (thus notice that e can be small but still big enough to be able to find those values $T_i^{(j)}$ and $V_i^{(j)}$ that satisfy such conditions). The recipient will derive the intended ciphertext by computing:

$$T_i^{(j)} = G(MID_X \mid\mid 0 \mid\mid \alpha_i^{(j)} \mid\mid i \mid\mid j) \text{ or } V_i^{(j)} = G(MID_X \mid\mid 1 \mid\mid \beta_i^{(j)} \mid\mid i \mid\mid j),$$

where $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, t\}$. As an example, we can set $m = 128$, $e_1 = 160$, and $e = 8$. In this case the ciphertext expansion per single bit of the plaintext is only $2 \cdot (1024 + 1024)$ bits which is twice the amount required by Cocks IBE for $n = 1024$. (In addition, extra 160 bits are needed for MID_X but these bits are transmitted only once per message.)

3.4 A Second Efficient Variant: Trade-off Between Ciphertext Expansion and Performance

We propose a second variant of UAnonIBE which provides an optimal trade-off between efficiency and ciphertext expansion. Our performance tests show that this variant is in practice as efficient as any of the previous variants and at the same time it provides the smallest ciphertext expansion (thus we recommend this version for practical systems).

We fix a new global parameter ℓ which is a small positive integer. Let $X = x^{(1)}x^{(2)}\dots x^{(t)}$ be the plaintext of t bits. For each plaintext X, the sender selects a random identifier $MID_X \in \{0,1\}^{e_1}$ which is sent along with the ciphertext. For bit $x^{(j)}$, the sender computes:

$$(Z_1^{(j)}, \alpha_1^{(j)}, \dots, \alpha_\ell^{(j)}) \text{ and } (Z_2^{(j)}, \beta_1^{(j)}, \dots, \beta_\ell^{(j)})$$

where $\alpha_i^{(j)}$, $\beta_i^{(j)}$ are in $\{0,1\}^e$, when $i < \ell$, and $\alpha_\ell^{(j)}$, $\beta_\ell^{(j)}$ are in $\{0,1\}^{e'}$, for some $e' > e$. The intended ciphertext is derived by the recipient by computing:

$$T_i^{(j)} = G(MID_X \mid\mid 0 \mid\mid \alpha_i^{(j)} \mid\mid i \mid\mid j) \text{ or } V_i^{(j)} = G(MID_X \mid\mid 1 \mid\mid \beta_i^{(j)} \mid\mid i \mid\mid j)$$

for $i < l$, and

$$T_i^{(j)} = G(MID_X \mid\mid 0 \mid\mid \alpha_\ell^{(j)} \mid\mid i \mid\mid j) \text{ or } V_i^{(j)} = G(MID_X \mid\mid 1 \mid\mid \beta_\ell^{(j)} \mid\mid i \mid\mid j)$$

for $i \geq \ell$. Note that in this variant of our basic scheme an arbitrary number of $T_i^{(j)}$ and $V_i^{(j)}$ can be generated (i.e., there is no fixed global parameter m).

Given the distribution of k_1, k_2, for a large enough ℓ, we expect $k_1 \leq \ell$ or $k_2 \leq \ell$ with high probability. When $k_1 \leq \ell$ or $k_2 \leq \ell$ the scheme is as efficient as the first variant. When $k_1 > \ell$ (or $k_2 > \ell$) the computational cost of finding a value for $\alpha_\ell^{(j)}$ (or $\beta_\ell^{(j)}$) is exponential in $k_1 - \ell$ ($k_2 - \ell$, respectively).

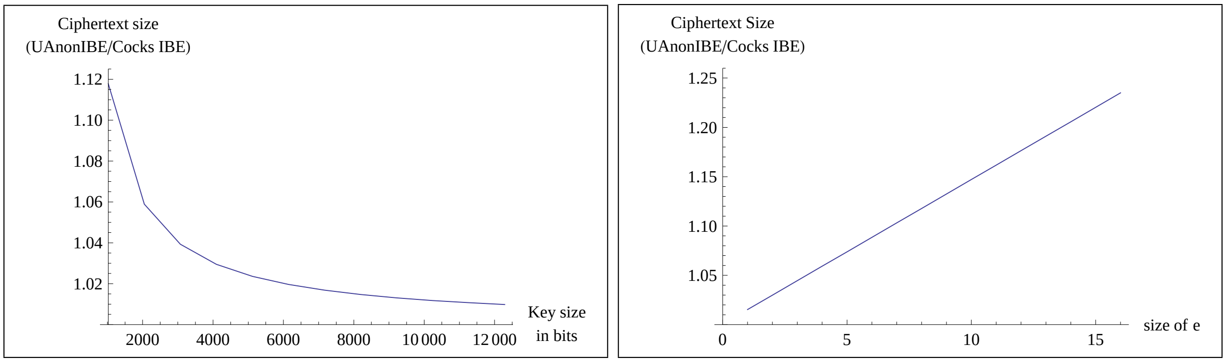

Figure 1: The two graphs show UAnonIBE's ciphertext size relative to Cocks' scheme. The first shows how the relative bandwidth overhead introduced by our solution decreases with the size of the master parameter, while the second shows how the size of the ciphertext increases, varying the size of e and fixing $e' = 10 \cdot e$, compared to Cocks' ciphertext.

As an example, we set the global parameter $\ell = 6$ and then $e_1 = 160$, $e = 8$, $e' = 80$, and $n = 1024$. The ciphertext expansion of this variant of UAnonIBE is $2 \cdot ((\ell - 1) \cdot e + e' + n)$, therefore, the ciphertext size for a single bit of the plaintext is now only $2 \cdot (120 + 1024)$ bits which is very close to the number of bits $(2 \cdot (1024))$ required by Cocks IBE (which is not anonymous). Note that for each message, the sender also transmits the random message identifier MID_X.

4 Optimizations and Implementation

An important aspect that should be considered in order to implement UAnonIBE efficiently is the value of the parameters ℓ, e (the size of $\alpha_1^{(j)}, \ldots, \alpha_{\ell-1}^{(j)}$) and e' (the size of $\alpha_\ell^{(j)}$). These values affect both the ciphertext expansion and the encryption time significantly, therefore they must be selected carefully. Choosing e or e' to be too small can reduce the probability of encrypting to an unacceptable level. Choosing ℓ to be too small can make the encryption process very slow. If we set $e = 8$ and $e' = 80$, we can find a suitable value for each $\alpha_i^{(j)}$, and therefore encrypt, with a probability of at least $1 - 2^{-80}$. We found that the value $\ell = 6$ is the best compromise between encryption time and ciphertext expansion. If we set $e = 8$, $e' = 80$ and $\ell = 6$, the ciphertext expansion for a 128-bit message is 3840 bytes more than a Cocks encryption for both $+a$ and $-a$: for a 1024-bit modulus N the encrypted message size is about 36KB instead of 32KB with Cocks IBE.

Size of e	2	4	6	8	10	Cocks IBE
Ciphertext size in bytes	33748	34708	35668	36628	37588	32768

We implemented the second efficient variant of UAnonIBE and compared it with the original Cocks IBE [12] and the scheme proposed by Boneh and Franklin based on pairings [9]. We have two goals in mind. The first is to show that Cocks IBE and our schemes are practical when used as hybrid encryption algorithms (following the KEM-DEM paradigm), even when compared with the Boneh-Franklin IBE, with the clear advantage compared to other IBE schemes of relying on a well-established assumption. Our second goal is to show that the efficiency of our scheme is comparable to that of the original scheme by Cocks.

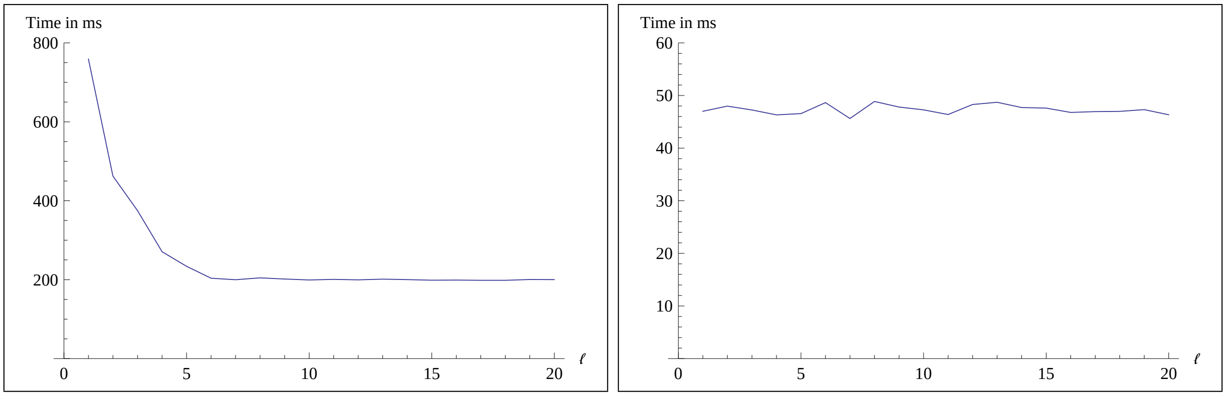

Figure 2: The two graphs show the average time required respectively to anonymize and de-anonymize a 128-bit message encrypted with Cocks IBE varying ℓ.

For our performance analysis, we set the size of the values $\alpha_i^{(j)}$ and $\beta_i^{(j)}$ with $1 \leq i < \ell$ to 8 bits and the size of $\alpha_\ell^{(j)}$ and $\beta_\ell^{(j)}$ to 80 bits. However, our tests showed that the size of those parameters have no measurable impact on the performance of the scheme. In order to calculate the optimal value for ℓ, we measured the time required to anonymize a key of 128 bit (for a total of 256 encrypted bits, considering both cases $+a$ and $-a$). Figure 2 summarizes our results. The value $\ell = 6$ seems to be optimal, since further increasing ℓ does not noticeably affect the time required to anonymize or decrypt a message.

Experimental Setup. We employed the MIRACL software package, developed by Shamus Software [25], to run our tests. MIRACL is a comprehensive library often used to implement cryptographic systems based on pairings. We used the optimized implementation of the Boneh-Franklin IBE provided by the library and we implemented Cocks IBE and our scheme with an RSA modulus of 1024 bits. The implementation of the Boneh-Franklin IBE uses a 512-bit prime p, Tate pairing and a small 160-bit subgroup q. The curve used is $y^2 = x^3 + x$ instead of $y^2 = x^3 + 1$ because it allows for a faster implementation. Those two settings should provide the same level of security according to NIST [21]. The tests were run on a machine that consisted of an Intel Pentium 4 2.8GHz with 512MB RAM. The system was running Linux kernel 2.6.20 with the compiler GCC 4.1.2. We implemented the cryptographic primitives using version 5.2.3 of the MIRACL library. Every source file was compiled with optimization '-02', as suggested in the MIRACL documentation. The table below shows average times over 1000 runs of the cryptographic operations specified in the first row. The symmetric key encrypted in our tests is of 128 bits.

	Extract	Encrypt	Decrypt	Anonymous	Universally Anonymous
Boneh-Franklin	9.1 ms	91.6 ms	85.4 ms	YES	NO
Cocks IBE	14.2 ms	115.3 ms	35.0 ms	NO	NO
Our Scheme	14.2 ms	319.4 ms	78.1 ms	YES	YES

In the table we also indicate whether a scheme is anonymous or not. Cocks IBE is not anonymous while Boneh-Franklin IBE is anonymous but not universally anonymous.

Remark. The main aspect characterizing universal anonymity is the ability to separate the role of the sender of encrypted messages from the role of the anonymzer. It is trivial to achieve this "separation of roles" with the Boneh-Franklin scheme or with any other anonymous scheme: it is enough to append to the ciphertext a claim about the recipient public key together with a "proof of consistency", i.e., a proof that the claim about the recipient key is valid. In order to restore anonymity, this extra information that is appended to the original ciphertext can simply be removed and this can be done without knowing any secret information. Depending on the underlying IBE scheme, this "proof of consistency" may or may not be efficient. In addition, this approach inevitably makes the non-anonymous scheme more expensive than the anonymous version while ideally the opposite should hold. One could try to find a variant of Boneh-Franklin IBE that is not anonymous but more efficient than the original scheme and then make it universally anonymous by applying the techniques in [19]. Even assuming that this is possible, the new scheme would be different and more expensive than the original one and still depending on pairing-based assumptions.

5 Conclusions

We proposed UAnonIBE: the first IBE providing universal anonymity (thus key-privacy) and secure under the standard quadratic residuosity assumption. The efficiency and ciphertext expansion of our scheme are comparable to those of Cocks IBE. We showed that Cocks IBE and our anonymous variant are suitable in practice whenever hybrid encryption (KEM-DEM paradigm) is employed. We believe our schemes are valid alternatives to decidedly more expensive schemes introduced in [10] (which, in addition, are anonymous but not universally anonymous).

Acknowledgments. We are grateful to Giovanni Di Crescenzo, Marc Joye, Ivano Repetto, and the anonymous reviewers for their insightful comments.

References

[1] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. MaloneLee, G. Neven, P. Paillier, and H. Shi. Searchable Encryption Revisited: Consistency Properties, Relation to Anonymous IBE, and Extensions. In *Advances in Cryptology, CRYPTO '05, volume 3621 of Lecture Notes in Computer Science*, pages 205–222. Springer-Verlag, 2005.

[2] G. Ateniese, J. Camenisch, and B. de Medeiros. Untraceable RFID Tags via Insubvertible Encryption. In *CCS '05: Proceedings of the 12th ACM conference on Computer and communications security*, pages 92–101, New York, NY, USA, 2005. ACM.

[3] G. Ateniese and P. Gasti. Universally Anonymous IBE Based on the Quadratic Residuosity Assumption. In M. Fischlin, editor, *CT-RSA*, volume 5473 of *Lecture Notes in Computer Science*, pages 32–47. Springer, 2009.

[4] L. Ballard, M. Green, B. de Medeiros, and F. Monrose. Correlation-Resistant Storage via Key-wordSearchable Encryption. In *Cryptology ePrint Archive, Report 2005/417*, 2005. Available at `http://eprint.iacr.org/2005/417`.

[5] M. Bellare, A. Boldyreva, A. Desai, and D. Pointcheval. Key-Privacy in Public-Key Encryption. In *Advances in Cryptology, ASIACRYPT '01, volume 2248 of Lecture Notes in Computer Science*, pages 566–582, London, UK, 2001. Springer-Verlag.

[6] K. Bentahar, P. Farshim, J. Malone-Lee, and N. Smart. Generic Constructions of Identity-Based and Certificateless KEMs. *Journal of Cryptology*, 21(2):178–199, April 2008.

[7] D. Boneh, X. Boyen, and H. Shacham. Short Group Signatures. In *Advances in Cryptology, CRYPTO '04, volume 3152 of Lecture Notes in Computer Science*, pages 41–55. Springer-Verlag, 2004.

[8] D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano. Public Key Encryption with Keyword Search. In *Advances in Cryptology, EUROCRYPT '04, volume 3027 of Lecture Notes in Computer Science*, pages 506–522, Interlaken, Switzerland, 2004.

[9] D. Boneh and M. Franklin. Identity-Based Encryption from the Weil Pairing. *SIAM Journal on Computing, volume 32-3*, pages 586–615, Philadelphia, PA, USA, 2003.

[10] D. Boneh, C. Gentry, and M. Hamburg. Space-Efficient Identity Based Encryption Without Pairings. In *FOCS '07: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 647–657, Washington, DC, USA, 2007. IEEE Computer Society.

[11] X. Boyen and B. Waters. Anonymous Hierarchical Identity-Based Encryption (Without Random Oracles). In *Advances in Cryptology, CRYPTO '06, volume 4117 of Lecture Notes in Computer Science*, pages 290–307. Springer-Verlag, 2006.

[12] C. Cocks. An Identity Based Encryption Scheme Based on Quadratic Residues. In *Proceedings of the 8th IMA International Conference on Cryptography and Coding*, pages 360–363, London, UK, 2001. Springer-Verlag.

[13] R. Cramer and V. Shoup. A Practical Public Key Cryptosystem Provably Secure against Adaptive Chosen Ciphertext Attack. In *Advances in Cryptology, CRYPTO '98, volume 1462 of Lecture Notes in Computer Science*, pages 13–25. Springer-Verlag, 1998.

[14] R. Cramer and V. Shoup. Design and Analysis of Practical Public-Key Encryption Schemes Secure against Adaptive Chosen Ciphertext Attack. *SIAM Journal on Computing, volume 33-1*, pages 167–226, Philadelphia, PA, USA, 2004.

[15] G. Di Crescenzo and V. Saraswat. Public Key Encryption with Searchable Keywords Based on Jacobi Symbols. In *Progress in Cryptology, INDOCRYPT '07, volume 3797 of Lecture Notes in Computer Science*, pages 282–296, Chennai, India, December 9-13, 2007.

[16] I. Damgård. On the Randomness of Legendre and Jacobi Sequences. In *Advances in Cryptology, CRYPTO '88, volume 403 of Lecture Notes in Computer Science*, pages 163–172, London, UK, 1990. Springer-Verlag.

[17] E. Fujisaki and T. Okamoto. Secure Integration of Asymmetric and Symmetric Encryption Schemes. In *Advances in Cryptology, CRYPTO '99, volume 1666 of Lecture Notes in Computer Science*, pages 537–554, London, UK, 1999. Springer-Verlag.

[18] S. Halevi. A Sufficient Condition for Key-Privacy. In *Cryptology ePrint Archive, Report 2005/05*, 2005. Available at `http://eprint.iacr.org/2005/005`.

[19] R. Hayashi and K. Tanaka. Universally Anonymizable Public-Key Encryption. In *Advences in Cryptologu, ASIACRYPT '05, volume 3788 of Lecture Notes in Computer Science*, pages 293–312, London, UK, 2005.

[20] K. Kurosawa and Y. Desmedt. A New Paradigm of Hybrid Encryption Scheme. In *Advances in Cryptology, CRYPTO '04, volume 3152 of Lecture Notes in Computer Science*, pages 426–442, London, UK, 2004.

[21] NIST. The Case for Elliptic Curve Cryptography. Available at `http://www.nsa.gov/ia/industry/crypto_elliptic_curve.cfm`.

[22] R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems Based on Pairing. In *Symposium on Cryptography and Information Security (SCIS 2000), Okinawa, Japan*, 2000.

[23] M. Scott. Authenticated ID-based Key Exchange and Remote Log-in With Insecure Token and PIN Number. In *Cryptology ePrint Archive, Report 2002/164*, 2002. Available at `http://eprint.iacr.org/2002/164`.

[24] A. Shamir. Identity-Based Cryptosystems and Signature Schemes. In *Advances in Cryptology, CRYPTO '84, volume 196 of Lecture Notes in Computer Science*, pages 47–53, New York, NY, USA, 1985. Springer-Verlag.

[25] Shamus Software. The MIRACL library. Available at `http://www.shamus.ie`.

[26] V. Shoup. A Proposal for an ISO Standard for Public Key Encryption (Version 2.1), manuscript, December 20, 2001. Available at `http://www.shoup.net/papers/iso-2_1.pdf`.

[27] V. Shoup. Using Hash Functions as a Hedge against Chosen Ciphertext Attack. In *Advances in Cryptology, EUROCRYPT '00, volume 1807 of Lecture Notes in Computer Science*, pages 275–288, 2000.

[28] M. R. Spiegel. *Theory and Problems of Probability and Statistics*. McGraw-Hill, 1992.