

Searchable Encryption Revisited: Consistency Properties, Relation to Anonymous IBE, and Extensions*

Michel Abdalla

Departement d'Informatique, École Normale Supérieure & CNRS, 45 Rue d'Ulm, 75230 Paris Cedex 05,
France
Michel.Abdalla@ens.fr

Mihir Bellare

Department of Computer Science & Engineering, University of California San Diego, 9500 Gilman Drive,
La Jolla, CA 92093, USA
mihir@cs.ucsd.edu

Dario Catalano

Dipartimento di Matematica e Informatica, Università di Catania, Viale Andrea Doria 6, 95125 Catania, Italy
catalano@dmi.unict.it

Eike Kiltz

CWI Amsterdam, P.O. Box 94079, 1090 GB, Amsterdam, The Netherlands
kiltz@cw.nl

Tadayoshi Kohno

Department of Computer Science and Engineering, University of Washington, P.O. Box 352350 Seattle,
WA 98195-2350, USA
yoshi@cs.washington.edu

Tanja Lange

Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513,
5600 MB Eindhoven, The Netherlands
tanja@hyperelliptic.org

John Malone-Lee

EMB Consultancy LLP, Bristol, UK
malone@compsci.bristol.ac.uk

Gregory Neven

Department of Electrical Engineering, Katholieke Universiteit Leuven, Kasteelpark Arenberg 10,
3001 Heverlee-Leuven, Belgium
Gregory.Neven@esat.kuleuven.be

Pascal Paillier

Cryptography & Innovation, Security Labs, Gemalto, 6, Rue de la Verrerie, 92197 Meudon Cedex, France
Pascal.Paillier@gemalto.com

Haixia Shi

NVIDIA Corporation, 2701 San Tomas Expressway, Santa Clara, CA 95050, USA
hshi@nvidia.com

Communicated by Arjen K. Lenstra

Received 1 September 2005 and revised 8 February 2007

Online publication 14 September 2007

Abstract. We identify and fill some gaps with regard to consistency (the extent to which false positives are produced) for public-key encryption with keyword search (PEKS). We define computational and statistical relaxations of the existing notion of perfect consistency, show that the scheme of Boneh et al. (*Advances in Cryptology—EUROCRYPT 2004*, ed. by C. Cachin, J. Camenisch, pp. 506–522, 2004) is computationally consistent, and provide a new scheme that is statistically consistent. We also provide a transform of an anonymous identity-based encryption (IBE) scheme to a secure PEKS scheme that, unlike the previous one, guarantees consistency. Finally, we suggest three extensions of the basic notions considered here, namely anonymous hierarchical identity-based encryption, public-key encryption with temporary keyword search, and identity-based encryption with keyword search.

Key words. Foundations, Random-oracle model, Anonymity, Identity-based encryption, Searchable encryption

1. Introduction

There has recently been interest in various forms of “searchable encryption” [8,16,18,24,26]. In this paper, we further explore one of the variants of this goal, namely public-key encryption with keyword search (PEKS) as introduced by Boneh, Di Crescenzo, Ostrovsky and Persiano [8].

The killer application envisaged by Boneh et al. is that of intelligent email routing. We consider emails as consisting of some header information, a body, and a list of keywords. Imagine Alice uses different electronic devices to read her email, including a pager, a PDA, and a desktop computer. Alice may prefer emails to be routed to her devices depending on the associated keywords. For example, she may like to receive emails with the keyword “urgent” on her pager, emails with the keyword “agenda” on her PDA, and all other emails on her desktop computer.

Existing mail server software could be updated to provide this type of service for plain, unencrypted email. When Bob sends an email to Alice encrypted under her public key, however, routing becomes much harder. One option would be for Bob to leave the list of keywords unencrypted; if Bob is a colleague of Alice however, he may not like the gateway to know that he is exchanging emails with her with the keyword “personal”. Alice is probably not willing to hand her decryption key to the gateway either. Rather, she would like to give the gateway some piece of trapdoor information that allows it to test whether the keyword “urgent” is among those in the list, without revealing any other information about the email to the gateway. This is exactly the type of functionality

* An extended abstract of this paper appears in *Advances in Cryptology—CRYPTO 2005*, ed. by V. Shoup, Santa Barbara, California, August 14–18, 2005, Lecture Notes in Computer Science, vol. 3621 (Springer, Berlin, 2005), pp. 205–222. This is the full version.

provided by a PEKS scheme. Bob can then use a standard public-key encryption scheme to encrypt the body of the email, and a PEKS scheme to separately encrypt each of the keywords.

The routing configuration of the email gateway need not be static. Alternatively, Alice could send the trapdoors for the keywords that she wants to receive at the time of login. This could be useful for checking email over a low-bandwidth connection: when Alice is at a conference, for example, she may want to download to her laptop only those emails tagged with keyword “urgent”.

As another application, Waters et al. [26] show how PEKS schemes can be used to let an untrusted logging device maintain an encrypted audit log of privacy-sensitive data (e.g. user actions on a computer system) that is efficiently searchable by authorized auditors only. The entries in the audit log are encrypted under the public key of a PEKS scheme, of which the corresponding secret key is unknown to the logging device. If the device is ever confiscated, or if the logbook leaks, privacy of users and their actions is maintained. The secret key is known only to a trusted audit escrow agent, who provides (less trusted) authorized investigators with trapdoors for the keywords they want to search for.

In this paper, we investigate some consistency-related issues and results of PEKS schemes, then consider the connection to anonymous identity-based encryption (IBE), and finally discuss some new extensions.

1.1. Consistency in PEKS

Any cryptographic primitive must meet two conditions. One is of course a security condition. The other, which we will here call a *consistency* condition, ensures that the primitive fulfills its function. For example, for public-key encryption, the security condition is privacy. (This could be formalized in many ways, e.g. IND-CPA or IND-CCA.) The consistency condition is that decryption reverses encryption, meaning that if M is encrypted under public key pk to result in ciphertext C , then decrypting C under the secret key corresponding to pk results in M being returned.

PEKS In a PEKS scheme, Alice can provide a *gateway* with a trapdoor t_w (computed as a function of her secret key) for any keyword w of her choice. A sender encrypts a keyword w' under Alice's public key pk to obtain a ciphertext C that is sent to the gateway. The latter can apply a test function Test to t_w, C to get back 0 or 1. The consistency condition as per [8] is that if $w = w'$ then $\text{Test}(t_w, C)$ returns 1 and if $w \neq w'$ it returns 0. The security condition is that the gateway learn nothing about w' beyond whether or not it equals w . (The corresponding formal notion will be denoted PEKS-IND-CPA.) The application setting is that C can be attached to an email (ordinarily encrypted for Alice under a different public key), allowing the gateway to route the email to different locations (e.g. Alice's desktop, laptop or pager) based on w while preserving privacy of the latter to the largest extent possible.

Consistency of BDOP-PEKS It is easy to see (cf. Proposition 3.1) that the main construction of [8] (a random oracle (RO) model, bilinear Diffie–Hellman (BDH) based

PEKS-IND-CPA secure PEKS scheme that we call *BDOP-PEKS*) fails to meet the consistency condition defined in [8] and stated above. (Specifically, there are distinct keywords w, w' such that $\text{Test}(t_w, C) = 1$ for any C that encrypts w' .) The potential problem this raises in practice is that email will be incorrectly routed.

New Notions of Consistency It is natural to ask if *BDOP-PEKS* meets some consistency condition that is weaker than theirs but still adequate in practice. To answer this, we provide some new definitions. Somewhat unusually for a consistency condition, we formulate consistency more like a security condition, via an experiment involving an adversary. The difference is that this adversary is not very “adversarial”: it is supposed to reflect some kind of worst case but not malicious behavior. However this turns out to be a difficult line to draw, definitionally, so that some subtle issues arise. One outcome of this approach is that it naturally gives rise to a hierarchy of notions of consistency, namely perfect, statistical and computational. The first asks that the advantage of any (even computationally unbounded) adversary be zero; the second that the advantage of any (even computationally unbounded) adversary be negligible; the third that the advantage of any polynomial-time adversary be negligible. We note that perfect consistency as per our definition coincides with consistency as per [8], and so our notions can be viewed as natural weakenings of theirs.

An Analogy There is a natural notion of *decryption error* for encryption schemes [17, Sect. 5.1.2]. A perfectly consistent PEKS is the analog of an encryption scheme with zero decryption error (the usual requirement). A statistically consistent PEKS is the analog of an encryption scheme with negligible decryption error (a less common but still often used condition [2,13]). However, computational consistency is a non-standard relaxation, for consistency conditions are typically not computational. This is not because one cannot define them that way (one could certainly define a computational consistency requirement for encryption) but rather because there has never been any motivation to do so. What makes PEKS different, as emerges from the results below, is that computational consistency is relevant and arises naturally.

Consistency of BDOP-PEKS, Revisited The counter-example (cf. Proposition 3.1) showing that *BDOP-PEKS* is not perfectly consistent extends to show that it is not statistically consistent either. However, we show (cf. Theorem 3.3) that *BDOP-PEKS* is computationally consistent. In the random-oracle model, this is not under any computational assumption: the limitation on the running time of the adversary is relevant because it limits the number of queries the adversary can make to the random oracle. When the random oracle is instantiated via a hash function, we would need to assume collision-resistance of the hash function. The implication of this result is that *BDOP-PEKS* is probably fine to use in practice, in that incorrect routing of email, while possible in principle, is unlikely to actually happen.

A Statistically Consistent PEKS Scheme We provide the first construction of a PEKS scheme that is *statistically* consistent. The scheme is in the random oracle model, and is also PEKS-IND-CPA secure assuming the BDH problem is hard. The motivation for the new scheme was largely theoretical. From a foundational perspective, we wanted to

know whether PEKS was an anomaly in the sense that only computational consistency is possible, or whether, like other primitives, statistical consistency could be achieved. However, it is also true that while computational consistency is arguably enough in an application, statistical might be preferable because the guarantee is unconditional.

1.2. PEKS and Anonymous IBE

BDOP-PEKS is based on the Boneh–Franklin IBE (*BF-IBE*) scheme [9]. It is natural to ask whether one might, more generally, build PEKS schemes from IBE schemes in some blackbox way. To this end, a transform of an IBE scheme into a PEKS scheme is suggested in [8]. Interestingly, they note that the property of the IBE scheme that appears necessary to provide PEKS-IND-CPA of the PEKS scheme is not the usual IBE-IND-CPA but rather anonymity. (An IBE scheme is anonymous if a ciphertext does not reveal the identity of the recipient [3].) While [8] stops short of stating and proving a formal result here, it is not hard to verify that their intuition is correct. Namely one can show that if the starting IBE scheme *IBE* meets an appropriate formal notion of anonymity (IBE-ANO-CPA, cf. Section 4.1) then $\text{PEKS} = \text{ibe-2-peks}(\text{IBE})$ is PEKS-IND-CPA, where *ibe-2-peks* denotes the transform suggested in [8].

Consistency in ibe-2-peks Unfortunately, we show (cf. Theorem 4.1) that there are IBE schemes for which the PEKS scheme outputted by *ibe-2-peks* is not even computationally consistent. This means that *ibe-2-peks* is not in general a suitable way to turn an IBE scheme into a PEKS scheme. (Although it might be in some cases, and in particular is when the starting IBE scheme is *BF-IBE*, for in that case the resulting PEKS scheme is *BDOP-PEKS*.)

new-ibe-2-peks We propose a randomized variant of the *ibe-2-peks* transform that we call *new-ibe-2-peks*, and prove that if an IBE scheme *IBE* is IBE-ANO-CPA and IBE-IND-CPA then the PEKS scheme *new-ibe-2-peks*(*IBE*) is PEKS-IND-CPA and computationally consistent (cf. Sect. 4.3). We do not know of a transform where the resulting PEKS scheme is statistically or perfectly consistent.

Anonymous IBE Schemes The above motivates finding anonymous IBE schemes. Towards this, we begin by extending Halevi’s condition for anonymity [19] to the IBE setting (cf. Sect. 4.4). Based on this, we are able to give a simple proof that the (random-oracle model) *BF-IBE* scheme [9] is IBE-ANO-CPA assuming the BDH problem is hard (cf. Theorem 4.4). (We clarify that a proof of this result is implicit in the proof of security of the *BF-IBE* based *BDOP-PEKS* scheme given in [8]. Our contribution is to have stated the formal definition of anonymity and provided a simpler proof via the extension of Halevi’s condition.) Towards answering the question of whether there exist anonymous IBE schemes in the standard (as opposed to random oracle) model, we present in Appendix A.1 an attack to show that Water’s IBE scheme [25] is not IBE-ANO-CPA.

1.3. Extensions

Anonymous HIBE We provide definitions of anonymity for hierarchical IBE (HIBE) schemes. Our definition can be parameterized by a level, so that we can talk of a HIBE

that is anonymous at level l . We note that the HIBE schemes of [7,15] are not anonymous, even at level 1. (That of [20] appears to be anonymous at both levels 1 and 2 but is very limited in nature and thus turns out not to be useful for our applications.) We modify the construction of Gentry and Silverberg [15] to obtain a HIBE that is (HIBE-IND-CPA and) anonymous at level 1. The construction is in the random oracle model and assumes BDH is hard.

PETKS In a PEKS scheme, once the gateway has the trapdoor for a certain keyword, it can test whether this keyword was present in any past ciphertexts or future ciphertexts. It may be useful to limit the period in which the trapdoor can be used. Here we propose an extension of PEKS that we call public-key encryption with temporary keyword search (PETKS) that allows this. A trapdoor here is created for a time interval $[s, e]$ and will only allow the gateway to test whether ciphertexts created in this time interval contain the keyword. We provide definitions of privacy and consistency for PETKS, and then show how to implement it with overhead that is only logarithmic in the total number of time periods. Our construction can use any HIBE that is anonymous at level 1. Using the above-mentioned HIBE we get a particular instantiation that is secure in the random-oracle model if BDH is hard.

IBEKS We define the notion of an identity-based encryption with keyword search scheme. This is just like a PEKS scheme except that encryption is performed given only the identity of the receiver and a master public-key, just like in an IBE scheme. We show how to implement IBEKS given any level-2 anonymous HIBE scheme. The first suitable implementation of the latter primitive was proposed in subsequent work by Boyen and Waters [11].

1.4. Remarks

peks-2-ibe Boneh et al. [8] showed how to transform a PEKS-IND-CPA PEKS scheme into an IBE-IND-CPA IBE scheme. We remark that their transform requires the starting PEKS scheme to be perfectly consistent. Unfortunately, no perfectly consistent PEKS schemes are known to date. If it is only statistically or computationally consistent, the resulting IBE scheme will only meet a corresponding statistical or computational relaxation of the consistency condition for IBE schemes. Thus, the resulting scheme will not be an IBE scheme as per the standard definition of the latter [9].

Limited PEKS Schemes Boneh et al. [8] also present a couple of PEKS schemes that avoid the RO model but are what they call *limited*. Both use a standard public-key encryption scheme as a building block. In the first scheme, the public key has size polynomial in the number of keywords that can be used. In the second scheme, the key and ciphertext have size polynomial in the number of trapdoors that can be securely issued to the gateway. Although these schemes are not very interesting due to their limited nature, one could ask about their consistency. In [1], we extend our definitions of consistency to this limited setting. Interestingly, we show that based on only a computational assumption about the underlying standard public-key encryption scheme (namely, that it is IND-CPA, or even just one-way), the first scheme is *statistically* consistent. We also show that the second scheme is computationally consistent under the same assumption

on the standard public-key encryption scheme, and present a variant that is statistically consistent.

Consistency of Other Searchable Encryption Schemes Of the other papers on searchable encryption of which we are aware [16,18,24,26], none formally define or rigorously address the notion of consistency for their respective types of searchable encryption schemes. Goh [16] and Golle, Staddon, and Waters [18] define consistency conditions analogous to BDOP’s “perfect consistency” condition, but none of the constructions in [16,18] satisfy their respective perfect consistency condition. Song, Wagner, and Perrig [24] and Waters et al. [26] do not formally state and prove consistency conditions for their respective searchable encryption schemes, but they, as well as Goh [16], do acknowledge and informally bound the non-zero probability of a false positive.

Subsequent Work In a preliminary version of our work, we raised various open problems that have subsequently been solved. The first one of these problems was to find a construction of an (IBE-IND-CPA and) IBE-ANO-CPA IBE scheme with a proof of security in the standard model (i.e., without random oracles). This problem was solved independently by Gentry [14] and by Boyen and Waters [11]. As a result, one can also obtain a PEKS-IND-CPA and computationally consistent PEKS scheme in the standard model due to Theorem 4.2.

Another interesting question that we raised was to find a HIBE scheme providing anonymity at the second level, even in the RO model. This open problem was solved by Boyen and Waters [11], who proposed a fully anonymous HIBE scheme in the standard model.

Finally, we raised the issue of building a searchable encryption scheme that allows for more advanced searching tools such as searches for simple boolean formulas on keywords (say $w_1 \wedge w_2 \vee w_3$). First steps in this direction have been taken [10,18,22] by schemes that allow for conjunctive combinations of keywords, range queries, and subset queries.

2. Some Definitions

Notation and Conventions If x is a string then $|x|$ denotes its length, and if S is a set then $|S|$ denotes its size. The empty string is denoted ε . Constructs in the RO model [5] might use multiple random oracles, but since one can always obtain these from a single one [5], formal definitions will assume just one RO. Unless otherwise indicated, an algorithm may be randomized. “PT” stands for polynomial time and “PTA” for polynomial-time algorithm or adversary. We denote by \mathbb{N} the set of positive integers, and by $k \in \mathbb{N}$ the security parameter. A function $\nu : \mathbb{N} \rightarrow [0, 1]$ is said to be *negligible* if for every $c \in \mathbb{N}$ there exists a $k_c \in \mathbb{N}$ such that $\nu(k) \leq k^{-c}$ for all $k > k_c$, and it is said to be *overwhelming* if the function $|1 - \nu(k)|$ is negligible. A *message space* MsgSp is a map, assigning to every $k \in \mathbb{N}$ a set of strings, such that $\{0, 1\}^k \subseteq \text{MsgSp}(k) \subseteq \{0, 1\}^*$ for all $k \in \mathbb{N}$ and the following conditions hold: first, there is a PTA that on input $1^k, M$ returns 1 if $M \in \text{MsgSp}(k)$ and 0 otherwise; second, $\{0, 1\}^{|M|} \subseteq \text{MsgSp}(k)$ for all $k \in \mathbb{N}$ and $M \in \text{MsgSp}(k)$.

PEKS A *public key encryption with keyword search* (PEKS) scheme [8] $\mathcal{PEKS} = (\text{KG}, \text{PEKS}, \text{Td}, \text{Test})$ consists of PTAs. Via $(pk, sk) \xleftarrow{\$} \text{KG}(1^k)$, where $k \in \mathbb{N}$ is the security parameter and KG is the randomized key-generation algorithm, the receiver produces its keys; via $C \xleftarrow{\$} \text{PEKS}^H(pk, w)$ a sender encrypts a keyword w to get a ciphertext; via $t_w \xleftarrow{\$} \text{Td}^H(sk, w)$ the receiver computes a trapdoor t_w for keyword w and provides it to the gateway; via $b \leftarrow \text{Test}^H(t_w, C)$ the gateway tests whether C encrypts w , where b is a bit with 1 meaning “accept” or “yes” and 0 meaning “reject” or “no”. Here H is a random oracle whose domain and/or range might depend on k and pk .

Consistency The requirement of [8] can be divided into two parts. The first, which we call *right keyword consistency*, is that $\text{Test}(t_w, C)$ always accepts when C encrypts w . More formally, for all $k \in \mathbb{N}$ and all $w \in \{0, 1\}^*$,

$$\Pr[\text{Test}^H(\text{Td}^H(sk, w), \text{PEKS}^H(pk, w)) = 1] = 1,$$

where the probability is taken over the choice of $(pk, sk) \xleftarrow{\$} \text{KG}(1^k)$, the random choice of H , and the coins of all the algorithms in the expression above. Since we will always require this, it is convenient henceforth to take it as an integral part of the PEKS notion and not mention it again, reserving the term “consistency” to only refer to what happens when the ciphertext encrypts a keyword different from the one for which the gateway is testing. In this regard, the requirement of [8], which we will call *perfect consistency*, is that $\text{Test}(t_w, C)$ always reject when C doesn’t encrypt w' . More formally, for all $k \in \mathbb{N}$ and all distinct $w, w' \in \{0, 1\}^*$,

$$\Pr[\text{Test}^H(\text{Td}^H(sk, w'), \text{PEKS}^H(pk, w)) = 1] = 0,$$

where the probability is taken over the choice of $(pk, sk) \xleftarrow{\$} \text{KG}(1^k)$, the random choice of H , and the coins of all the algorithms in the expression above. (We note that [8] provide informal rather than formal statements, but it is hard to interpret them in any way other than what we have done.)

Privacy Privacy for a PEKS scheme [8] asks that an adversary should not be able to distinguish between the encryption of two challenge keywords of its choice, even if it is allowed to obtain trapdoors for any non-challenge keywords. Formally, we associate to an adversary \mathcal{A} and a bit $b \in \{0, 1\}$ the following experiment:

<p>Experiment $\text{Exp}_{\mathcal{PEKS}, \mathcal{A}}^{\text{peks-ind-cpa-b}}(k)$</p> <p>$WSet \leftarrow \emptyset$; $(pk, sk) \xleftarrow{\\$} \text{KG}(1^k)$ pick random oracle H $(w_0, w_1, state) \xleftarrow{\\$} \mathcal{A}^{\text{TRAPD}(\cdot), H}(\text{find}, pk)$ $C \xleftarrow{\\$} \text{PEKS}^H(pk, w_b)$ $b' \xleftarrow{\\$} \mathcal{A}^{\text{TRAPD}(\cdot), H}(\text{guess}, C, state)$ if $\{w_0, w_1\} \cap WSet = \emptyset$ then return b' else return 0</p>	<p>Oracle $\text{TRAPD}(w)$</p> <p>$WSet \leftarrow WSet \cup \{w\}$ $t_w \xleftarrow{\\$} \text{Td}^H(sk, w)$ return t_w</p>
--	---

The PEKS-IND-CPA-*advantage* of \mathcal{A} is defined as

$$\mathbf{Adv}_{\mathcal{PEKS}, \mathcal{A}}^{\text{peks-ind-cpa}}(k) = \Pr[\mathbf{Exp}_{\mathcal{PEKS}, \mathcal{A}}^{\text{peks-ind-cpa-1}}(k) = 1] - \Pr[\mathbf{Exp}_{\mathcal{PEKS}, \mathcal{A}}^{\text{peks-ind-cpa-0}}(k) = 1].$$

A scheme \mathcal{PEKS} is said to be PEKS-IND-CPA-*secure* if the above advantage is a negligible function in k for all PTAs \mathcal{A} .

Parameter Generation Algorithms and the BDH Problem All pairing based schemes will be parameterized by a *pairing parameter generator*. This is a PTA \mathcal{G} that on input 1^k returns the description of an additive cyclic group \mathbb{G}_1 of prime order p , where $2^k < p < 2^{k+1}$, the description of a multiplicative cyclic group \mathbb{G}_2 of the same order, and a non-degenerate bilinear pairing $e: \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$. See [9] for a description of the properties of such pairings. We use \mathbb{G}_1^* to denote $\mathbb{G}_1 \setminus \{0\}$, i.e. the set of all group elements except the neutral element. We define the advantage of an adversary \mathcal{A} in solving the bilinear Diffie–Hellman (BDH) problem relative to a pairing parameter generator \mathcal{G} as

$$\begin{aligned} \mathbf{Adv}_{\mathcal{G}, \mathcal{A}}^{\text{bdh}}(k) &= \Pr \left[\mathcal{A}(1^k, (\mathbb{G}_1, \mathbb{G}_2, p, e), P, aP, bP, cP) \right. \\ &= e(P, P)^{abc}: \left. \begin{array}{l} (\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\$} \mathcal{G}(1^k); \\ P \xleftarrow{\$} \mathbb{G}_1^*; a, b, c \xleftarrow{\$} \mathbb{Z}_p^* \end{array} \right]. \end{aligned}$$

We say that the BDH problem is hard relative to this generator if $\mathbf{Adv}_{\mathcal{G}, \mathcal{A}}^{\text{bdh}}$ is a negligible function in k for all PTAs \mathcal{A} .

3. Consistency in PEKS

We show that the $\mathcal{BDOP}\text{-}\mathcal{PEKS}$ scheme is not perfectly consistent, introduce new notions of statistical and computational consistency, and show that although $\mathcal{BDOP}\text{-}\mathcal{PEKS}$ continues to fail the former it does meet the latter. We then provide a new PEKS scheme that is statistically consistent.

3.1. Perfect Consistency of $\mathcal{BDOP}\text{-}\mathcal{PEKS}$

Figure 1 presents the $\mathcal{BDOP}\text{-}\mathcal{PEKS}$ scheme. It is based on a pairing parameter generator \mathcal{G} .

Proposition 3.1. *The $\mathcal{BDOP}\text{-}\mathcal{PEKS}$ scheme is not perfectly consistent.*

Proof. Since the number of possible keywords is infinite, there will certainly exist distinct keywords $w, w' \in \{0, 1\}^*$ such that $H_1(w) = H_1(w')$. The trapdoors for such keywords will be the same, and so $\text{Test}^{H_1, H_2}(\text{Td}(sk, w), \text{PEKS}^{H_1, H_2}(pk, w'))$ will always return 1. \square

It is tempting to say that, since H_1 is a random oracle, the probability of a collision is small, and thus the above really does not matter. Whether or not this is true depends on how one wants to define consistency, which is the issue we explore next.

$\text{KG}(1^k)$ $(\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\$} \mathcal{G}(1^k); P \xleftarrow{\$} \mathbb{G}_1^*; s \xleftarrow{\$} \mathbb{Z}_p^*$ $pk \leftarrow (\mathbb{G}_1, \mathbb{G}_2, p, e, P, sP); sk \leftarrow (pk, s)$ $\text{return } (pk, sk)$	$\text{Td}^{H_1}(sk, w)$ $\text{parse } sk \text{ as } (pk = (\mathbb{G}_1, \mathbb{G}_2, p, e, P, sP), s)$ $t_w \leftarrow (pk, sH_1(w)); \text{return } t_w$
$\text{PEKS}^{H_1, H_2}(pk, w)$ $\text{parse } pk \text{ as } (\mathbb{G}_1, \mathbb{G}_2, p, e, P, sP)$ $r \xleftarrow{\$} \mathbb{Z}_p^*; T \leftarrow e(H_1(w), sP)^r$ $C \leftarrow (rP, H_2(T)); \text{return } C$	$\text{Test}^{H_1, H_2}(t_w, C)$ $\text{parse } t_w \text{ as } ((\mathbb{G}_1, \mathbb{G}_2, p, e, P, sP), X)$ $\text{parse } C \text{ as } (U, V); T \leftarrow e(X, U)$ $\text{if } V = H_2(T) \text{ then return } 1$ $\text{else return } 0$

Fig. 1. Algorithms constituting the $\mathcal{BDOP}\text{-PEKS}$ scheme. \mathcal{G} is a pairing parameter generator and $H_1: \{0, 1\}^* \rightarrow \mathbb{G}_1$ and $H_2: \mathbb{G}_2 \rightarrow \{0, 1\}^k$ are random oracles.

3.2. New Notions of Consistency

We consider a possible relaxation of perfect consistency and argue that it is inadequate because it is too weak. We then motivate and present our approach and definitions.

A Possible Relaxation of Perfect Consistency One way to obtain a relaxed definition of perfect consistency is by analogy with the definition of encryption with negligible decryption error [17, Sect. 5.1.2]. This results in asking that there exist a negligible function $\nu(\cdot)$ such that for all k and all distinct keywords w, w' ,

$$\forall w \neq w': \Pr[\text{Test}^H(pk, \text{Td}^H(sk, w')), \text{PEKS}^H(pk, w)] = 1] \leq \nu(k), \quad (1)$$

where the probability is taken over the choice of $(pk, sk) \xleftarrow{\$} \text{KG}(1^k)$, the random choice of H , and the coins of all the algorithms in the expression above. Now, since we are fixing w, w' before taking the probability, and the latter includes the choice of H_1 in the $\mathcal{BDOP}\text{-PEKS}$ scheme, the probability that $H_1(w) = H_1(w')$ is at most 2^{-k} . Our “attack” of Proposition 3.1 therefore no longer applies. And in fact (using the techniques of our proof of Theorem 3.3) one can show that the BDOP scheme *does* meet the above condition. However, (1) is in our view an incorrect definition of consistency because it does not allow w, w' to depend on public quantities related to the receiver, such as its public key, the hash functions being used, or queries to them if they are random oracles. Our claim is that, as a result, the condition is too weak to guarantee that email is correctly routed by the gateway.

Our Definitions To define consistency, we take a different approach. Namely, we imagine the existence of an adversary \mathcal{U} that wants to make consistency fail. More precisely, let $\text{PEKS} = (\text{KG}, \text{PEKS}, \text{Td}, \text{Test})$ be a PEKS scheme. We associate to an adversary \mathcal{U} the following experiment:

Experiment $\text{Exp}_{\text{PEKS}, \mathcal{U}}^{\text{peks-consist}}(k)$
 $(pk, sk) \xleftarrow{\$} \text{KG}(1^k);$ pick random oracle H
 $(w, w') \xleftarrow{\$} \mathcal{U}^H(pk); C \xleftarrow{\$} \text{PEKS}^H(pk, w); t_{w'} \xleftarrow{\$} \text{Td}^H(sk, w')$
 if $w \neq w'$ and $\text{Test}^H(t_{w'}, C) = 1$ then return 1 else return 0.

We define the advantage of \mathcal{U} as

$$\mathbf{Adv}_{\text{PEKS}, \mathcal{U}}^{\text{peks-consist}}(k) = \Pr[\mathbf{Exp}_{\text{PEKS}, \mathcal{U}}^{\text{peks-consist}}(k) = 1],$$

where the probability is taken over all possible coin flips of all the algorithms involved, and over all possible choices of random oracle H . The scheme is said to be *perfectly consistent* if this advantage is 0 for all (computationally unrestricted) adversaries \mathcal{U} , *statistically consistent* if it is negligible for all (computationally unrestricted) adversaries \mathcal{U} , and *computationally consistent* if it is negligible for all PTAs \mathcal{U} . We remark that we have purposely re-used the term perfect consistency, for in fact the above notion of perfect consistency coincides with the one from [8] recalled above.

Stronger Notions? In giving the adversary \mathcal{U} the public key and access to the random oracle, our definition is already quite liberal. One could however, consider an even more liberal (i.e. stronger) definition in which the adversary gets a trapdoor oracle and/or a test oracle under trapdoors for keywords of its choice. To be able to tell whether or not this would be appropriate, we must ask whether in “real-life” there could be an occasion in which the keywords chosen by a sender could depend on information provided by these oracles. Given that the answer is not cut-and-dry and since we believe that our current definition is already quite strong, we opted here not to consider these stronger variants of our definition.

3.3. Statistical and Computational Consistency of $\mathcal{BDOP}\text{-PEKS}$

Having formally defined the statistical and computational consistency requirements for PEKS schemes, we return to evaluating the consistency of $\mathcal{BDOP}\text{-PEKS}$. We first observe that Proposition 3.1 extends to show:

Proposition 3.2. *The $\mathcal{BDOP}\text{-PEKS}$ scheme is not statistically consistent.*

Proof. Recall that in the proof of Proposition 3.1 we show that there exist two distinct keywords $w, w' \in \{0, 1\}^*$ such that $H_1(w) = H_1(w')$, and that, for these two keywords, $\text{Test}(\text{Td}(sk, w'), \text{PEKS}(pk, w))$ will always return 1. A computationally unbounded adversary can find two such keywords by exhaustive search. \square

On the positive side, the following means that $\mathcal{BDOP}\text{-PEKS}$ is probably fine in practice:

Theorem 3.3. *The $\mathcal{BDOP}\text{-PEKS}$ scheme is computationally consistent.*

Proof. Let \mathcal{U} be a PTA. Let (w, w') denote the pair of keywords that \mathcal{U} returns in the consistency experiment, and assume without loss of generality that $w \neq w'$. Let $r \in \mathbb{Z}_p^*$ denote the value chosen at random by $\text{PEKS}^{H_1, H_2}(pk, w)$. Let $T = e(H_1(w), sP)^r$ and let $T' = e(H_1(w'), sP)^r$. Note that \mathcal{U} wins exactly when $w \neq w'$ and $H_2(T) = H_2(T')$. Let w_1, \dots, w_{q_1} be the queries of \mathcal{U} to H_1 and let $WSet = \{w_1, \dots, w_{q_1(k)}\} \cup \{w, w'\}$. Let $T_1, \dots, T_{q_2(k)}$ be the queries of \mathcal{U} to H_2 and let $TSet = \{T_1, \dots, T_{q_2(k)}\} \cup \{T, T'\}$. Let E_1 be the event that there exist distinct $v, v' \in WSet$ such that $H_1(v) = H_1(v')$, and

let E_2 be the event that there exist distinct $x, x' \in TSet$ such that $H_2(x) = H_2(x')$. If $[\cdot]$ denotes the probability in the consistency experiment, then

$$\mathbf{Adv}_{\mathcal{PEKS}, \mathcal{U}}^{\text{peks-consist}}(k) \leq \Pr[E_1] + \Pr[E_2] + \Pr[\mathbf{Exp}_{\mathcal{BDOP-PEKS}, \mathcal{U}}^{\text{peks-consist}}(k) = 1 \wedge \overline{E_1} \wedge \overline{E_2}]. \quad (2)$$

Our definition of \mathcal{G} required that $|\mathbb{G}_1| > 2^k$, and hence the first and second terms are respectively upper bounded via $(q_1 + 2)^2/|\mathbb{G}_1| < (q_1 + 2)^2/2^k$ and $(q_2 + 2)^2/2^k$. Now we claim that if $H_1(w) \neq H_1(w')$, then $T \neq T'$. Under this claim, the last term of (2) is 0, since if $\overline{E_1}$ occurs, then $H_1(w) \neq H_1(w')$ and $T \neq T'$, and if $\overline{E_2}$ also occurs, then $H_2(T) \neq H_2(T')$. To justify our claim above, note that if $H_1(w) \neq H_1(w')$, then $H_1(w) = \alpha P$ and $H_1(w') = \alpha' P$ for some distinct $\alpha, \alpha' \in \mathbb{Z}_p$. Setting $g = e(P, P)^{rs}$, we can rewrite T, T' as $T = g^\alpha$ and $T' = g^{\alpha'}$. Since $e(P, P)$ is a generator of \mathbb{G}_2 , since \mathbb{G}_2 is of prime order p , and since p does not divide rs , g must also be a generator of \mathbb{G}_2 . Thus $T \neq T'$. \square

3.4. A Statistically Consistent PEKS Scheme

We present the first PEKS scheme that is (PEKS-IND-CPA and) *statistically* consistent. To define the scheme, we first introduce the function $f(k) = k^{\lg(k)}$. (Any function that is super-polynomial but sub-exponential would suffice. This choice is made for concreteness.) The algorithms constituting our scheme $\mathcal{PEKS-STAT}$ are then depicted in Fig. 2.

The scheme uses ideas from the $\mathcal{BDOP-PEKS}$ scheme [8] as well as from the $\mathcal{BF-IBE}$ scheme [9], but adds some new elements. Note that the encryption algorithm is trivial, returning the keyword as the ciphertext, when the keyword has length more than $f(k)$. If not, the processing is more complex, depending on some random choices and numerous random oracles. In particular the random choice of “session” key K , and the fact that the random oracle H_2 is length-increasing, are important.

The first thing we stress about the scheme is that the algorithms are PT. This is because PT means in the length of the inputs, and the input of (say) the encryption algorithm includes w as well as 1^k , so it can test whether $|w| \geq f(k)$ in polynomial time. Now the following says that the scheme is private:

Proposition 3.4. *The $\mathcal{PEKS-STAT}$ scheme is PEKS-IND-CPA-secure assuming that the BDH problem is hard relative to generator \mathcal{G} .*

Before providing the proof, let us give some intuition. While sending w in the clear looks at first glance like it violates privacy, the reason it does not is that this only happens when w has length at least $f(k)$, and the privacy adversary is $\text{poly}(k)$ time and thus cannot even write down such a keyword in order to query it to its challenge oracle. (This is where we use the fact that $f(k)$ is super-polynomial. We will use the fact that it is sub-exponential in the proof of statistical consistency.) The privacy adversary is thus effectively restricted to attacking the scheme only on keywords of size at most $f(k)$. Here, privacy can be reduced to solving the BDH problem using techniques used to prove IBE-IND-CPA of the $\mathcal{BF-IBE}$ scheme [9] and to prove anonymity of the same scheme (cf. Theorem 4.4).

<p>KG(1^k)</p> <p>$(\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\\$} \mathcal{G}(1^k); P \xleftarrow{\\$} \mathbb{G}_1^*$</p> <p>$s \xleftarrow{\\$} \mathbb{Z}_p^*; pk \leftarrow (1^k, P, sP, \mathbb{G}_1, \mathbb{G}_2, p, e)$</p> <p>$sk \leftarrow (pk, s); \text{return } (pk, sk)$</p>	<p>Td$^{H_1}(sk, w)$</p> <p>parse sk as $(pk = (1^k, P, sP, \mathbb{G}_1, \mathbb{G}_2, p, e), s)$</p> <p>$t_w \leftarrow (pk, sH_1(w), w)$</p> <p>return t_w</p>
<p>PEKS$^{H_1, H_2, H_3, H_4}(pk, w)$</p> <p>parse pk as $(1^k, P, sP, \mathbb{G}_1, \mathbb{G}_2, p, e)$</p> <p>if $w \geq f(k)$ then return w</p> <p>$r \xleftarrow{\\$} \mathbb{Z}_p^*; T \leftarrow e(sP, H_1(w))^r$</p> <p>$K_1 \leftarrow H_4(T); K_2 \leftarrow H_2(T)$</p> <p>$K \xleftarrow{\\$} \{0, 1\}^k; c \leftarrow K_1 \oplus K$</p> <p>$t \leftarrow H_3(K w)$</p> <p>return (rP, c, t, K_2)</p>	<p>Test$^{H_1, H_2, H_3, H_4}(t_w, C)$</p> <p>parse t_w as $((1^k, P, sP, \mathbb{G}_1, \mathbb{G}_2, p, e), sH_1(w), w)$</p> <p>if $w \geq f(k)$ then</p> <p style="padding-left: 20px;">if $C = w$ then return 1 else return 0</p> <p>if C cannot be parsed as (rP, c, t, K_2)</p> <p>then return 0</p> <p>$T \leftarrow e(rP, sH_1(w))$</p> <p>$K \leftarrow c \oplus H_4(T)$</p> <p>if $K_2 \neq H_2(T)$ then return 0</p> <p>if $t = H_3(K w)$ then return 1 else return 0</p>

Fig. 2. Algorithms constituting the PEKS scheme PEKS-STAT . Here $f(k) = k^{\lg(k)}$, \mathcal{G} is a pairing parameter generator and $H_1: \{0, 1\}^* \rightarrow \mathbb{G}_1$, $H_2: \mathbb{G}_2 \rightarrow \{0, 1\}^{3k}$, $H_3: \{0, 1\}^* \rightarrow \{0, 1\}^k$, and $H_4: \{0, 1\}^* \rightarrow \{0, 1\}^k$ are random oracles.

Proof of Proposition 3.4. Let \mathcal{B} be a PTA attacking the PEKS-IND-CPA security $\text{PEKS-STAT} = (\text{KG}, \text{PEKS}, \text{Td}, \text{Test})$. Say it makes at most q queries to its $\text{TRAPD}(\cdot, \cdot)$ oracle and at most q_i queries to H_i for $i = 1, 2, 3$. (These are actually functions of k , but we drop the argument to simplify notation.) We construct a PTA \mathcal{A} attacking the BDH relative to \mathcal{G} such that

$$\text{Adv}_{\mathcal{G}, \mathcal{A}}^{\text{bdh}}(k) \geq \frac{1}{e(1+q) \cdot (q_2 + q_4)} \cdot \left(\frac{1}{2} \cdot \text{Adv}_{\text{PEKS}, \mathcal{B}}^{\text{peks-ind-cpa}}(k) - \frac{q_3}{2^k} \right). \quad (3)$$

Our adversary \mathcal{A} is shown in Fig. 3. We show that \mathcal{A} outputs the correct answer $T = e(P, P)^{r \cdot s \alpha}$ with probability at least the quantity on the right-hand-side of (3).

Let $t(k)$ be a polynomial which bounds the running time of \mathcal{B} . So there is an integer N such that $t(k) < f(k)$ for all $k \geq N$. Notice that the PEKS algorithm of the PEKS scheme in Fig. 2 returns w in the clear when $|w| \geq f(k)$. However, the keywords output by \mathcal{B} in the find stage have length at most $t(k)$, so if $k \geq N$, the encryption is done by the code for the case $|w| < f(k)$ shown in PEKS. Since it suffices to prove (3) for all $k \geq N$, we assume that the encryption is done by the code for the case $|w| < f(k)$ shown in PEKS.

Let $\text{Pr}_1[\cdot]$ denote the probability over the experiment for $\text{Adv}_{\mathcal{G}, \mathcal{A}}^{\text{bdh}}(k)$ as defined in Sect. 2. Let E_1 denote the event in this experiment that \mathcal{A} aborts in simulating the trapdoor oracle. Let E_2 denote the event that $d[w_b] = 0$ (which also causes \mathcal{A} to abort). Let E_3 denote the event that $\mathcal{Q} = \emptyset$ (which also causes \mathcal{A} to abort). Let E_4 denote the event that \mathcal{B} issues a query $H_2(e(rP, sH_1(w_b)))$ or $H_4(e(rP, sH_1(w_b)))$. Let $\text{Pr}_2[\cdot]$ denote the probability over $\text{Exp}_{\text{PEKS}, \mathcal{B}}^{\text{peks-ind-cpa-b}}$ for a random choice for $b \in \{0, 1\}$, and let b' denote the output of \mathcal{B} in this experiment. Let E_5 be the event that \mathcal{B} issues a query $H_2(e(rP, sH_1(w_b)))$ or $H_4(e(rP, sH_1(w_b)))$ to its oracles in this experiment. Let E_6 denote the event that \mathcal{B} issues a query $K||w_b$ to its oracle H_3 , where K is the random

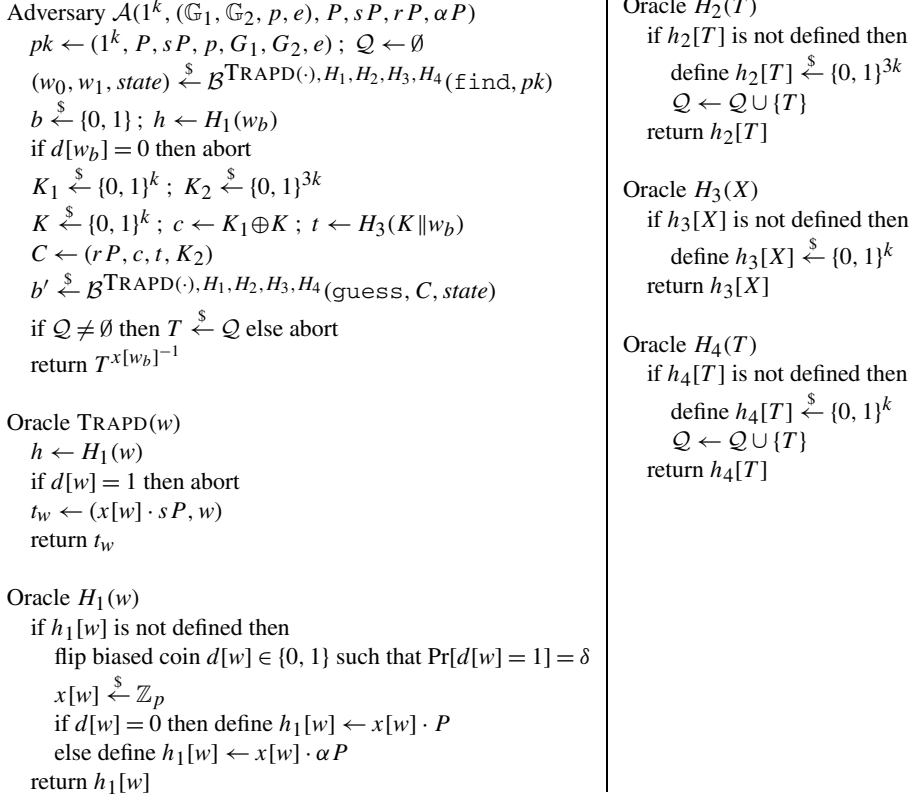


Fig. 3. Adversary \mathcal{A} attacking the BDH problem.

k -bit string that $\text{Exp}_{\text{PEKS}, \mathcal{B}}^{\text{peks-ind-cpa-b}}$ used in PEKS when replying \mathcal{B} 's challenge after the first stage. Equation (3) follows from the following claims. \square

Claim 1. $\text{Adv}_{\mathcal{G}, \mathcal{A}}^{\text{bdh}}(k) \geq \Pr_1[\neg E_1 \wedge \neg E_2 \wedge \neg E_3 \wedge E_4] / (q_2 + q_4)$.

Proof. In the above simulation if none of the events E_1, E_2 and E_3 happens, then \mathcal{A} will randomly choose an element $T \xleftarrow{\$} \mathcal{Q}$ and return $T^{x[w_b]^{-1}}$. However, by definition of event E_4 , one of the elements in \mathcal{Q} is equal to $e(P, P)^{sr\alpha \cdot x[w_b]}$, thus \mathcal{A} has at least the probability of $1/|\mathcal{Q}| \geq 1/(q_2 + q_4)$ to give the correct answer to the BDH problem. \square

Claim 2. $\Pr[\neg E_1 \wedge \neg E_2 \wedge \neg E_3 \wedge E_4] = \Pr[E_4 | \neg E_1 \wedge \neg E_2] \cdot \Pr[\neg E_1 \wedge \neg E_2]$.

Proof. Notice that when event E_4 happens, the set \mathcal{Q} must contain at least one element, thus E_3 is always false. Therefore we have $\Pr_1[\neg E_1 \wedge \neg E_2 \wedge \neg E_3 \wedge E_4] = \Pr_1[\neg E_1 \wedge \neg E_2 \wedge E_4]$. The claim follows by conditioning off of the event $\neg E_1 \wedge \neg E_2$. \square

Claim 3. $\Pr_1[E_4 | \neg E_1 \wedge \neg E_2] = \Pr_2[E_5]$.

Proof. Under the condition that \mathcal{A} does not abort, the simulation is perfect, i.e. all \mathcal{A} 's answers to the simulated oracles $\text{TRAPD}(sk, \cdot)$, $H_1(\cdot) \dots H_4(\cdot)$ have exactly the same distribution as those in the real PEKS-IND-CPA experiment. \square

Claim 4. $\Pr_2[E_5] \geq 1/2 \cdot \text{Adv}_{\text{PEKS}, \mathcal{B}}^{\text{peks-ind-cpa}}(k) - q_3 \cdot 2^{-k}$.

Proof. First observe that

$$\begin{aligned} \Pr_2[b = b'] &= \Pr_2[b = b' \wedge E_5] + \Pr_2[b = b' \wedge \neg E_5 \wedge E_6] + \Pr_2[b = b' \wedge \neg E_5 \wedge \neg E_6] \\ &\leq \Pr_2[E_5] + \Pr_2[E_6] + \Pr_2[b = b' \mid \neg E_5 \wedge \neg E_6] \cdot \Pr_2[\neg E_5 \wedge \neg E_6] \\ &\leq \Pr_2[E_5] + q_3 \cdot 2^{-k} + \Pr_2[b = b' \mid \neg E_5 \wedge \neg E_6] \cdot \Pr_2[\neg E_5 \wedge \neg E_6] \quad (4) \\ &\leq \Pr_2[E_5] + q_3 \cdot 2^{-k} + 1/2. \quad (5) \end{aligned}$$

Equation (4) comes from the fact that, by assumption, \mathcal{B} makes at most q_3 queries to H_3 . Equation (5) comes from the fact that, if E_5 and E_6 both do not occur, \mathcal{B} learns no information from the ciphertext. Rearranging gives

$$\Pr_2[E_5] \geq \Pr_2[b = b'] - 1/2 - q_3 \cdot 2^{-k} = 1/2 \cdot \text{Adv}_{\text{PEKS}, \mathcal{B}}^{\text{peks-ind-cpa}}(k) - q_3 \cdot 2^{-k}.$$

The last equality follows from the standard result that $\text{Adv}_{\text{PEKS}, \mathcal{B}}^{\text{peks-ind-cpa}}(k) = 2 \cdot \Pr_2[b = b'] - 1$. \square

Claim 5. $\Pr[\neg E_1 \wedge \neg E_2] \geq 1/(e(q+1))$ for $\delta = 1/(q+1)$.

Proof. Since for every keyword w the biased coin $d[w]$ is flipped independently, and $\Pr[d[w] = 1] = \delta$ for all w , let \mathcal{Q}_T be the set of queries issued by \mathcal{B} to the $\text{TRAPD}(sk, \cdot)$ oracle, then

$$\Pr[\neg E_1 \wedge \neg E_2] = \delta \cdot \prod_{w \in \mathcal{Q}_T} (1 - \delta) = \delta \cdot (1 - \delta)^{|\mathcal{Q}_T|} \geq \delta \cdot (1 - \delta)^q.$$

The last quantity is maximized at $\delta = 1/(q+1)$ with value at least $1/e(q+1)$. \square

Let us move to the more interesting claim, namely consistency:

Proposition 3.5. *The PEKS-STAT scheme is statistically consistent.*

Before providing the proof, let us give some intuition. The main issue is that the computationally unbounded consistency adversary \mathcal{U} can easily find any collisions that exist for the random-oracle hash functions. Let w, w' denote the keywords output by the adversary \mathcal{U} . We proceed via a case analysis. One can show that if either w or w' have length at least $f(k)$ then Test will not be wrong. The interesting case is when w, w' both have length at most $f(k)$. Let (rP, c, t, K_2) denote the challenge ciphertext formed by

encrypting w . Let $T = e(rP, H_1(w))$ and let $K = c \oplus H_4(T)$ be the underlying session key. Let $T' = e(rP, H_1(w'))$ and let $K' = c \oplus H_4(T')$. Now consider two cases.

The first case is that $H_1(w) \neq H_1(w')$. Properties of pairings imply $T \neq T'$. Now we claim that this means $K_2 = H_2(T) \neq H_2(T')$ with high probability, and thus Test will correctly reject, meaning \mathcal{U} does not win. This is *not* merely because H_2 is random, for remember the adversary is not computationally bounded and can search for, and find, any collisions that exist. The reason is that H_2 is with high probability an injective function and collisions for it simply do not exist. The reason for this is that its domain is \mathbb{G}_2 which has size $p < 2^{k+1}$ (our definition of a pairing parameter generator required this) but H_2 outputs $3k$ bits, and thus a union bound can be used to show that H_2 is injective except with probability $4 \cdot 2^{-k}$.

The second case, which is the harder one, is that $H_1(w) = H_1(w')$ (again, we cannot prevent \mathcal{U} from finding collisions in H_1), and this is where we will use the fact that $f(k)$ is sub-exponential. Here the idea is that at the time it chooses w, w' , adversary \mathcal{U} does not know the value of the session key K that is randomly chosen later. We divide pairs (V, V') of strings of length at most $f(k)$ (candidate keywords) into two classes. A pair is *heavy* if there are “lots” of session keys L such that $H_3(L \parallel V) = H_3(L \parallel V')$, and *light* otherwise, where “lots” is defined as $2^{k/2}$. Now we again consider two cases. If (w, w') is light then the randomly chosen K has only a $2^{-k/2}$ chance of being a session key for which $H_3(K \parallel w) = H_3(K \parallel w')$ and thus Test will most likely reject, so \mathcal{U} does not win. Next we use an occupancy problem based counting argument to show that the probability (over H_3) that a particular pair (V, V') of keywords is heavy is *double* exponentially small in k . But the number of choices of keyword pairs is $2^{O(f(k))}$ which is sub-double-exponentially small by choice of $f(k)$, and thus a union bound allows us to conclude that (w, w') is not likely to be heavy.

Proof of Proposition 3.5. Let \mathcal{U} be a computationally unbounded adversary algorithm. We show that there is a constant $c > 0$ such that

$$\text{Adv}_{\text{PEKS}, \mathcal{S}, \mathcal{U}}^{\text{peks-consist}}(k) \leq O(2^{-ck}).$$

Consider the experiment $\text{Exp}_{\text{PEKS}, \mathcal{S}, \mathcal{U}}^{\text{peks-consist}}(k)$. Let w, w' denote the keywords output by \mathcal{U} and assume they are distinct, since otherwise \mathcal{U} does not win. Let WIN be the event that the experiment outputs 1. Let r, \mathbf{K} be the random choices made by $\text{PEKS}_{H_1, H_2, H_3, H_4}(pk, w)$ in the experiment. Then we let

$$\begin{aligned} \mathbf{T} &= e(rP, sH_1(w)), & \mathbf{T}' &= e(rP, sH_1(w')), \\ \mathbf{c} &= \mathbf{K} \oplus H_4(\mathbf{T}), & \mathbf{K}' &= \mathbf{c} \oplus H_4(\mathbf{T}'), \\ \mathbf{K}_2 &= H_2(\mathbf{T}), & \mathbf{K}'_2 &= H_2(\mathbf{T}'), \\ \mathbf{t} &= H_3(\mathbf{K} \parallel w), & \mathbf{t}' &= H_3(\mathbf{K}' \parallel w'). \end{aligned}$$

The random choices of H_1, H_2, H_3, H_4, r and \mathbf{K} determine all these random variables. Let BAD be the event that $\text{Test}(t_{w'}, (\mathbf{C}, \mathbf{c}, \mathbf{t}, \mathbf{K}_2)) = 1$. Let BIG be the event that either w or w' has length greater than or equal to $f(k)$. Then

$$\begin{aligned} \text{Adv}_{\text{PEKS}, \mathcal{S}, \mathcal{U}}^{\text{peks-consist}}(k) &= \Pr[\text{BAD}] \leq \Pr[\text{BAD} \wedge \text{BIG}] + \Pr[\text{BAD} \wedge \neg \text{BIG}] \\ &\leq \Pr[\text{BAD} \mid \text{BIG}] + \Pr[\text{BAD} \wedge \neg \text{BIG}]. \end{aligned}$$

Suppose BIG holds. If $|w'| \geq f(k)$ then $\text{Test}(t_{w'}, \mathbf{C})$ will return 1 only if $\mathbf{C} = w'$. But this will not be the case because either $|w| \geq f(k)$ and $\mathbf{C} = w \neq w'$, or $|w| < f(k)$ and, for large enough k , $|\mathbf{C}| < f(k) \leq |w'|$. On the other hand if $|w| \geq f(k)$ and $|w'| < f(k)$ then $\mathbf{C} = w$ and the latter cannot be parsed as an appropriate 4-tuple (rP, c, t, K_2) , so Test will return 0. We conclude that $\Pr[\text{BAD} \mid \text{BIG}] = 0$ for all large enough k . We now want to bound

$$\Pr[\text{BAD} \wedge \neg \text{BIG}] = \underbrace{\Pr[\text{BAD} \wedge \neg \text{BIG} \wedge H_1(w) \neq H_1(w')]}_{p_1} + \underbrace{\Pr[\text{BAD} \wedge \neg \text{BIG} \wedge H_1(w) = H_1(w')]}_{p_2}.$$

We bound p_1, p_2 in turn. We let S be the set of all distinct pairs (g, g') of elements in \mathbb{G}_1 . So p_1 is at most the sum, over all $(g, g') \in S$, of the product terms

$$\Pr[H_2(e(rP, g)) = H_2(e(rP, g')) \mid (H_1(w), H_1(w')) = (g, g')] \times \Pr[(H_1(w), H_1(w')) = (g, g')].$$

Properties of pairings tell us that $g \neq g'$ implies $e(rP, g) \neq e(rP, g')$. So due to the randomness of H_2 , the first term of each product above is 2^{-3k} . However, there are at most p^2 choices for the pair (g, g') , and we know that $p < 2^{k+1}$. Thus we have

$$p_1 \leq p^2 \cdot 2^{-3k} \leq 2^{2k+2-3k} = 4 \cdot 2^{-k}.$$

(As we discussed above, the intuition here is that with probability at least $1 - 4 \cdot 2^{-k}$ the function H_2 is injective.) We now proceed to bound p_2 . In this argument, we regard H_1 as fixed. (Formally, imagine that we condition on a particular choice of H_1 . This suffices since what follows holds for all values of this choice.) Let U be the set of all pairs (V, V') of distinct keywords of length at most $f(k)$ each such that $H_1(V) = H_1(V')$. For any $(V, V') \in U$ we let

$$\text{Keys}(V, V') = \{A \in \{0, 1\}^k : H_3(A \parallel V) = H_3(A \parallel V')\}.$$

We say that (V, V') is *heavy* if $|\text{Keys}(V, V')| \geq 2^{k/2}$, and *light* otherwise. We let $\text{Lt}(V, V')$ denote the event that (V, V') is light and $\text{Hw}(V, V')$ the event that (V, V') is heavy, where the probability is over the choice of H_3 only. Then $p_2 \leq p_L + p_H$ where

$$p_L = \sum_{(V, V') \in U} \Pr[\text{BAD} \wedge (w, w') = (V, V') \wedge \text{Lt}(V, V')],$$

$$p_H = \sum_{(V, V') \in U} \Pr[\text{BAD} \wedge (w, w') = (V, V') \wedge \text{Hw}(V, V')].$$

We bound these in turn. We have

$$p_L = \sum_{(V, V') \in U} \Pr[\text{BAD} \mid (w, w') = (V, V') \wedge \text{Lt}(V, V')]$$

$$\begin{aligned}
 & \times \Pr[(w, w') = (V, V') \wedge \text{Lt}(V, V')] \\
 & \leq \sum_{(V, V') \in U} \frac{2^{k/2}}{2^k} \cdot \Pr[(w, w') = (V, V') \wedge \text{Lt}(V, V')] \\
 & = 2^{-k/2} \cdot \sum_{(V, V') \in U} \Pr[(w, w') = (V, V') \wedge \text{Lt}(V, V')] \\
 & \leq 2^{-k/2}.
 \end{aligned} \tag{6}$$

Equation (6) is justified by the definition of the **Test**, the fact that **K** is chosen at random from $\{0, 1\}^k$ and the fact that (V, V') is light. Now we turn to bounding p_H . \square

Claim. For any $(V, V') \in U$,

$$\Pr[\text{Hw}(V, V')] \leq O(2^{-2^{k/2}}),$$

where the probability is only over the choice of H_3 .

Note the bound of the claim is double-exponentially small. We prove the claim later. Using it we can conclude via the union bound:

$$\begin{aligned}
 p_H &= \sum_{(V, V') \in U} \Pr[\text{BAD} \wedge (w, w') = (V, V') \wedge \text{Hw}(V, V')] \\
 &\leq \sum_{(V, V') \in U} \Pr[\text{Hw}(V, V')] \leq 2^{2+2f(k)} \cdot O(2^{-2^{k/2}}) \leq O(2^{-g(k)}),
 \end{aligned}$$

where $g(k) = 2^{k/2} - 2 - 2f(k) = \Omega(2^{k/2})$. So certainly $2^{-g(k)}$ is $O(2^{-k})$.

Proof of Claim. We use an occupancy problem approach:

$$\begin{aligned}
 \Pr[\text{Hw}(V, V')] &= \sum_{i=2^{k/2}}^{2^k} \binom{2^k}{i} \cdot (2^{-k})^i \cdot (1 - 2^{-k})^{2^k - i} \leq \sum_{i=2^{k/2}}^{2^k} \binom{2^k}{i} \cdot (2^{-k})^i \\
 &\leq \sum_{i=2^{k/2}}^{2^k} \left(\frac{2^k \cdot e}{i}\right)^i \cdot (2^{-k})^i \leq \sum_{i=2^{k/2}}^{2^k} \left(\frac{e}{i}\right)^i \leq \sum_{i=2^{k/2}}^{\infty} \left(\frac{e}{i}\right)^i.
 \end{aligned}$$

Let $x = e2^{-k/2}$. For $k \geq 6$, we have $x \leq 1/2$. So the above is at most

$$\sum_{i=2^{k/2}}^{\infty} x^i = x^{2^{k/2}} \cdot \sum_{i=0}^{\infty} x^i = x^{2^{k/2}} \frac{1}{1-x} \leq \frac{2}{2^{2^{k/2}}},$$

as desired. \square

4. PEKS and Anonymous IBE

We formally define anonymity of IBE schemes and investigate the relation between PEKS and anonymous IBE.

4.1. Definitions

IBE Schemes An *identity-based encryption* (IBE) scheme [9,23] $IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ consists of four PTAs. Via $(pk, msk) \xleftarrow{\$} \text{Setup}(1^k)$ the master generates master keys for security parameter $k \in \mathbb{N}$; via $usk[id] \xleftarrow{\$} \text{KeyDer}^H(msk, id)$ the master computes the secret key for identity id ; via $C \xleftarrow{\$} \text{Enc}^H(pk, id, M)$ a sender encrypts a message M to identity id to get a ciphertext; via $M \leftarrow \text{Dec}^H(usk, C)$ the possessor of secret key usk decrypts ciphertext C to get back a message. Here H is a random oracle with domain and range possibly depending on k and pk . Associated to the scheme is a message space MsgSp obeying the conventions discussed in Sect. 2. For consistency, we require that for all $k \in \mathbb{N}$, all identities id and messages $M \in \text{MsgSp}(k)$ we have $\Pr[\text{Dec}^H(\text{KeyDer}^H(msk, id), \text{Enc}^H(pk, id, M)) = M] = 1$, where the probability is taken over the choice of $(pk, msk) \xleftarrow{\$} \text{Setup}(1^k)$, the random choice of H , and the coins of all the algorithms in the expression above.

Privacy and Anonymity Privacy (IBE-IND-CPA) follows [9] while anonymity (IBE-ANO-CPA) is a straightforward adaptation of [3] to IBE schemes. Let $IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ be an IBE scheme with associated message space MsgSp . To an adversary \mathcal{A} and bit $b \in \{0, 1\}$, we associate the following experiments:

<p>Experiment $\text{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ind-cpa-b}}(k)$</p> <p>$IDSet \leftarrow \emptyset$; $(pk, msk) \xleftarrow{\\$} \text{Setup}(1^k)$ pick random oracle H $(id, M_0, M_1, state)$ $\xleftarrow{\\$} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{find}, pk)$ if $\{M_0, M_1\} \not\subseteq \text{MsgSp}(k)$ then return 0 $C \xleftarrow{\\$} \text{Enc}^H(pk, id, M_b)$ $b' \xleftarrow{\\$} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{guess}, C, state)$ if $id \notin IDSet$ and $M_0 = M_1$ then return b' else return 0</p>	<p>Experiment $\text{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-cpa-b}}(k)$</p> <p>$IDSet \leftarrow \emptyset$; $(pk, msk) \xleftarrow{\\$} \text{Setup}(1^k)$ pick random oracle H $(id_0, id_1, M, state)$ $\xleftarrow{\\$} \mathcal{A}^{\text{KEYDER}, H}(\text{find}, pk)$ if $M \notin \text{MsgSp}(k)$ then return 0 $C \xleftarrow{\\$} \text{Enc}^H(pk, id_b, M)$ $b' \xleftarrow{\\$} \mathcal{A}^{\text{KEYDER}, H}(\text{guess}, C, state)$ if $\{id_0, id_1\} \cap IDSet = \emptyset$ then return b' else return 0</p>
---	---

where the oracle $\text{KEYDER}(id)$ is defined as

$$IDSet \leftarrow IDSet \cup \{id\}; usk[id] \xleftarrow{\$} \text{KeyDer}^H(msk, id); \text{Return } usk[id].$$

For $\text{prop} \in \{\text{ind}, \text{ano}\}$, we define the advantage of \mathcal{A} in the corresponding experiment as

$$\text{Adv}_{IBE, \mathcal{A}}^{\text{ibe-prop-cpa}}(k) = \Pr[\text{Exp}_{IBE, \mathcal{A}}^{\text{ibe-prop-cpa-1}}(k) = 1] - \Pr[\text{Exp}_{IBE, \mathcal{A}}^{\text{ibe-prop-cpa-0}}(k) = 1].$$

IBE scheme IBE is said to be *IBE-IND-CPA-secure* (resp., *IBE-ANO-CPA-secure*) if the respective advantage function is negligible for all PTAs \mathcal{A} .

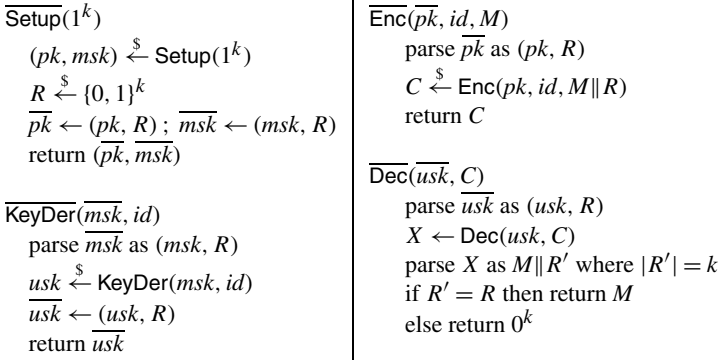


Fig. 4. IBE scheme for proof of Theorem 4.1.

4.2. The ibe-2-peks Transform

The ibe-2-peks transform suggested in [8] takes input an IBE scheme $IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ and returns a PEKS scheme $\mathcal{PEKS} = (\text{KG}, \text{Td}, \text{PEKS}, \text{Test})$ as follows. The public key pk and secret key sk of the receiver in the PEKS scheme are the master public and secret keys, respectively, of the IBE scheme (i.e., $\text{KG} = \text{Setup}$). The trapdoor t_w associated to keyword w is the secret key that the IBE scheme would assign to the identity w (i.e., $\text{Td}(sk, w) = \text{KeyDer}(sk, w)$). A keyword w is PEKS-encrypted by IBE-encrypting the message 0^k for the identity w (i.e., $\text{PEKS}(pk, w) = \text{Enc}(pk, w, 0^k)$). Finally, testing is done by checking that the ciphertext decrypts to 0^k (i.e., $\text{Test}(t_w, C)$ returns 1 iff $\text{Dec}(t_w, C) = 0^k$).

We know that $\mathcal{BF}\text{-IBE}$ is anonymous (Theorem 4.4), that $\mathcal{BDOP}\text{-PEKS} = \text{ibe-2-peks}(\mathcal{BF}\text{-IBE})$, and that $\mathcal{BDOP}\text{-PEKS}$ is not statistically consistent (Proposition 3.2). Thus, we can conclude that the ibe-2-peks transform does not necessarily yield a statistically consistent PEKS scheme. Unfortunately, as the following theorem shows, the ibe-2-peks transform does not necessarily yield a computationally consistent PEKS scheme either (under the minimal assumption of the existence of some IBE-IND-CPA- and IBE-ANO-CPA-secure IBE scheme). As a result, ibe-2-peks is not in general a suitable way to obtain a PEKS scheme.

Theorem 4.1. *Assume there exist IBE-ANO-CPA-secure and IBE-IND-CPA-secure IBE schemes. Then there exists a IBE-ANO-CPA-secure and IBE-IND-CPA-secure IBE scheme \overline{IBE} such that the PEKS scheme \mathcal{PEKS} derived from \overline{IBE} via ibe-2-peks is not computationally consistent.*

Proof (Sketch). The proof of Theorem 4.1 is quite simple and its details are omitted here. Instead, we only provide the general intuition behind it. In order to show that ibe-2-peks does not necessarily yield a computationally consistent PEKS scheme, we first assume the existence of a IBE-IND-CPA- and IBE-ANO-CPA-secure IBE scheme $IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ and then build an IBE scheme $\overline{IBE} = (\overline{\text{Setup}}, \overline{\text{KeyDer}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ as shown in Fig. 4. It is easy to see that the IBE-IND-CPA-

and IBE-ANO-CPA-security of \overline{IBE} follows from simple reductions from the security of IBE . Now, let \mathcal{PEKS} denote the PEKS scheme outputted by ibe-2-peks on input \overline{IBE} . Clearly, \mathcal{PEKS} is not computationally consistent as its test algorithm outputs 1 with overwhelming probability, when given the trapdoor for the wrong keyword. The only case in which it outputs 0 when given the wrong trapdoor is when the last k bits of the decryption of the ciphertext C with the wrong trapdoor matches the random value R in the public key \overline{pk} , but this only happens with negligible probability due to the IBE-IND-CPA security of the IBE scheme IBE . \square

4.3. The new-ibe-2-peks Transform

The negative result in Theorem 4.1 raises the question: Does the existence of IBE schemes imply the existence of computationally consistent PEKS schemes? We answer that in the affirmative by presenting a revision of the ibe-2-peks transform, called new-ibe-2-peks , that transforms any IBE-IND-CPA- and IBE-ANO-CPA-secure IBE scheme into a PEKS-IND-CPA-secure and computationally consistent PEKS scheme. It is similar to ibe-2-peks except that instead of always using 0^k as the message encrypted, the PEKS-encryption algorithm chooses and encrypts a random message R and appends R in the clear to the ciphertext. In more detail, the new-ibe-2-peks transform takes input an IBE scheme $IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ and returns a PEKS scheme $\mathcal{PEKS} = (\text{KG}, \text{Td}, \text{PEKS}, \text{Test})$ as follows. The public key pk and secret key sk of the receiver in the PEKS scheme are the master public and secret keys, respectively, of the IBE scheme. (I.e. $\text{KG} = \text{Setup}$.) The trapdoor associated to keyword w is the secret key that the IBE scheme would assign to the identity w . (I.e. $\text{Td}(sk, w) = \text{KeyDer}(sk, w)$.) PEKS-encryption of keyword w is done as follows: $\text{PEKS}(pk, w)$ picks $R \xleftarrow{\$} \{0, 1\}^k$, lets $C \xleftarrow{\$} \text{Enc}(pk, w, R)$, and returns (C, R) as the ciphertext. Finally, $\text{Test}(t_w, (C, R))$ returns 1 iff $\text{Dec}(t_w, C) = R$.

Intuitively, this construction avoids the problem of oddly-behaving Dec algorithms by making sure that the *only* way to ruin the consistency of the PEKS scheme is by correctly guessing the value encrypted by a ciphertext, using the secret key of a different identity, which should not be possible for an IBE-IND-CPA-secure IBE scheme. Hence, the consistency of the resulting PEKS scheme is due to the data privacy property of the IBE scheme, while the data privacy property of the PEKS scheme comes from the anonymity of the IBE scheme. The formal result statement and proof follow.

Theorem 4.2. *Let IBE be an IBE scheme and let \mathcal{PEKS} be the PEKS scheme derived from IBE via new-ibe-2-peks . If IBE is IBE-IND-CPA-secure, then \mathcal{PEKS} is computationally consistent. Further, if IBE is IBE-ANO-CPA-secure, then \mathcal{PEKS} is PEKS-IND-CPA-secure.*

Proof. Let \mathcal{U} be any PTA attacking the computational consistency of \mathcal{PEKS} , and consider the following PTA \mathcal{A} attacking the IBE-IND-CPA-security of IBE . In its `find` stage, given master public key pk , adversary \mathcal{A} runs $\mathcal{U}(pk)$ to get keywords w, w' . It returns w as the challenge identity and $R_0, R_1 \xleftarrow{\$} \{0, 1\}^k$ as the challenge messages. In the `guess` stage, given challenge ciphertext C (that encrypts R_b under identity w for

challenge bit $b \in \{0, 1\}$, \mathcal{A} uses its key-derivation oracle to obtain a trapdoor $t_{w'}$ for w' . If $\text{Dec}(t_{w'}, C) = R_1$ then it returns 1 else it returns 0. It is easy to see that

$$\begin{aligned} \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ind-cpa-1}}(k) = 1] &\geq \Pr[\mathbf{Exp}_{PEKS, \mathcal{U}}^{\text{peks-consist}}(k) = 1], \\ \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ind-cpa-0}}(k) = 1] &\leq 2^{-k}. \end{aligned}$$

Thus $\mathbf{Adv}_{PEKS, \mathcal{U}}^{\text{peks-consist}}(k) \leq \mathbf{Adv}_{IBE, \mathcal{A}}^{\text{ibe-ind-cpa}}(k) + 2^{-k}$, proving the first claim of the theorem.

Let \mathcal{B} be any PTA attacking the PEKS-IND-CPA-security of $PEKS$, and consider the following PTA \mathcal{A} attacking the IBE-ANO-CPA-security of IBE . In its `find` stage, given master public key pk , adversary \mathcal{A} runs $\mathcal{B}(\text{find}, pk)$ to get challenge keywords w_0, w_1 , which it returns along with a message $R \xleftarrow{\$} \{0, 1\}^k$. In the `guess` stage, given challenge ciphertext C (that encrypts R under identity w_b for challenge bit $b \in \{0, 1\}$), \mathcal{A} runs \mathcal{B} , in its `guess` stage, with challenge ciphertext (C, R) , to get its guess bit b' , which \mathcal{A} returns. In both stages, \mathcal{A} answers any trapdoor-oracle queries of \mathcal{B} via its key-derivation oracle. It is easy to see that for $b = 0, 1$,

$$\Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-cpa-b}}(k) = 1] = \Pr[\mathbf{Exp}_{PEKS, \mathcal{B}}^{\text{peks-ind-cpa-b}}(k) = 1].$$

Thus $\mathbf{Adv}_{PEKS, \mathcal{B}}^{\text{peks-ind-cpa}}(k) \leq \mathbf{Adv}_{IBE, \mathcal{A}}^{\text{ibe-ano-cpa}}(k)$, proving the second claim of the theorem. \square

4.4. A Sufficient Condition for Anonymity

Halevi [19] provides a simple sufficient condition for an IND-CPA public-key encryption scheme to meet the notion of anonymity (a.k.a. key-privacy) of [3]. The condition is that even a computationally unbounded adversary, given public keys pk_0, pk_1 and the encryption of a random message under pk_b , have only a negligible advantage in determining the random challenge bit b . Towards finding anonymous IBE schemes (a task motivated by Theorem 4.2) we extend Halevi's condition to identity-based encryption. In the process we also extend it in two other ways: first to handle the random oracle model (the standard model is a special case) and second to weaken the statistical (i.e. information-theoretic) requirement of [19] to a computational one. (The application of this paper does not need the last extension, but it may be useful in other contexts.)

We begin by defining a relevant (new) notion of security that we call IBE-ANO-RE-CPA. Let $IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ be an IBE scheme with associated message space MsgSp . We associate to an adversary \mathcal{A} and bit $b \in \{0, 1\}$ the following experiment:

Experiment $\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-b}}(k)$

$IDSet \leftarrow \emptyset$; $(pk, msk) \xleftarrow{\$} \text{Setup}(1^k)$
 pick random oracle H
 $(id_0, id_1, M, state) \xleftarrow{\$} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{find}, pk)$
 if $M \notin \text{MsgSp}(k)$ then return 0
 $R \xleftarrow{\$} \{0, 1\}^{|M|}$; $C \xleftarrow{\$} \text{Enc}^H(pk, id_b, R)$
 $b' \xleftarrow{\$} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{guess}, C, state)$
 if $\{id_0, id_1\} \cap IDSet = \emptyset$ then return b'
 else return 0

Oracle $\text{KEYDER}(id)$

$IDSet \leftarrow IDSet \cup \{id\}$
 $usk[id] \xleftarrow{\$} \text{KeyDer}^H(msk, id)$
 return $usk[id]$

The IBE-ANO-RE-CPA-*advantage* of an adversary \mathcal{A} in violating the anonymity of the scheme IBE is defined as

$$\mathbf{Adv}_{IBE, \mathcal{A}}^{\text{ibe-ano-re}}(k) = \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-1}}(k) = 1] - \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-0}}(k) = 1].$$

A scheme IBE is said to be IBE-ANO-RE-CPA-*secure* if the above advantage is a negligible function in k for all PTAs \mathcal{A} .

Lemma 4.3. *Let IBE be an IBE scheme that is IBE-IND-CPA and IBE-ANO-RE-CPA-secure. Then it is also IBE-ANO-CPA-secure.*

Proof of Lemma 4.3. The proof is a simple hybrid argument. Let \mathcal{A} be a PTA attacking the IBE-ANO-CPA-security of IBE . It is easy to construct PTAs $\mathcal{A}_1, \mathcal{A}_3$ attacking the IBE-IND-CPA-security of IBE , and PTA \mathcal{A}_2 attacking the IBE-ANO-RE-CPA-security of IBE , such that

$$\begin{aligned} \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-cpa-1}}(k) = 1] - \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-1}}(k) = 1] &\leq \mathbf{Adv}_{IBE, \mathcal{A}_1}^{\text{ibe-ind-cpa}}(k), \\ \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-1}}(k) = 1] - \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-0}}(k) = 1] &\leq \mathbf{Adv}_{IBE, \mathcal{A}_2}^{\text{ibe-ano-re}}(k), \\ \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-re-0}}(k) = 1] - \Pr[\mathbf{Exp}_{IBE, \mathcal{A}}^{\text{ibe-ano-cpa-0}}(k) = 1] &\leq \mathbf{Adv}_{IBE, \mathcal{A}_3}^{\text{ibe-ind-cpa}}(k). \end{aligned}$$

Summing concludes the proof. We omit the details, save to remark that we use here the second convention about message spaces noted in Sect. 2. \square

4.5. Anonymity of \mathcal{BF} - IBE

The Boneh–Franklin BasicIdent IBE scheme [9] is shown in Fig. 5. We apply Lemma 4.3 to give a simple proof that it is IBE-ANO-CPA.

Theorem 4.4. *The \mathcal{BF} - IBE scheme is IBE-ANO-CPA-secure assuming that the BDH is hard relative to generator \mathcal{G} .*

Proof. Given Lemma 4.3, and given that the \mathcal{BF} - IBE scheme is IBE-IND-CPA-secure [9], it suffices to show that the scheme is IBE-ANO-RE-CPA-secure. Notice that the ciphertext C in Fig. 5 has two parts, namely $U = rP$ and $V = M \oplus H_2(T)$. The value

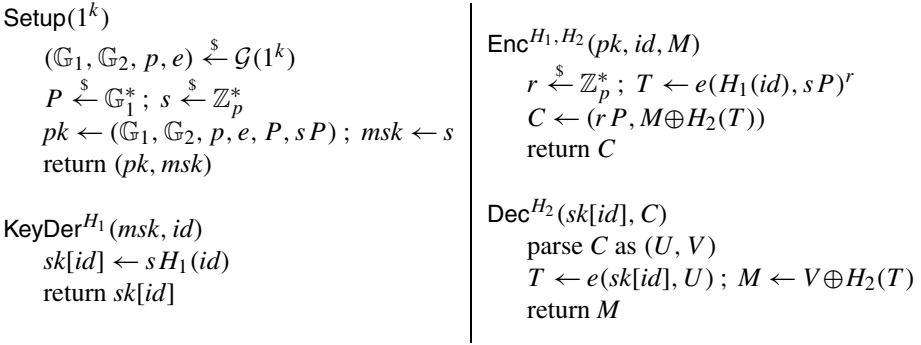


Fig. 5. Algorithms of the IBE scheme $\mathcal{BF}\text{-IBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$. Here \mathcal{G} is a pairing parameter generator and $H_1 : \{0, 1\}^* \rightarrow \mathbb{G}_1^*$ and $H_2 : \mathbb{G}_2 \rightarrow \{0, 1\}^k$ are random oracles. The message space is defined by $\text{MsgSp}(k) = \{0, 1\}^k$ for all $k \in \mathbb{N}$.

U is chosen uniformly at random from \mathbb{G}_1^* by the encryption algorithm. If the message M is chosen uniformly at random from $\{0, 1\}^k$, then V is also uniformly distributed in $\{0, 1\}^k$ and independent of the $H_2(T)$. Thus in both the 0- and 1-worlds of the IBE-ANO-RE-CPA-security game, the challenge ciphertext C has exactly the same distribution. Therefore any adversary against IBE-ANO-RE-CPA-security will have 0 advantage. □

5. Anonymous HIBE

5.1. Definitions

HIBE Schemes A *hierarchical identity-based encryption* (HIBE) scheme [7,15,20] is a generalization of an IBE scheme in which an identity is a vector of strings $id = (id_1, \dots, id_l)$ with the understanding that when $l = 0$ this is the empty vector $()$. The number of components in this vector is called the level of the identity and is denoted $|id|$. If $0 \leq i \leq l$ then $id|_i = (id_1, \dots, id_i)$ denotes the vector containing the first i components of id . If $|id'| \geq l + 1$ ($l \geq 0$) and $id'|_l = id$ then we say that id is an ancestor of id' , or equivalently, that id' is a descendant of id . If the level of id' is $l + 1$ then id is a parent of id' , or, equivalently, id' is a child of id . For any id with $|id| \geq 1$ we let $\text{par}(id) = id|_{|id|-1}$ denote its parent. Two nodes $id = (id_1, \dots, id_l)$ and $id' = (id'_1, \dots, id'_l)$ at level l are said to be siblings iff $id|_{l-1} = id'|_{l-1}$. Moreover, if $id_1 < id'_1$ in lexicographic order, then id is a left sibling of id' and id' is a right sibling of id . An identity at level one or more can be issued a secret key by its parent. (And thus an identity can issue keys for any of its descendants if necessary.)

Formally a HIBE scheme $\mathcal{HIBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ consists of four PTAs. Via $(pk, msk = usk[()]) \xleftarrow{\$} \text{Setup}(1^k)$, where $k \in \mathbb{N}$ is a security parameter, the root generates master keys, with the secret key being associated to the (unique) identity $()$ at level 0. Via $usk[id] \xleftarrow{\$} \text{KeyDer}^H(usk[\text{par}(id)], id)$ the parent of an identity id with $|id| \geq 1$ can compute a secret key for id . Note that by iteratively applying the KeyDer

algorithm a user id can derive secret keys for any of its descendants id' ; we occasionally use the notation $usk[id'] \stackrel{\$}{\leftarrow} \text{KeyDer}^H(usk[id], id')$ to denote this process. Via $C \stackrel{\$}{\leftarrow} \text{Enc}^H(pk, id, M)$ a sender encrypts a message M to identity id to get a ciphertext; via $M \leftarrow \text{Dec}^H(usk[id], C)$ the identity id decrypts ciphertext C to get back a message. Here H is a random oracle with domain and range possibly depending on k and pk . Associated to the scheme is a message space MsgSp obeying the conventions discussed in Sect. 2. For consistency, we require that for all $k \in \mathbb{N}$, all identities id with $|id| \geq 1$ and all messages $M \in \text{MsgSp}(k)$,

$$\Pr[\text{Dec}^H(\text{KeyDer}^H(usk[\text{par}(id)], id), \text{Enc}^H(pk, id, M)) = M] = 1,$$

where the probability is taken over the choice of $(pk, usk[()]) \stackrel{\$}{\leftarrow} \text{Setup}(1^k)$, the random choice of H , and the coins of all the algorithms in the expression above.

Privacy and Anonymity The notion of privacy for HIBE schemes is analogous to that for IBE schemes (IBE-IND-CPA) but using identity vectors rather than identity strings and where the adversary is not allowed to query the KEYDER oracle for the secret key of any ancestor of the identity under attack. Since we will deal with schemes where privacy holds only up to some level, the notion is parameterized by a maximum depth function $d : \mathbb{N} \rightarrow \mathbb{N}$, and all identities id (in queries or challenges) must have $|id| \leq d(k)$. To allow a fine-grained treatment of anonymity we introduce the concept of anonymity at a set $L(k)$ of levels, meaning that in an experiment the adversary \mathcal{A} is challenged to distinguish two distinct identities differing only at levels $l \in L(k)$. (Here for each k , $L(k)$ is a finite set of integers. For ease of notation, we will write l rather than $\{l\}$ when $L(k) = \{l\}$ is a singleton set.)

Formally, let $\mathcal{HIBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ be an identity-based encryption scheme with message space MsgSp , let $d : \mathbb{N} \rightarrow \mathbb{N}$ be the maximum depth, and let L be a set of levels. Let $\text{diff}(\cdot, \cdot)$ be the function that returns the set of coordinates at which the input identities differ, and $\text{anc}(\cdot)$ the function returning the set of ancestors of the input identity. To any bit $b \in \{0, 1\}$ and any adversary \mathcal{A} , we associate the experiments:

Experiment $\text{Exp}_{\mathcal{HIBE}, \mathcal{A}}^{\text{hibe-ind-cpa-b}[d]}(k)$

$IDSet \leftarrow \emptyset$; $(pk, msk) \stackrel{\$}{\leftarrow} \text{Setup}(1^k)$
 pick random oracle H
 $(id, M_0, M_1, state)$
 $\stackrel{\$}{\leftarrow} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{find}, pk)$
 if $|M_0| \neq |M_1|$ or $|id| > d(k)$
 or $\{M_0, M_1\} \not\subseteq \text{MsgSp}(k)$
 then return 0
 $C \stackrel{\$}{\leftarrow} \text{Enc}^H(pk, id, M_b)$
 $b' \stackrel{\$}{\leftarrow} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{guess}, C, state)$
 if $IDSet \cap \text{anc}(id) = \emptyset$
 then return b' else return 0

Experiment $\text{Exp}_{\mathcal{HIBE}, \mathcal{A}}^{\text{hibe-ano-cpa-b}[L, d]}(k)$

$IDSet \leftarrow \emptyset$; $(pk, msk) \stackrel{\$}{\leftarrow} \text{Setup}(1^k)$
 pick random oracle H
 $(id_0, id_1, M, state)$
 $\stackrel{\$}{\leftarrow} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{find}, pk)$
 if $|id_0| \neq |id_1|$ or $|id_0| > d(k)$
 or $|id_1| > d(k)$ or $M \notin \text{MsgSp}(k)$
 then return 0
 $C \stackrel{\$}{\leftarrow} \text{Enc}(pk, id_b, M)$
 $b' \stackrel{\$}{\leftarrow} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{guess}, C, state)$
 if $IDSet \cap (\text{anc}(id_0) \cup \text{anc}(id_1)) = \emptyset$
 and $\text{diff}(id_0, id_1) \subseteq L(k)$
 then return b' else return 0

where the oracle $\text{KEYDER}(\cdot)$ is defined as

if $|id| > d(k)$ then return \perp ; $IDSet \leftarrow IDSet \cup \{id\}$; return $\text{KeyDer}(msk, id)$.

We define the advantage of \mathcal{A} in the corresponding experiments as

$$\begin{aligned} \mathbf{Adv}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ind-cpa}[d]}(k) &= \Pr[\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ind-cpa-1}[d]}(k) = 1] \\ &\quad - \Pr[\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ind-cpa-0}[d]}(k) = 1] \\ \mathbf{Adv}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-cpa}[L, d]}(k) &= \Pr[\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-cpa-1}[L, d]}(k) = 1] \\ &\quad - \Pr[\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-cpa-0}[L, d]}(k) = 1]. \end{aligned}$$

The scheme $\mathcal{H}IBE$ is said to be $\text{HIBE-IND-CPA}[d]$ -secure (resp. $\text{HIBE-ANO-CPA}[L, d]$ -secure) if the respective advantage function is negligible for all PTAs \mathcal{A} .

5.2. A Sufficient Condition for Anonymity

We further extend Lemma 4.3 to the hierarchical case. To this end, we introduce a new notion $\text{HIBE-ANO-RE-CPA}[L, d]$ as follows. Let $\mathcal{H}IBE = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ be a HIBE scheme with message space MsgSp , let L be a set of levels, and let d be the maximum hierarchy depth. To an adversary \mathcal{A} and a bit b , we associate the following experiment:

<p>Experiment $\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-re-b}[L, d]}(k)$</p> <p>$IDSet \leftarrow \emptyset$; $(pk, msk) \xleftarrow{\\$} \text{Setup}(1^k)$</p> <p>pick random oracle H</p> <p>$(id_0, id_1, M, state) \xleftarrow{\\$} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{find}, pk)$</p> <p>if $id_0 \neq id_1$ or $id_0 > d(k)$ or $id_1 > d(k)$ or $M \notin \text{MsgSp}(k)$ then return 0</p> <p>$R \xleftarrow{\\$} \{0, 1\}^{ M }$; $C \xleftarrow{\\$} \text{Enc}^H(pk, id_b, R)$</p> <p>$b' \xleftarrow{\\$} \mathcal{A}^{\text{KEYDER}(\cdot), H}(\text{guess}, C, state)$</p> <p>if $IDSet \cap (\{id_0, id_1\} \cup \text{anc}(id_0) \cup \text{anc}(id_1)) = \emptyset$ and $\text{diff}(id_0, id_1) \subseteq L(k)$ then return b' else return 0</p>	<p>Oracle $\text{KEYDER}(id)$</p> <p>if $id > d(k)$ then return \perp</p> <p>$IDSet \leftarrow IDSet \cup \{id\}$</p> <p>return $\text{KeyDer}^H(msk, id)$</p>
--	--

The $\text{HIBE-ANO-RE-CPA}[L, d]$ -advantage of an adversary \mathcal{A} in violating the level- L anonymity of the scheme $\mathcal{H}IBE$ with depth $d(k)$ is defined as

$$\begin{aligned} \mathbf{Adv}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-re}[L, d]}(k) &= \Pr[\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-re-1}[L, d]}(k) = 1] \\ &\quad - \Pr[\mathbf{Exp}_{\mathcal{H}IBE, \mathcal{A}}^{\text{hibe-ano-re-0}[L, d]}(k) = 1]. \end{aligned}$$

A scheme $\mathcal{H}IBE$ is said to be $\text{HIBE-ANO-RE-CPA}[L, d]$ -secure if this advantage is a negligible function in k for all PTAs \mathcal{A} . The following lemma follows from a hybrid argument similar to that of Lemma 4.3.

Lemma 5.1. *Let \mathcal{HIBE} be a HIBE scheme that is HIBE-IND-CPA[d] and HIBE-ANO-RE-CPA[L, d]-secure for some set of levels L and hierarchy depth d . Then \mathcal{HIBE} is also HIBE-ANO-CPA[L, d]-secure.*

5.3. Construction

The HIBE scheme of [20] appears to be anonymous, but supports only two levels of identities, and is only resistant against limited collusions at the second level, and hence is not usable for our constructions that follow. Since the HIBE of [15], here denoted $\mathcal{GS}\text{-HIBE}$, is equivalent to the Boneh–Franklin IBE scheme [9] when restricted to the first level, and since the latter is provably anonymous as per Theorem 4.4, one could hope that $\mathcal{GS}\text{-HIBE}$ is level-1 anonymous, but this turns out not to be true, and the HIBE of [7] is not level-1 anonymous either. To see why, consider the following. The $\mathcal{GS}\text{-HIBE}$ encryption of a message M under identity $id = (id_1, \dots, id_l)$ is a tuple

$$(rP, rH_1(id|_2), \dots, rH_1(id|_l), H_2(e(rP, H_1(id_1))) \oplus m) \tag{7}$$

where H_1, H_2 are random oracles, P is a generator of a pairing group that is part of pk , and r is chosen at random from \mathbb{Z}_p by the encryption algorithm. Anonymity is violated because an adversary can decide whether a given ciphertext (C_1, C_2, C_3) is intended for $id = (id_1, id_2)$ or $id' = (id'_1, id_2)$ by checking whether $e(C_2, P)$ equals $e(C_1, H_1(id))$ or $e(C_1, H_1(id'))$.

The lack of anonymity in $\mathcal{GS}\text{-HIBE}$ stems from the fact that the hashes in the first l components of the ciphertext depend on the first component of the recipient’s identity. In Fig. 6, we present a modified $m\mathcal{GS}\text{-HIBE}$ scheme that uses a different random oracle $H_{1,l}$ for each level l , and that computes ciphertexts as

$$(rP, rH_{1,2}(id_2), \dots, rH_{1,l}(id_l), H_2(e(rP, H_{1,1}(id_1))) \oplus m).$$

The following implies in particular that $m\mathcal{GS}\text{-HIBE}$ is the first full HIBE scheme providing anonymity at any level. The restriction on d is inherited from [15]. We note that, subsequently to our work, Boyen and Waters [11] proposed a HIBE scheme that is anonymous at all levels in the standard (i.e., non-random-oracle) model.

Theorem 5.2. *For any $d(k) = O(\log(k))$, the $m\mathcal{GS}\text{-HIBE}$ scheme is HIBE-ANO-CPA[1, d]-secure and HIBE-IND-CPA[d]-secure in the random oracle model assuming the BDH problem is hard relative to the generator \mathcal{G} .*

We split up the proof in the following two lemmas. The proof of the first is given in Appendix B, and recycles ideas from [9,15]. We use Lemma 5.1 to prove the second lemma.

Lemma 5.3. *For any $d(k) = O(\log(k))$, the $m\mathcal{GS}\text{-HIBE}$ scheme is HIBE-IND-CPA[d]-secure in the random oracle model assuming the BDH problem is hard relative to the generator \mathcal{G} .*

Lemma 5.4. *For any $d(k) = O(\log(k))$, the $m\mathcal{GS}\text{-HIBE}$ scheme is HIBE-ANO-CPA[1, d]-secure in the random oracle model assuming the BDH problem is hard relative to the generator \mathcal{G} .*

<pre> Setup(1^k) $(\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\\$} \mathcal{G}(1^k)$; $P \xleftarrow{\\$} \mathbb{G}_1^*$ $s_0 \xleftarrow{\\$} \mathbb{Z}_p^*$; $S_0 \leftarrow 0$; $Q_0 \leftarrow s_0 P$ $pk \leftarrow (\mathbb{G}_1, \mathbb{G}_2, p, e, P, Q_0)$ $msk \leftarrow (pk, (), S_0, s_0)$ return (pk, msk) KeyDer$^{H_{1,1}, \dots, H_{1,l}}(usk, id)$ parse id as (id_1, \dots, id_{l+1}) parse usk as $(pk, id_l, S_l, Q_1, \dots, Q_{l-1}, s_l)$ parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, Q_0)$ $S_{l+1} \leftarrow S_l + s_l H_{1,l+1}(id_{l+1})$ $Q_l \leftarrow s_l P$; $s_{l+1} \xleftarrow{\\$} \mathbb{Z}_p^*$ return $(pk, id, S_{l+1}, Q_1, \dots, Q_l, s_{l+1})$ </pre>	<pre> Enc$^{H_{1,1}, \dots, H_{1,l}, H_2}(pk, id, M)$ parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, Q_0)$ parse id as (id_1, \dots, id_l) $r \xleftarrow{\\$} \mathbb{Z}_p^*$; $C_1 \leftarrow rP$ for $i = 2, \dots, l$ do $C_i \leftarrow rH_{1,i}(id_i)$ $C_{l+1} \leftarrow M \oplus H_2(e(rH_{1,1}(id_1), Q_0))$ return (C_1, \dots, C_{l+1}) Dec$^{H_2}(usk, C)$ parse usk as $(pk, id, S_l, Q_1, \dots, Q_{l-1}, s_l)$ parse id as (id_1, \dots, id_l) parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, Q_0)$ parse C as (C_1, \dots, C_{l+1}) $\kappa \leftarrow e(S_l, C_1) \cdot \prod_{i=2}^l e(Q_{i-1}, C_i)^{-1}$ return $C_{l+1} \oplus H_2(\kappa)$ </pre>
--	---

Fig. 6. Algorithms of the $mGS\text{-}\mathcal{HIBE}$ scheme. \mathcal{G} is a pairing parameter generator and $H_{1,i}: \{0, 1\}^* \rightarrow \mathbb{G}_1^*$ and $H_2: \mathbb{G}_2 \rightarrow \{0, 1\}^k$ are random oracles.

Proof. Given Lemmas 5.1 and 5.3, it suffices to show that $mGS\text{-}\mathcal{HIBE}$ is HIBE-ANORE-CPA[1, d]-secure. In the challenge ciphertext $(C_1^*, \dots, C_{l+1}^*)$, the first component C_1 is chosen uniformly at random from \mathbb{G}_1^* . Component C_i^* for $2 \leq i \leq l$ is uniquely defined by C_1^* and the i -th component of the identity, which is the same for both challenge identities since they can only differ at level 1. Finally, if the message M is chosen uniformly at random from $\{0, 1\}^k$, then the last component C_{l+1}^* is also uniformly distributed over $\{0, 1\}^k$, independent of $H_2(e(rH_{1,1}(id_1), Q_0))$. Hence, the challenge ciphertext is identically distributed in both worlds, and the advantage of any adversary is 0. \square

6. Public-Key Encryption with Temporary Keyword Search

In a PEKS scheme, once the gateway has the trapdoor for a certain keyword, it can test whether this keyword was present in past ciphertexts, and can test its presence in any future ciphertexts. It may be useful to limit the period in which the trapdoor can be used. Here we propose an extension of PEKS that allows this. We call it public-key encryption with temporary keyword search (PETKS) or temporarily searchable encryption for short. A trapdoor here is created for a time interval $[s, e]$ and will only allow the gateway to test whether ciphertexts created in this time interval contain the keyword.

6.1. Definitions

PETKS Schemes *Public-key encryption with temporary keyword search* (PETKS) is a generalization of PEKS in which a trapdoor can be issued for any desired window of time rather than forever. Formally, the scheme $\mathcal{PETKS} = (\text{KG}, \text{Td}, \text{PETKS}, \text{Test}, N)$ consists of four PTAs and a function $N: \mathbb{N} \rightarrow \mathbb{N}$. Via $(pk, sk) \xleftarrow{\$} \text{KG}(1^k)$, the receiver generates its public and secret key; via $C \xleftarrow{\$} \text{PETKS}^H(pk, w, i)$ a sender encrypts a keyword w in time period $i \in [0, N(k) - 1]$ to get a ciphertext; via $t_w \xleftarrow{\$} \text{Td}^H(sk, w, s, e)$

the receiver computes a trapdoor t_w for keyword w in period $[s, e]$ where $0 \leq s \leq e \leq N(k) - 1$, and provides it to the gateway; via $b \leftarrow \text{Test}^H(t_w, C)$ the gateway tests whether C encrypts w , where b is a bit with 1 meaning “accept” or “yes” and 0 meaning “reject” or “no”. Here H is a random oracle whose domain and/or range might depend on k and pk . We require that for all $k \in \mathbb{N}$, all s, e, i with $0 \leq s \leq i \leq e \leq N(k) - 1$, and all $w \in \{0, 1\}^*$,

$$\Pr[\text{Test}^H(\text{Td}^H(sk, w, s, e), \text{PETKS}^H(pk, w, i)) = 1] = 1,$$

where the probability is taken over the choice of $(pk, sk) \xleftarrow{\$} \text{KG}(1^k)$, the random choice of H , and the coins of all the algorithms in the expression above.

Consistency Consistency for PETKS schemes requires that no user \mathcal{U} can output keywords w, w' and time period indices $s, e, i \in [1, N(k) - 1]$ such that $w \neq w'$ or $i \notin [s, e]$, yet still an encryption of w for time period i tests positively under a trapdoor for keyword w' and time period $[s, e]$. We define the advantage $\text{Adv}_{\text{PETKS}, \mathcal{U}}^{\text{petks-consist}}(k)$ as the probability that \mathcal{U} succeeds in doing so. Just like for standard PEKS schemes, we distinguish between perfect, statistical and computational consistency.

Privacy Privacy for a PETKS scheme asks that an adversary be unable to distinguish between the encryption of two challenge keywords of its choice in a time period $i \in [0, N(k) - 1]$ of its choice, even if it is allowed not only to obtain trapdoors for non-challenge keywords issued for any time interval, but also is allowed to obtain trapdoors for *any* keywords (even the challenge ones), issued for time intervals not containing i . The formal experiment and the definition of PETKS-IND-CPA-advantage and security are otherwise analogous to those of standard PEKS schemes, and hence are omitted here.

6.2. Constructions for PETKS Schemes

Constructions with Linear Complexity PETKS is reminiscent of forward-security [4,12], and, as in these works, there are straightforward solutions with keys or trapdoors of length linear in $N(k)$. One such solution is to use a standard PEKS scheme and generate a different key pair (pk_i, sk_i) for each time period $i \in [0, N(k) - 1]$. Let $pk = (pk_0, \dots, pk_{N(k)-1})$ be the PETKS public key and $sk = (sk_0, \dots, sk_{N(k)-1})$ be the PETKS secret key. During time period i , the sender encrypts a keyword w by encrypting w under pk_i using the PEKS scheme. The trapdoor for a keyword w in the interval $[s, e]$ consists of all PEKS trapdoors for w of periods s, \dots, e . A somewhat more efficient solution is to let the PETKS master key pair be a single key pair for the standard PEKS scheme, and append the time period to the keyword (making sure that the string is uniquely decodable, e.g. by using a special separator symbol) when encrypting or computing trapdoors. This scheme achieves short public and secret keys, but still has trapdoor length linear in $N(k)$, because the PETKS trapdoor still contains PEKS trapdoors for all time periods s, \dots, e . Note that both these construction only work for polynomially bounded $N(k)$.

The hibe-2-petks transform We now present a transformation *hibe-2-petks* of a HIBE scheme into a PETKS scheme that yields a PETKS scheme with complexity logarithmic in $N(k)$ for all parameters. The construction is very similar to the generic construction of forward-secure encryption from binary-tree encryption [12]. The number of time periods is $N(k) = 2^{t(k)}$ for some polynomially bounded function $t : \mathbb{N} \rightarrow \mathbb{N}$. If $i \in [0, N(k) - 1]$, then let $i_1, \dots, i_{t(k)}$ denote its binary representation as a $t(k)$ -bit string. Intuitively, our construction instantiates a HIBE of depth $t(k) + 1$ with keywords as the first level of the identity tree and the time structure on the lower levels. The trapdoor for keyword w and interval of time periods $[s, e]$ consists of the user secret keys of all identities from $(w, s_1, \dots, s_{t(k)})$ to $(w, e_1, \dots, e_{t(k)})$, but taking advantage of the hierarchical structure to include entire subtrees of keys.

More precisely, let $\mathcal{HIBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ be a HIBE scheme. Then we associate to it a PETKS scheme $\mathcal{PETKS} = \text{hibe-2-petks}(\mathcal{HIBE}, t(k)) = (\text{KG}, \text{Td}, \text{PETKS}, \text{Test}, N)$ such that $N(k) = 2^{t(k)}$, $\text{KG}(1^k) = \text{Setup}(1^k)$ and $\text{PETKS}(pk, w, i) = (i, R, C)$ where $R \xleftarrow{\$} \{0, 1\}^k$ and $C \leftarrow \text{Enc}(pk, (w, i_1, \dots, i_{t(k)}), R)$. The trapdoor algorithm $\text{Td}(sk, w, s, e)$ first constructs a set T of identities as follows. Let j be the smallest index so that $s_j \neq e_j$. Then T is the set containing $(w, s_1, \dots, s_{t(k)})$, $(w, e_1, \dots, e_{t(k)})$, the right siblings of all nodes on the path from (w, s_1, \dots, s_{j+1}) to $(w, s_1, \dots, s_{t(k)})$, and the left siblings of all nodes on the path from (w, e_1, \dots, e_{j+1}) to $(w, e_1, \dots, e_{t(k)})$. If j does not exist, meaning $s = e$, then $T \leftarrow \{(w, s_1, \dots, s_{t(k)})\}$. The trapdoor t_w is the set of tuples $((w, i_1, \dots, i_r), \text{KeyDer}(sk, (w, i_1, \dots, i_r)))$ for all $(i_1, \dots, i_r) \in T$. To test a ciphertext (i, R, C) , the Test algorithm looks up a tuple $((w, i_1, \dots, i_r), usk[(w, i_1, \dots, i_r)])$ in t_w . It returns 0 when no such tuple is found. Otherwise, it derives $usk[(w, i_1, \dots, i_{t(k)})]$ using repetitive calls to the KeyDer algorithm, and returns 1 iff $\text{Dec}(usk[(w, i_1, \dots, i_{t(k)})], C) = R$.

Theorem 6.1. *Let \mathcal{HIBE} be a HIBE scheme, and let $\mathcal{PETKS} = \text{hibe-2-petks}(\mathcal{HIBE}, t(k))$ for some polynomially bounded function $t : \mathbb{N} \rightarrow \mathbb{N}$. If \mathcal{HIBE} is HIBE-ANO-CPA[1, $t(k) + 1$]-secure, then \mathcal{PETKS} is PETKS-IND-CPA-secure. Furthermore, if \mathcal{HIBE} is HIBE-IND-CPA[$t(k) + 1$]-secure, then \mathcal{PETKS} is computationally consistent.*

We split the proof of the theorem over the following two lemmas.

Lemma 6.2. *Let \mathcal{HIBE} be a HIBE scheme, and let $\mathcal{PETKS} = \text{hibe-2-petks}(\mathcal{HIBE}, t(k))$ for some polynomially bounded function $t : \mathbb{N} \rightarrow \mathbb{N}$. If \mathcal{HIBE} is HIBE-ANO-CPA[1, $t(k) + 1$]-secure, then \mathcal{PETKS} is PETKS-IND-CPA-secure.*

Proof. Let $\mathcal{HIBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ be a level-1 anonymous (HIBE-ANO-CPA[1, $t(k) + 1$]-secure) HIBE scheme, and let $\mathcal{PETKS} = \text{hibe-2-petks}(\mathcal{HIBE}, t(k)) = (\text{KG}, \text{Td}, \text{PETKS}, \text{Test}, N)$ be the associated PETKS scheme. Given an adversary \mathcal{A} breaking the PETKS-IND-CPA security of \mathcal{PETKS} , we construct an adversary \mathcal{B} breaking the HIBE-ANO-CPA[1, $t(k) + 1$] security of \mathcal{HIBE} as follows. On input public parameters pk , \mathcal{B} runs \mathcal{A} on inputs (find, pk) . When \mathcal{A} queries its TRAPD oracle for the trapdoor of keyword w for interval $[s, e]$, then \mathcal{B} constructs a set T exactly as the

Td algorithm does, and constructs the corresponding trapdoor by querying its KEYDER oracle for the user secret keys corresponding to all identities in T .

When \mathcal{A} outputs challenge keywords w_0, w_1 and time period i , \mathcal{B} outputs challenge identities $id_0 = (w_0, i_1, \dots, i_{t(k)}), id_1 = (w_1, i_1, \dots, i_{t(k)})$ and a randomly chosen message M of length k . Note that identities id_0 and id_1 differ on level 1, but are otherwise equal, as required for level-1 anonymity. Upon receiving challenge ciphertext C , adversary \mathcal{B} sends (i, R, C) to \mathcal{A} and runs it until \mathcal{A} outputs a bit b' (responding to \mathcal{A} 's oracle queries the same way as before). Adversary \mathcal{B} outputs the same bit b' .

It is easy to see that, due to the ordered structure of the time tree, adversary \mathcal{B} does not need to corrupt any ancestors of its challenge identities. Therefore, adversary \mathcal{B} succeeds whenever \mathcal{A} does, and we have

$$\mathbf{Adv}_{\mathcal{PETKS}, \mathcal{A}}^{\text{petks-ind-cpa}}(k) \leq \mathbf{Adv}_{\mathcal{HIBE}, \mathcal{B}}^{\text{hibe-ano-cpa}[1, t(k)+1]}(k)$$

for all $k \in \mathbb{N}$, from which the lemma follows. □

Lemma 6.3. *Let \mathcal{HIBE} be a HIBE scheme, and let $\mathcal{PETKS} = \text{hibe-2-petks}(\mathcal{HIBE}, t(k))$ for some polynomially bounded function $t : \mathbb{N} \rightarrow \mathbb{N}$. If \mathcal{HIBE} is HIBE-IND-CPA[$t(k) + 1$]-secure, then \mathcal{PETKS} is computationally consistent.*

Proof. Let \mathcal{A} be an adversary of the consistency of \mathcal{PETKS} . We construct an HIBE-IND-CPA adversary \mathcal{B} of \mathcal{HIBE} as follows.

Adversary $\mathcal{B}^{\text{KEYDER}(\cdot)}(\text{find}, pk)$

$(w, w', s, e, i) \xleftarrow{\$} \mathcal{A}(pk)$
 $R, R' \xleftarrow{\$} \{0, 1\}^k$ (where $\{0, 1\}^k$ is the message space of \mathcal{HIBE})
 $id \leftarrow (w, i_1, \dots, i_{t(k)})$
 $M_0 \leftarrow R; M_1 = R'$
 $state \leftarrow (pk, w, w', R, R', s, e, i)$
 return $(id, M_0, M_1, state)$

Adversary $\mathcal{B}^{\text{KEYDER}(\cdot)}(\text{guess}, C, state)$

parse C as (i, R, C')
 $t_w \xleftarrow{\$} \text{KEYDER}((w', i_1, \dots, i_{t(k)}))$; $X \leftarrow \text{Dec}(t_w, C')$
 if $X = R'$ then return 1 else return 0

Since, by construction, $\text{Test}(t_w, C)$ returns 0 whenever $i \notin [s, e]$, we can assume that $w' \neq w$ and $i \in [s, e]$. Then, exactly as in Theorem 4.2, we have

$$\Pr[\mathbf{Exp}_{\mathcal{HIBE}, \mathcal{B}}^{\text{hibe-ind-cpa-1}[t(k)+1]}(k) = 1] \geq \Pr[\mathbf{Exp}_{\mathcal{PETKS}, \mathcal{A}}^{\text{petks-consist}}(k) = 1], \tag{8}$$

$$\Pr[\mathbf{Exp}_{\mathcal{HIBE}, \mathcal{B}}^{\text{hibe-ind-cpa-0}[t(k)+1]}(k) = 1] \leq 2^{-l}. \tag{9}$$

Equations (8) and (9) give us

$$\mathbf{Adv}_{\mathcal{PETKS}, \mathcal{A}}^{\text{petks-consist}}(k) \leq \mathbf{Adv}_{\mathcal{HIBE}, \mathcal{B}}^{\text{hibe-ind-cpa}[t(k)+1]}(k) + 2^{-l}.$$

The result follows. □

Complexity Since the $mGS\text{-HIBE}$ has user secret keys and ciphertexts of size linear in the depth of the tree, our resulting PETKS scheme has public and secret keys of size $O(1)$, ciphertexts of size $O(\log N(k))$ and trapdoors of size $O(\log^2 N(k))$. We note that in this case a user can decrypt ciphertexts intended for any of its descendants directly, without needing to derive the corresponding secret key first. This makes the call to the KeyDer algorithm in the Test algorithm superfluous, thereby improving the efficiency of Test . Note that since the $mGS\text{-HIBE}$ scheme is only secure for tree depths $d(k) = O(\log(k))$, the derived PETKS scheme is restricted to a polynomial number of time periods.

Unbounded Time Periods Using the techniques of [21], one can create a variant of our scheme with efficiency depending on the number of *elapsed* time periods, rather than the maximal number of time periods $N(k)$. This means that there is no efficiency penalty for overestimating $N(k)$, so that a sufficiently high value can be chosen when setting up the system. However, for security reasons the number of time periods remains limited to a maximum of $N(k) \leq 2^{d(k) - \lceil \log d(k) \rceil - 1}$ periods, where $d(k)$ is the maximum depth of the underlying HIBE scheme.

7. Identity-Based Encryption with Keyword Search

In this section, we show how to combine the concepts of identity-based encryption and PEKS to obtain identity-based encryption with keyword search (IBEKS) or ID-based searchable encryption for short. Like in IBE schemes, this allows to use any string as a recipient's public key for the PEKS scheme.

7.1. Definitions

IBEKS Schemes An identity-based encryption with keyword search scheme $\text{IBEKS} = (\text{Setup}, \text{KeyDer}, \text{Td}, \text{IBEKS}, \text{Test})$ is made up of five algorithms. Via $(pk, msk) \xleftarrow{\$} \text{Setup}(1^k)$, where $k \in \mathbb{N}$ is the security parameter, the master generates the master keys; via $usk[id] \xleftarrow{\$} \text{KeyDer}^H(msk, id)$, the master computes the secret key for identity id ; via $C \xleftarrow{\$} \text{IBEKS}^H(pk, id, w)$, a sender encrypts a keyword w to identity id to get a ciphertext; via $t_w \xleftarrow{\$} \text{Td}^H(usk[id], w)$, the receiver computes a trapdoor t_w for keyword w and identity id and provides it to the gateway; via $b \leftarrow \text{Test}^H(t_w, C)$, the gateway tests whether C encrypts w , where b is a bit with 1 meaning “accept” or “yes” and 0 meaning “reject” or “no”. As usual H is a random oracle whose domain and/or range might depend on k and pk . For correctness, we require that for all $k \in \mathbb{N}$, all identities id , and all $w \in \{0, 1\}^*$,

$$\Pr[\text{Test}^H(\text{Td}^H(\text{KeyDer}^H(msk, id), w), \text{IBEKS}^H(pk, id, w)) = 1] = 1,$$

where the probability is taken over the choice of $(pk, msk) \xleftarrow{\$} \text{Setup}(1^k)$, the random choice of H , and the coins of all algorithms in the expression above.

Consistency The notion of consistency for IBEKS is similar to the one given for PEKS. The advantage of a user \mathcal{U} is defined as the probability that, on input the master public key pk , it can output keywords w, w' and identities id, id' such that $w \neq w'$ or $id \neq id'$, yet still an encryption of w under identity id tests positively under a trapdoor derived for keyword w' and identity id' . We again distinguish between perfect, statistical and computational consistency. Note that this definition also considers it a consistency problem if a trapdoor for identity id' tests positively for a ciphertext intended for identity $id \neq id'$. This type of problems is easily avoided by having the KeyDer, Td and IBEKS algorithms include the intended identity into the user secret keys, trapdoors and ciphertexts, respectively.

Privacy We define privacy for IBEKS schemes says that an adversary should not be able to distinguish between the encryption of two different challenge keywords w_0, w_1 of its choice for any identity id of its choice. Moreover, this should be the case even if the adversary is allowed to obtain trapdoors for non-challenge keywords issued for any identity and to obtain trapdoors for w_0, w_1 for identities other than id . The advantage function $\text{Adv}_{\text{IBEKS}, \mathcal{A}}^{\text{ibeks-ind-cpa}}(k)$ of an adversary \mathcal{A} and the notion of IBEKS-IND-CPA security are defined analogously to standard PEKS schemes.

7.2. A Generic Transformation from Anonymous HIBE Schemes

We now propose a generic transform, called *hibe-2-ibeks*, to convert any HIBE scheme with two levels into an IBEKS scheme. To obtain an IBEKS that is IBEKS-IND-CPA-secure, it is sufficient to start with a HIBE that is anonymous at level 2. Moreover, if the underlying HIBE is HIBE-IND-CPA[2]-secure, then the resulting IBEKS is also computationally consistent.

The hibe-2-ibeks Transform Given a HIBE scheme $\mathcal{HIBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ with two levels, *hibe-2-ibeks* returns the IBEKS scheme $\text{IBEKS} = (\text{Setup}, \overline{\text{KeyDer}}, \text{IBEKS}, \text{Td}, \text{Test})$ such that $\overline{\text{KeyDer}}(msk, id) = (usk, id)$ where $usk \xleftarrow{\$} \text{KeyDer}(msk, id)$, $\text{IBEKS}(pk, id, w) = (id, R, C)$ where $R \xleftarrow{\$} \{0, 1\}^k$ and $C = \text{Enc}(pk, (id, w), R)$, $\text{Td}(\overline{usk} = (usk, id), w) = (id, t_w)$ where $t_w \xleftarrow{\$} \text{KeyDer}(usk, (id, w))$ and $\text{Test}(\overline{t_w} = (id, t_w), (id', R, C))$ returns 1 iff $\text{Dec}(t_w, C) = R$ and $id = id'$.

Theorem 7.1. *Let \mathcal{HIBE} be a HIBE scheme and let $\text{IBEKS} = \text{hibe-2-ibeks}(\mathcal{HIBE})$. If \mathcal{HIBE} is HIBE-IND-CPA[2]-secure, then IBEKS is computationally consistent. Furthermore, if \mathcal{HIBE} is HIBE-ANO-CPA[2, 2]-secure, then IBEKS is IBEKS-IND-CPA-secure.*

The proof of Theorem 7.1 follows from Lemmas 7.2 and 7.3.

Lemma 7.2. *Let \mathcal{HIBE} be a HIBE scheme and let $\text{IBEKS} = \text{hibe-2-ibeks}(\mathcal{HIBE})$. If \mathcal{HIBE} is HIBE-ANO-CPA[2, 2]-secure, then IBEKS is IBEKS-IND-CPA-secure.*

Proof. Given an adversary \mathcal{A} breaking the IBEKS-IND-CPA-security of IBEKS , we construct an HIBE-ANO-CPA[2, 2]-adversary \mathcal{B} breaking \mathcal{HIBE} as follows. On input

a public key pk , algorithm \mathcal{B} runs \mathcal{A} on the same input, answering \mathcal{A} 's $\text{KEYDER}(\cdot)$ queries by forwarding the output of its own $\text{KEYDER}(\cdot)$ oracle, and answering \mathcal{A} 's $\text{TRAPD}(id, w)$ oracle queries by querying its own $\text{KEYDER}(\cdot)$ oracle for the secret key corresponding to identity (id, w) . When \mathcal{A} outputs a challenge identity id^* and two challenge keywords w_0^*, w_1^* , adversary \mathcal{B} chooses a random message $M^* \in \{0, 1\}^k$ and outputs M^* as the challenge message and $id_0^* = (id^*, w_0^*)$ and $id_1^* = (id^*, w_1^*)$ as the challenge identities, which in fact differ only in the second entry. Let C^* be the challenge ciphertext that \mathcal{B} receives at the beginning of its `guess` phase. Adversary \mathcal{B} returns (M^*, C^*) to \mathcal{A} , and continues to run \mathcal{A} (answering TRAPD queries the same way as before) until it outputs a bit b' . Algorithm \mathcal{B} then outputs the same bit b' as its own output.

It is clear from the construction that \mathcal{B} 's simulation of \mathcal{A} 's environment is perfect. Since \mathcal{A} cannot query its TRAPD oracle on keywords (id^*, w_0^*) and (id^*, w_1^*) , \mathcal{B} will not be forced to query its KEYDER on identities id_0^* and id_1^* , and hence wins the game whenever \mathcal{A} does. Therefore, we have that

$$\text{Adv}_{\text{IBEXS}, \mathcal{A}}^{\text{ibeks-ind-cpa}}(k) \leq \text{Adv}_{\mathcal{HIBE}, \mathcal{B}}^{\text{hibe-ano-cpa}[2,2]}(k),$$

from which the theorem follows for CPA security. This proves the lemma. \square

Lemma 7.3. *Let \mathcal{HIBE} be a HIBE scheme and let $\text{IBEXS} = \text{hibe-2-ibeks}(\mathcal{HIBE})$. If \mathcal{HIBE} is $\text{HIBE-IND-CPA}[2]$ -secure, then IBEXS is computationally consistent.*

Proof. Let \mathcal{A}_1 be an adversary of the consistency of IBEXS . We construct an $\text{HIBE-IND-CPA}[2]$ adversary \mathcal{B}_1 of \mathcal{HIBE} as follows.

Adversary $\mathcal{B}_1^{\text{KEYDER}(\cdot)}$ (find, pk)

$(w, w', id, id') \xleftarrow{\$} \mathcal{A}_1(pk)$; $R, R' \xleftarrow{\$} \{0, 1\}^l$ (where $\{0, 1\}^l$ is the message space of \mathcal{HIBE})

$id = (id', w)$

$w_0 = R$; $w_1 = R'$

$state = (pk, w, w', R, R')$

return $(id, w_0, w_1, state)$

Adversary $\mathcal{B}_1^{\text{KEYDER}(\cdot)}$ (guess, $C, state$)

$t_{w'} \xleftarrow{\$} \text{KEYDER}((id', w'))$; $X \leftarrow \text{Dec}(t_{w'}, C)$

if $X = R'$ then return 1 else return 0

Since, by construction, $\text{Test}(t_w, C)$ returns 0 whenever $id \neq id'$, we can assume that $w' \neq w$ and $id' = id$. Thus, exactly as in Theorem 4.2, we have

$$\Pr[\text{Exp}_{\mathcal{HIBE}, \mathcal{B}_1}^{\text{hibe-ind-cpa-1}[2]}(k) = 1] \geq \Pr[\text{Exp}_{\text{IBEXS}, \mathcal{A}}^{\text{ibeks-consist}}(k) = 1], \quad (10)$$

$$\Pr[\text{Exp}_{\mathcal{HIBE}, \mathcal{B}_1}^{\text{hibe-ind-cpa-0}[2]}(k) = 1] \leq 2^{-l}. \quad (11)$$

Equations (10) and (11) give us

$$\text{Adv}_{\text{IBEXS}, \mathcal{A}_1}^{\text{ibeks-consist}}(k) \leq \text{Adv}_{\mathcal{HIBE}, \mathcal{B}_1}^{\text{hibe-ind-cpa}[2]}(k) + 2^{-l}.$$

The result follows. \square

7.3. Concrete Instantiations

Neither the $\mathcal{GS}\text{-HIBE}$ scheme of [15] nor the $m\mathcal{GS}\text{-HIBE}$ scheme of Fig. 6 are anonymous at the second level. For the $\mathcal{GS}\text{-HIBE}$ scheme, consider an adversary \mathcal{A} who outputs challenge identities $id = (id_1, id_2)$ and $id' = (id_1, id'_2)$ for any $id_1, id_2, id'_2 \in \{0, 1\}^*$ such that $id_2 \neq id'_2$, and any challenge message $M \in \{0, 1\}^k$. When given the challenge ciphertext $C = (C_1, C_2, C_3)$, \mathcal{A} checks whether $e(C_1, H_1(id)) = e(P, C_2)$. (See (7) for how ciphertexts are created in the $\mathcal{GS}\text{-HIBE}$ scheme.) If the test succeeds, then \mathcal{A} returns 0, otherwise it returns 1. It is easy to see that the advantage of \mathcal{A} is $\text{Adv}_{\mathcal{GS}\text{-HIBE}, \mathcal{A}}^{\text{hibe-ano-cpa}[2,2]}(k) \geq 1 - 2^{-k}$. A similar attack can be mounted on the $m\mathcal{GS}\text{-HIBE}$ scheme by checking whether $e(C_1, H_{1,2}(id_2)) = e(P, C_2)$.

In Appendix A.3, we show that the recently introduced HIBE scheme by Boneh et al. [7] is not level-2 anonymous either (and actually, not anonymous at any level). Subsequent to our work, Boyen and Waters [11] proposed a fully anonymous HIBE scheme that, when used to instantiate our generic construction, immediately yields an IBEKS scheme with security and consistency in the standard model.

7.4. Identity-Based Encryption with Temporary Keyword Search

The ideas of Sects. 6 and 7 can be further combined to create an identity-based encryption scheme with temporary keyword search. This can be constructed from a level-2 anonymous HIBE scheme by putting the users' identities at the first level of the hierarchy, the keywords at the second, and a binary tree of time frames on the levels below.

Acknowledgements

We thank Nigel Smart for suggesting the concept of temporarily searchable encryption. Second and tenth authors were supported in part by NSF grants ANR-0129617 and CCR-0208842 and by an IBM Faculty Partnership Development Award. The fourth author was supported by the research program Sentinels (<http://www.sentinels.nl>). Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs. Fifth author was supported by an IBM Ph.D. Fellowship. Eighth author is a Postdoctoral Fellow of the Research Foundation—Flanders (FWO), and was supported in part by the Flemish Government under GOA Mefisto 2006/06 and Ambiorix 2005/11, and by the European Commission through the IST Project PRIME. The rest of the authors were supported in part by the European Commission through the IST Program under Contract IST-2002-507932 ECRYPT.

Appendix A. Attacks Against the Anonymity of Existing Schemes

A.1. Waters' IBE Scheme

We recall Waters' IBE scheme [25] $\mathcal{W}\text{-IBE} = (\text{Setup}, \text{KeyDer}, \text{Enc}, \text{Dec})$ in Fig. 7. Associated with $\mathcal{W}\text{-IBE}$ is a polynomial n . It is assumed that all user identities are $n(k)$ -bit (e.g. 160-bit) strings (for instance obtained by hashing the actual identity using a

<p>Setup(1^k)</p> <p>$(\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\\$} \mathcal{G}(1^k)$ $P, Q \xleftarrow{\\$} \mathbb{G}_1^*$; $\alpha \xleftarrow{\\$} \mathbb{Z}_p$; $P_1 \leftarrow \alpha P$; $Q_1 \leftarrow \alpha Q$ $U[0 \dots n] \xleftarrow{\\$} \mathbb{G}_1^{n+1}$; $E \leftarrow e(P, Q)$ $pk \leftarrow (\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, U, E)$ $msk \leftarrow (pk, Q_1)$ return (pk, msk)</p> <p>KeyDer(msk, id)</p> <p>parse msk as $((\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, U, E), Q_1)$ $r \xleftarrow{\\$} \mathbb{Z}_p$; $V \leftarrow U[0] + \sum_{i=1}^n id[i]U[i]$ $usk[id] \leftarrow (Q_1 + rV, rP)$ return $usk[id]$</p>	<p>Enc(pk, id, M)</p> <p>parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, U, E)$ $V \leftarrow U[0] + \sum_{i=1}^n id[i]U[i]$ $t \xleftarrow{\\$} \mathbb{Z}_p$; $T \leftarrow E^t$ $C \leftarrow (T \cdot M, tP, tV)$ return C</p> <p>Dec($usk[id], C$)</p> <p>parse $usk[id]$ as (S_1, S_2), C as (C_1, C_2, C_3) $T' \leftarrow e(S_1, C_2) \cdot e(S_2, C_3)^{-1}$ return $T'^{-1} \cdot C_1$</p>
--	---

Fig. 7. The algorithms constituting \mathcal{W} -IBE. Identities are represented as bit strings $id = id[1, \dots, n] \in \{0, 1\}^n$.

collision-resistant hash function), which are written as $id = id[1]id[2] \dots id[n]$, where each $id[i]$ ($1 \leq i \leq n$) is a bit $id[i] \in \{0, 1\}$. (We drop the argument k to n when k is understood.) The message space is defined by $\text{MsgSp}(k) = \{0, 1\}^k$, and messages are encoded as elements of \mathbb{G}_2 in the scheme.

We now describe a PTA \mathcal{A} against the IBE-ANO-CPA-security of \mathcal{W} -IBE. In the find stage it gets input a public key $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, U, E)$, and returns any two distinct n -bit strings id_0, id_1 as challenge identities, along with any k -bit challenge message. In the guess phase, given a challenge ciphertext $C = (C_1, C_2, C_3)$ formed by encrypting M under id_b , where $b \in \{0, 1\}$ is the challenge bit, it computes $V' \leftarrow U[0] + \sum_{i=1}^n id_1[i]U[i]$. If $e(P, C_3) = e(C_2, V')$ then it returns 1 else it returns 0. It is easy to see that $\text{Adv}_{\mathcal{W}\text{-IBE}, \mathcal{A}}^{\text{ibe-ano-cpa}}(k) \geq 1 - 2^{-k}$.

A.2. Boneh–Boyen’s IBE Scheme

The IBE scheme by Boneh and Boyen [6], here referred to as \mathcal{BB} -IBE, is depicted in Fig. 8. An identity is represented by a vector of $n(k)$ symbols $id[1, \dots, n] \in \Sigma^n$ where Σ is an alphabet of size s . In the original scheme, these are obtained as the output of an admissible hash function, but we ignore this here as it is irrelevant to the attack.

Consider a PTA \mathcal{A} that, on input $pk = (\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, U)$, outputs any two distinct identities $id_0, id_1 \in \Sigma^n$, and any message $M \in \{0, 1\}^k$. Let $i \in \{1, \dots, n\}$ be an index so that $id_0[i] \neq id_1[i]$. When \mathcal{A} is given the challenge ciphertext $C = (C_1, \dots, C_{n+2})$, it checks whether $e(C_2, U[i, id_0[i]]) = e(P, C_{i+2})$. If so, then \mathcal{A} returns 0, else it returns 1. It is easily verified that $\text{Adv}_{\mathcal{BB}\text{-IBE}, \mathcal{A}}^{\text{ibe-ano-cpa}}(k) \geq 1 - 2^{-k}$.

A.3. Boneh–Boyen–Goh’s HIBE Scheme

The recently proposed \mathcal{BBG} -HIBE scheme [7], depicted in Fig. 9, is not anonymous at any single level, and therefore not at any set of multiple levels either. This can be seen from the following adversary \mathcal{A} that breaks the anonymity at level l . On input $pk = (\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, Q_2, U)$, adversary \mathcal{A} outputs challenge identities

<p>Setup(1^k)</p> $(\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\$} \mathcal{G}(1^k)$ $P, Q \xleftarrow{\$} \mathbb{G}_1^*; \alpha \xleftarrow{\$} \mathbb{Z}_p;$ $P_1 \leftarrow \alpha P; Q_1 \leftarrow \alpha Q$ $U[1 \dots n, 1 \dots s] \xleftarrow{\$} \mathbb{G}_1^{n \times s}$ $pk \leftarrow (\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, U)$ $msk \leftarrow (pk, Q_1)$ <p>return (pk, msk)</p> <p>KeyDer(msk, id)</p> <p>parse msk as $((\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, U, Q), Q_1)$</p> $r_1, \dots, r_n \xleftarrow{\$} \mathbb{Z}_p; V \leftarrow \sum_{i=1}^n r_i U[i, id[i]]$ $usk[id] \leftarrow (Q_1 + V, r_1 P, \dots, r_n P)$ <p>return $usk[id]$</p>	<p>Enc(pk, id, M)</p> <p>parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, U, Q)$</p> $t \xleftarrow{\$} \mathbb{Z}_p; T \leftarrow e(P_1, Q)^t$ $C \leftarrow (T \cdot M, tP, tU[1, id[1]], \dots,$ $tU[n, id[n]])$ <p>return C</p> <p>Dec($usk[id], C$)</p> <p>parse $usk[id]$ as $(S_1, S_2, \dots, S_{n+1})$</p> <p>parse C as (C_1, \dots, C_{n+2})</p> $T' \leftarrow e(S_1, C_2) \cdot \prod_{i=1}^n e(S_{i+1}, C_{i+2})^{-1}$ <p>return $T'^{-1} \cdot C_1$</p>
--	---

Fig. 8. The algorithms constituting $\mathcal{BB}\text{-}\mathcal{IBE}$. Identities are represented as vectors of symbols $id = id[1, \dots, n] \in \Sigma^n$, where $|\Sigma| = s$.

<p>Setup(1^k)</p> $(\mathbb{G}_1, \mathbb{G}_2, p, e) \xleftarrow{\$} \mathcal{G}(1^k)$ $P \xleftarrow{\$} \mathbb{G}_1; \alpha \xleftarrow{\$} \mathbb{Z}_p; P_1 \leftarrow \alpha P$ $Q, Q_2 \xleftarrow{\$} \mathbb{G}_1; Q_1 \leftarrow \alpha Q$ $U[1 \dots d(k)] \xleftarrow{\$} \mathbb{G}_1^{d(k)}$ $pk \leftarrow (\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, Q_2, U)$ $msk \leftarrow (pk, Q_1, 0, \dots, 0)$ <p>return (pk, msk)</p> <p>KeyDer(usk, id)</p> $l \leftarrow id $ <p>parse usk as $(pk, A, B, S_l, \dots, S_{d(k)})$</p> <p>parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, Q_2, U)$</p> <p>parse id as $(id_1, \dots, id_l); r \xleftarrow{\\$} \mathbb{Z}_p$</p> $A' \leftarrow A + id_l S_l$ $+ r(id_1 U[1] + \dots + id_l U[l] + Q_2)$ $B' \leftarrow B + rP$ <p>for $l+1 \leq i \leq d(k)$ do $S'_i \leftarrow S_i + rU[i]$</p> <p>return $(pk, A', B', S'_{l+1}, \dots, S'_{d(k)})$</p>	<p>Enc(pk, id, M)</p> <p>parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, Q_2, U)$</p> <p>parse id as (id_1, \dots, id_l)</p> $t \xleftarrow{\$} \mathbb{Z}_p; T \leftarrow e(P_1, Q)^t$ $C_3 \leftarrow t(id_1 U[1] + \dots + id_l U[l] + Q_2)$ $C \leftarrow (T \cdot M, tP, C_3)$ <p>return C</p> <p>Dec(usk, C)</p> <p>parse usk as $(pk, A, B, S_l, \dots, S_{d(k)})$</p> <p>parse pk as $(\mathbb{G}_1, \mathbb{G}_2, p, e, P, P_1, Q, Q_2, U)$</p> <p>parse C as (C_1, C_2, C_3)</p> $T' \leftarrow e(A, C_2) \cdot e(B, C_3)^{-1}$ <p>return $T'^{-1} \cdot C_1$</p>
--	--

Fig. 9. The algorithms constituting $\mathcal{BBG}\text{-}\mathcal{H}\mathcal{IBE}$ with maximum hierarchy depth $d(k)$. An identity at level l is represented as a vector $id = (id_1, \dots, id_l) \in \mathbb{Z}_p^l$.

$(id_1, \dots, id_{l-1}, id_l)$ and $(id_1, \dots, id_{l-1}, id'_l)$ for any $id_1, \dots, id_l, id'_l \in \mathbb{Z}_p$ such that $id_l \neq id'_l$, and any challenge message $M \in \{0, 1\}^k$. When given the challenge ciphertext $C = (C_1, C_2, C_3)$, \mathcal{A} checks whether $e(C_2, id_1 U[1] + \dots + id_l U[l] + Q_2) = e(P, C_3)$.

If this is the case, then \mathcal{A} returns 0, otherwise it returns 1. It is easily verified that $\text{Adv}_{\text{BBG-HIBE}, \mathcal{A}}^{\text{hibe-ano-cpa}[l, d]}(k) \geq 1 - 2^{-k}$.

Appendix B. Proof of Lemma 5.3

Suppose that there is an adversary \mathcal{A} of $mGS\text{-HIBE}$ that breaks its HIBE-IND-CPA[d] security. We will show how to use \mathcal{A} in the construction of a simulator \mathcal{B} that solves the bilinear Diffie–Hellman problem. Let $n_{1,i}$ be the number of queries that \mathcal{A} makes to the $H_{1,i}$ oracle, let n_2 be the number of queries to the H_2 oracle, let n_e be the number of queries to the key extraction oracle, and let $n_h = \sum_{i=1}^{d(k)} n_{1,i} + n_2$ be the total number of hash queries.

The simulator is given as input (P, aP, bP, cP) . It sets $Q_0 \leftarrow bP$ as the public key and then runs $\mathcal{A}(\text{find}, pk)$. The simulator responds to \mathcal{A} 's queries as described below. To maintain consistency between queries it keeps lists $L_{1,1}, \dots, L_{1,d(k)}$, L_2 and L_3 . All lists are initially empty. At the very beginning the simulator chooses $n_{1,i}^* \xleftarrow{\$} \{1, \dots, n_{1,i}\}$ and $s_i^*, x_i^* \xleftarrow{\$} \mathbb{Z}_p^*$, and it computes $Q_i^* \leftarrow s_i^*(bP)$ for $1 \leq i \leq d(k)$.

For the description of the simulation we distinguish between $H_{1,1}$ queries and $H_{1,i}$ queries for $i \geq 2$. Without loss of generality, we assume that before querying the KEYDER oracle to obtain the secret key of $id = (id_1, \dots, id_l)$, adversary \mathcal{A} first queried $H_{1,i}(id_i)$ for all $1 \leq i \leq l$.

$H_{1,1}$ Queries: To respond to a query id_1 , proceed as follows.

- If $L_{1,1}$ contains $(id_1, P_1, *)$ for some P_1 , respond with P_1 .
- If this is the $n_{1,1}^*$ -th call to the $H_{1,1}$ oracle, let $id_1^* \leftarrow id_1$, add (id_1^*, aP, \perp) to $L_{1,1}$ and respond with aP .
- Else, randomly choose an integer $x_1 \xleftarrow{\$} \mathbb{Z}_p^*$, add (id_1, x_1P, x_1) to $L_{1,1}$ and reply with x_1P .

$H_{1,i}$ Queries, $i \geq 2$: To respond to a query id_i , proceed as follows.

- If $L_{1,i}$ contains $(id_i, P_i, *)$ for some P_i then respond with P_i .
- If this is the $n_{1,i}^*$ -th query to the $H_{1,i}$ oracle, let $id_i^* \leftarrow id_i$, add (id_i^*, x_i^*P, x_i^*) to $L_{1,i}$ and respond with x_i^*P .
- Else, choose an integer $x_i \xleftarrow{\$} \mathbb{Z}_p^*$ and compute $P_i \leftarrow x_iP - s_{i-1}^{*-1}(aP + \sum_{j=2}^{i-1} s_{j-1}^* x_j^* P)$. If $P_i = 0$, then abort; else, add (id_i, P_i, x_i) to $L_{1,i}$ and reply with P_i .

H_2 Queries: To respond to a query κ , proceed as follows.

- If $(\kappa, K) \in L_2$ for some K , respond with K .
- Else, choose K uniformly at random from $\{0, 1\}^n$, respond with K and add (κ, K) to L_2 .

KEYDER Queries: To respond to a query $id = (id_1, \dots, id_l)$, proceed as follows.

- If $(id_1, \dots, id_l) = (id_1^*, \dots, id_l^*)$, then \mathcal{B} aborts.
- Let j be the largest integer $1 \leq j \leq l$ so that $(id|_j, S_j, Q_1, \dots, Q_{j-1}, s_j) \in L_3$, or let $j = 0$ if such element does not exist.

- For $i = j + 1, \dots, l$, do the following:
 - Find $(id_i, P_i, x_i) \in L_{1,i}$.
 - If $i = 1$ and $id_1 = id_1^*$, then add (id_1^*, \perp, \perp) to L_3 . If $i = 1$ and $id_1 \neq id_1^*$, then compute $S_1 \leftarrow x_1(bP)$, choose $s_1 \xleftarrow{\$} \mathbb{Z}_p^*$, and add (id_1, S_1, s_1) to L_3 .
 - If $i > 1$ and $S_{i-1} \neq \perp$, then look up (id_i, P_i, x_i) in $L_{1,i}$, compute $S_i \leftarrow S_{i-1} + s_{i-1}P_i$, $Q_{i-1} \leftarrow s_{i-1}P$, choose $s_i \xleftarrow{\$} \mathbb{Z}_p^*$, and add $(id|i, S_i, Q_1, \dots, Q_{i-1}, s_i)$ to L_3 .
 - If $i > 1$ and $S_{i-1} = \perp$ and $id_i = id_i^*$, then compute $Q_{i-1} \leftarrow s_{i-1}^*(bP)$ and add $(id|i, \perp, Q_1, \dots, Q_{i-1}, \perp)$ to L_3 .
 - If $i > 1$, $S_{i-1} = \perp$ and $id_i \neq id_i^*$, then look up (id_i, P_i, x_i) in $L_{1,i}$, compute $S_i \leftarrow s_{i-1}^*x_iP$, let $Q_{i-1} \leftarrow s_{i-1}^*(bP)$, choose $s_i \xleftarrow{\$} \mathbb{Z}_p^*$, and add $(id|i, S_i, Q_1^*, \dots, Q_{i-1}^*, s_i)$ to L_3 .
- Find $(id, S_l, Q_1, \dots, Q_{l-1}, s_l) \in L_3$ and return $(id, S_l, Q_1, \dots, Q_{l-1}, s_l)$.

At some point \mathcal{A} outputs $(id = (id_1, id_2, \dots, id_l), M_0, M_1, state)$. Without loss of generality, we assume that the adversary submitted id_i to the $H_{1,i}$ oracle before for all $1 \leq i \leq l$. If $id \neq (id_1^*, \dots, id_l^*)$, then \mathcal{B} aborts. Otherwise, he sets $C_1^* \leftarrow cP$, $C_2^* \leftarrow x_2^*(cP), \dots, C_l \leftarrow x_l^*(cP)$, he chooses C_{l+1}^* uniformly at random from $\{0, 1\}^n$, and lets $C^* \leftarrow (C_1^*, C_2^*, \dots, C_{l+1}^*)$. He then proceeds to run $\mathcal{A}(guess, C^*, state)$. Once \mathcal{A} completes its attack by outputting its guess b' , the simulator chooses a random element (κ, K) from L_2 and outputs κ as its solution to the bilinear Diffie–Hellman problem.

We first show that our simulator \mathcal{B} provides a real attack environment for \mathcal{A} as long as \mathcal{B} doesn't abort. The public key pk given to \mathcal{A} is correctly distributed because the challenge elements aP, bP, cP are random elements from \mathbb{G}_1^* . The responses to $H_{1,i}$ queries are uniformly distributed over \mathbb{G}_1^* due to the independent random choices of x_i (when simulating queries $H_{1,i}(id_i)$, $id_i \neq id_i^*$, $1 \leq i \leq d(k)$), of x_i^* (which is used to simulate $H_{1,i}(id_i^*)$ queries, $2 \leq i \leq d(k)$) and due to the uniform distribution of aP (which is used to simulate $H_{1,1}(id_1^*)$). Responses to H_2 queries are easily seen to be correctly distributed. The way KEYDER queries are handled requires a bit more explanation. For all level-1 identities $id_1 \neq id_1^*$, the returned secret key (S_1, s_1) contains the unique group element S_1 such that $e(Q_0, H_{1,1}(id_1)) = e(S_1, P)$ and a uniformly distributed scalar s_1 , as in the real game. For all descendants of $id_1 \neq id_1^*$, the secret keys are derived from (S_1, s_1) exactly as in the real scheme. Now consider identity $(id_1^*, \dots, id_{i-1}^*, id_i)$ with $id_i \neq id_i^*$, for which a tuple $(S_i, Q_1, \dots, Q_{i-1}, s_i)$ is returned as the secret key. The values Q_1, \dots, Q_{i-2} are inherited from the ancestors, as in the real scheme; Q_{i-1} is a random group element due to the random choice of s_{i-1}^* ; and s_i is a random element in \mathbb{Z}_p^* . The simulated value $S_i = s_{i-1}^*x_i(bP)$ is then the unique group element such that $e(H_{1,1}(id_1^*), Q_0) = e(S_i, P) \cdot \prod_{j=2}^{i-1} e(H_{1,j}(id_j^*), Q_{j-1})^{-1} \cdot e(H_{1,i}(id_i), Q_{i-1})^{-1}$, as required by the scheme. This can be seen from:

$$\begin{aligned}
 & e(H_{1,1}(id_1^*), Q_0) \cdot \prod_{j=2}^{i-1} e(H_{1,j}(id_j^*), Q_{j-1}) \cdot e(H_{1,i}(id_i), Q_{i-1}) \\
 &= e(aP, bP) \cdot \prod_{j=2}^{i-1} e(x_j^*P, s_{j-1}^*bP)
 \end{aligned}$$

$$\begin{aligned}
& \cdot e\left(x_i P - s_{i-1}^* \left(aP + \sum_{j=2}^{i-1} s_{j-1}^* x_j^* P\right), s_{i-1}^* bP\right) \\
& = e(aP, bP) \cdot e\left(\sum_{j=2}^{i-1} s_{j-1}^* x_j^* P, bP\right) \cdot e\left(s_{i-1}^* x_i P - aP - \sum_{j=2}^{i-1} s_{j-1}^* x_j^* P, bP\right) \\
& = e(S_3, P).
\end{aligned}$$

The secret keys of descendants of these nodes are derived from $(S_i, Q_1, \dots, Q_{i-1}, s_i)$ as dictated by the scheme, and hence are correctly distributed as well.

The only part of \mathcal{A} 's environment left to analyze is the challenge ciphertext $C^* = (C_1^*, \dots, C_{l+1}^*)$. The first component $C_1^* = cP$ is uniformly distributed over \mathbb{G}_1^* , and the second to l -th components are the unique group elements such that $e(C_i^*, P) = e(C_1^*, H_{1,i}(id_i^*))$ for $2 \leq i \leq l$. The last component C_{l+1}^* however may deviate from the distribution in a real game, depending on \mathcal{A} 's H_2 queries. In the following, we show that this does not harm our analysis, intuitively because the only way \mathcal{A} can distinguish between the real and the simulated game is by making an H_2 query that helps \mathcal{B} solve the BDH problem.

Let s_0 be the master secret key of the scheme in a real HIBE-IND-CPA[d] attack on $mGS\text{-}HIBE$, and let $D \leftarrow e(s_0 H_{1,1}(id_1), C_1^*)$. Let ASK be the event that \mathcal{A} queries the H_2 oracle on point D . Let $\Pr_{\mathcal{R}}[\cdot]$ denote the probability of an event taking place in a real attack on $mGS\text{-}HIBE$, and let $\Pr_{\mathcal{B}}[\cdot]$ denote the probability in the environment simulated by \mathcal{B} . We argue that $\Pr_{\mathcal{R}}[ASK] = \Pr_{\mathcal{B}}[ASK]$, as long as \mathcal{B} doesn't abort. Let ASK_i be the event that \mathcal{A} queries $H_2(D)$ within the first i queries to H_2 . Obviously, $\Pr_{\mathcal{R}}[ASK_0] = \Pr_{\mathcal{B}}[ASK_0] = 0$. Now assume that $\Pr_{\mathcal{R}}[ASK_{i-1}] = \Pr_{\mathcal{B}}[ASK_{i-1}]$. We have that

$$\begin{aligned}
\Pr_{\mathcal{R}}[ASK_i] &= \Pr_{\mathcal{R}}[ASK_i | ASK_{i-1}] \cdot \Pr_{\mathcal{R}}[ASK_{i-1}] \\
&\quad + \Pr_{\mathcal{R}}[ASK_i | \neg ASK_{i-1}] \cdot \Pr_{\mathcal{R}}[\neg ASK_{i-1}] \\
&= \Pr_{\mathcal{R}}[ASK_{i-1}] + \Pr_{\mathcal{R}}[ASK_i | \neg ASK_{i-1}] \cdot \Pr_{\mathcal{R}}[\neg ASK_{i-1}].
\end{aligned}$$

We know that $\Pr_{\mathcal{R}}[ASK_{i-1}] = \Pr_{\mathcal{B}}[ASK_{i-1}]$, so we only have to show that $\Pr_{\mathcal{R}}[ASK_i | \neg ASK_{i-1}] = \Pr_{\mathcal{B}}[ASK_i | \neg ASK_{i-1}]$. Given that $\neg ASK_{i-1}$ and that \mathcal{B} 's simulation didn't abort, the simulated public key, the oracle responses and the first l components of the ciphertext provided by \mathcal{B} are distributed exactly as in a real attack, as we explained before. Moreover, since \mathcal{A} did not query for $H_2(D)$ yet, from \mathcal{A} 's point of view the last ciphertext component C_{l+1}^* is a random string in $\{0, 1\}^n$, both in the real attack and in the simulated environment. Since all the information on which \mathcal{A} can base its decision for its next H_2 query is identically distributed in both environments, the probability that \mathcal{A} chooses to query D is the same in both environments as well. Hence, we have that $\Pr_{\mathcal{R}}[ASK_i] = \Pr_{\mathcal{B}}[ASK_i]$, and by induction that $\Pr_{\mathcal{R}}[ASK] = \Pr_{\mathcal{B}}[ASK]$.

The probability that \mathcal{A} wins a real attack against $mGS\text{-}HIBE$ can be written as

$$\Pr_{\mathcal{R}}[\mathcal{A} \text{ WINS}] = \Pr_{\mathcal{R}}[\mathcal{A} \text{ WINS} \wedge ASK] + \Pr_{\mathcal{R}}[\mathcal{A} \text{ WINS} \wedge \neg ASK]$$

$$= \Pr_{\mathcal{R}}[\mathcal{A} \text{ WINS} \wedge \text{ASK}] + \frac{1}{2} \leq \Pr_{\mathcal{R}}[\text{ASK}] + \frac{1}{2},$$

where the second equation is true because in the event $\neg \text{ASK}$, the distribution of the challenge ciphertext is completely independent of M_0, M_1 , and hence the probability that \mathcal{A} guesses correctly is $1/2$. Since $\Pr_{\mathcal{R}}[\text{ASK}] = \Pr_{\mathcal{B}}[\text{ASK}]$ and moreover

$$\Pr_{\mathcal{R}}[\mathcal{A} \text{ WINS}] = \frac{1}{2} \cdot \mathbf{Adv}_{mGS-\mathcal{H}^{\text{IBE}}, \mathcal{A}}^{\text{hibe-ind-cpa}[d]}(k) + \frac{1}{2},$$

it follows that

$$\Pr_{\mathcal{B}}[\text{ASK}] \geq \frac{1}{2} \cdot \mathbf{Adv}_{mGS-\mathcal{H}^{\text{IBE}}, \mathcal{A}}^{\text{hibe-ind-cpa}[d]}(k).$$

Now we only have to relate \mathcal{B} 's advantage in solving the BDH problem to $\Pr_{\mathcal{B}}[\text{ASK}]$. In the game simulated by \mathcal{B} , the probability that \mathcal{B} guesses the correct identities such that $id = (id_1^*, \dots, id_l^*)$ $1/\prod_{i=1}^l n_{1,i} \geq n_h^{-d(k)}$; the probability that \mathcal{B} guesses the correct H_2 query is $1/n_2 \geq n_h^{-1}$; and the probability that \mathcal{B} aborts when answering $H_{1,i}$ queries is $\sum_{i=1}^{d(k)} n_{1,i}/(p-1) \leq n_h/2^k$. The advantage of \mathcal{B} in solving the BDH problem is

$$\mathbf{Adv}_{\mathcal{G}, \mathcal{B}}^{\text{bdh}}(k) \geq \frac{1}{n_h^{d(k)+1}} \cdot \left(1 - \frac{n_h}{2^k}\right) \cdot \Pr_{\mathcal{B}}[\text{ASK}]$$

and hence

$$\mathbf{Adv}_{mGS-\mathcal{H}^{\text{IBE}}, \mathcal{A}}^{\text{hibe-ind-cpa}[d]}(k) \leq 2 \cdot n_h^{d(k)+1} \cdot \mathbf{Adv}_{\mathcal{G}, \mathcal{B}}^{\text{bdh}}(k) + \frac{n_h}{2^k},$$

from which the theorem follows.

References

- [1] M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, H. Shi, Searchable encryption revisited: consistency properties, relation to anonymous IBE, and extensions, *Cryptology ePrint Archive*, 2005. <http://eprint.iacr.org/>
- [2] M. Ajtai, C. Dwork, A public-key cryptosystem with worst-case/average-case equivalence, in *29th Annual ACM Symposium on Theory of Computing*, El Paso, TX, USA, May 4–6, 1997 (ACM Press, New York, 1997), pp. 284–293
- [3] M. Bellare, A. Boldyreva, A. Desai, D. Pointcheval, Key-privacy in public-key encryption, in *Advances in Cryptology—ASIA CRYPTO 2001*, ed. by C. Boyd, Gold Coast, Australia, December 9–13, 2001. Lecture Notes in Computer Science, vol. 2248 (Springer, Berlin, 2001), pp. 566–582
- [4] M. Bellare, S.K. Miner, A forward-secure digital signature scheme, in *Advances in Cryptology—CRYPTO'99*, ed. by M.J. Wiener, Santa Barbara, CA, USA, August 15–19, 1999. Lecture Notes in Computer Science, vol. 1666 (Springer, Berlin, 1994), pp. 431–448
- [5] M. Bellare, P. Rogaway, Random oracles are practical: a paradigm for designing efficient protocols, in *ACM CCS 93: 1st Conference on Computer and Communications Security*, Fairfax, VA, November 3–5, 1993 (ACM Press, New York, 1993), pp. 62–73
- [6] D. Boneh, X. Boyen, Secure identity based encryption without random oracles, in *Advances in Cryptology—CRYPTO 2004*, ed. by M. Franklin, Santa Barbara, CA, August 15–19, 2004. Lecture Notes in Computer Science, vol. 3152 (Springer, Berlin, 2004), pp. 443–459
- [7] D. Boneh, X. Boyen, E.-J. Goh, Hierarchical identity based encryption with constant size ciphertext, in *Advances in Cryptology—EUROCRYPT 2005*, ed. by R. Cramer, Aarhus, Denmark, May 22–26, 2005. Lecture Notes in Computer Science, vol. 3494 (Springer, Berlin, 2005), pp. 440–456

- [8] D. Boneh, G. Di Crescenzo, R. Ostrovsky, G. Persiano, Public key encryption with keyword search, in *Advances in Cryptology—EUROCRYPT 2004*, ed. by C. Cachin, J. Camenisch, Interlaken, Switzerland, May 2–6, 2004. Lecture Notes in Computer Science, vol. 3027 (Springer, Berlin, 2004), pp. 506–522
- [9] D. Boneh, M.K. Franklin, Identity based encryption from the Weil pairing, *SIAM J. Comput.* **32**(3), 586–615 (2003)
- [10] D. Boneh, Brent R. Waters, Conjunctive, subset, and range queries on encrypted data, in *TCC 2007: 4th Theory of Cryptography Conference*, ed. by S.P. Vadhan, Amsterdam, The Netherlands, February 21–24, 2007. Lecture Notes in Computer Science, vol. 4392 (Springer, Berlin, 2007). Also available at <http://eprint.iacr.org/>, Report 2006/287
- [11] X. Boyen, B. Waters, Anonymous hierarchical identity-based encryption (without random oracles), in *Advances in Cryptology—CRYPTO 2006*, ed. by C. Dwork, Santa Barbara, CA, August 20–24, 2006. Lecture Notes in Computer Science, vol. 4117 (Springer, Berlin, 2006), pp. 290–307
- [12] R. Canetti, S. Halevi, J. Katz, A forward-secure public-key encryption scheme, in *Advances in Cryptology—EUROCRYPT 2003*, ed. by E. Biham, Warsaw, Poland, May 4–8, 2003. Lecture Notes in Computer Science, vol. 2656 (Springer, Berlin, 2003), pp. 255–271
- [13] C. Dwork, M. Naor, O. Reingold, Immunizing encryption schemes from decryption errors, in *Advances in Cryptology—EUROCRYPT 2004*, ed. by C. Cachin, J. Camenisch, Interlaken, Switzerland, May 2–6, 2004. Lecture Notes in Computer Science, vol. 3027 (Springer, Berlin, 2004), pp. 342–360
- [14] C. Gentry, Practical identity-based encryption without random oracles, in *Advances in Cryptology—EUROCRYPT 2006*, ed. by S. Vaudenay, St. Petersburg, Russia, May 28–June 1, 2006. Lecture Notes in Computer Science, vol. 4004 (Springer, Berlin, 2006), pp. 445–464
- [15] C. Gentry, A. Silverberg, Hierarchical ID-based cryptography, in *Advances in Cryptology—ASIACRYPT 2002*, ed. by Y. Zheng, Queenstown, New Zealand, December 1–5, 2002. Lecture Notes in Computer Science, vol. 2501 (Springer, Berlin, 2002), pp. 548–566
- [16] E.-J. Goh, Secure indexes, Cryptology ePrint Archive, Report 2003/216, 2003. <http://eprint.iacr.org/>
- [17] O. Goldreich, *Foundations of Cryptography: Basic Applications*, vol. 2 (Cambridge University Press, Cambridge, 2004)
- [18] P. Golle, J. Staddon, B.R. Waters, Secure conjunctive keyword search over encrypted data, in *ACNS 04: 2nd International Conference on Applied Cryptography and Network Security*, ed. by M. Jakobsson, M. Yung, J. Zhou, Yellow Mountain, China, June 8–11, 2004. Lecture Notes in Computer Science, vol. 3089 (Springer, Berlin, 2004), pp. 31–45
- [19] S. Halevi, A sufficient condition for key-privacy, Cryptology ePrint Archive, Report 2005/005, 2005. <http://eprint.iacr.org/>
- [20] J. Horwitz, B. Lynn, Toward hierarchical identity-based encryption, in *Advances in Cryptology—EUROCRYPT 2002*, ed. by L.R. Knudsen, Amsterdam, The Netherlands, April 28–May 2, 2002. Lecture Notes in Computer Science, vol. 2332 (Springer, Berlin, 2002), pp. 466–481
- [21] T. Malkin, D. Micciancio, S.K. Miner, Efficient generic forward-secure signatures with an unbounded number of time periods, in *Advances in Cryptology—EUROCRYPT 2002*, ed. by L.R. Knudsen, Amsterdam, The Netherlands, April 28–May 2, 2002. Lecture Notes in Computer Science, vol. 2332 (Springer, Berlin, 2002), pp. 400–417
- [22] D.J. Park, K. Kim, P.J. Lee, Public key encryption with conjunctive field keyword search, in *WISA 04: 5th International Workshop on Information Security Applications*, ed. by C.H. Lim, M. Yung, Jeju Island, Korea, August 23–25, 2004. Lecture Notes in Computer Science, vol. 3325 (Springer, Berlin, 2004), pp. 73–86
- [23] A. Shamir, Identity-based cryptosystems and signature schemes, in *Advances in Cryptology—CRYPTO’84*, ed. by G.R. Blakley, D. Chaum, Santa Barbara, CA, August 19–23, 1985. Lecture Notes in Computer Science, vol. 196 (Springer, Berlin, 1985), pp. 47–53
- [24] D.X. Song, D. Wagner, A. Perrig, Practical techniques for searches on encrypted data, in *2000 IEEE Symposium on Security and Privacy*, Oakland, CA, May 2000 (IEEE Computer Society Press, Reading, 2000), pp. 44–55
- [25] B.R. Waters, Efficient identity-based encryption without random oracles, in *Advances in Cryptology—EUROCRYPT 2005*, ed. by R. Cramer, Aarhus, Denmark, May 22–26, 2005. Lecture Notes in Computer Science, vol. 3494 (Springer, Berlin, 2005), pp. 114–127
- [26] B.R. Waters, D. Balfanz, G. Durfee, D.K. Smetters, Building an encrypted and searchable audit log, in *ISOC Network and Distributed System Security Symposium—NDSS 2004*, San Diego, CA, February 4–6, 2004 (The Internet Society, Reston, 2004)