From Selective to Full Security: Semi-generic Transformations in the Standard Model

Michel Abdalla[1], Dario Fiore[2,⋆], and Vadim Lyubashevsky[1]

[1] Département d'Informatique, École normale supérieure, France.
{Michel.Abdalla,Vadim.Lyubashevsky}@ens.fr
[2] Department of Computer Science, New York University, USA.
fiore@cs.nyu.edu

Abstract. In this paper, we propose an efficient, standard model, semi-generic transformation of selective-secure (Hierarchical) Identity-Based Encryption schemes into fully secure ones. The main step is a procedure that uses admissible hash functions (whose existence is implied by collision-resistant hash functions) to convert any selective-secure *wild-carded* identity-based encryption (WIBE) scheme into a fully secure (H)IBE scheme. Since building a selective-secure WIBE, especially with a selective-secure HIBE already in hand, is usually much less involved than directly building a fully secure HIBE, this transform already significantly simplifies the latter task. This black-box transformation easily extends to schemes secure in the Continual Memory Leakage (CML) model of Brakerski et al. (FOCS 2010), which allows us obtain a new fully secure IBE in that model. We furthermore show that if a selective-secure HIBE scheme satisfies a particular security notion, then it can be generically transformed into a selective-secure WIBE. We demonstrate that several current schemes already fit this new definition, while some others that do not obviously satisfy it can still be easily modified into a selective-secure WIBE.

1 Introduction

The concept of identity-based encryption (IBE) is a generalization of the standard notion of public-key encryption in which the sender can encrypt messages to a user based only on the identity of the latter and a set of user-independent public parameters. In these systems, there exists a trusted authority, called private key generator, that is responsible for generating decryption keys for all identities in the system. Since being introduced by Shamir in 1984 [28], IBE has received a lot of attention due to the fact that one no longer needs to maintain a separate public key for each user. Despite being an attractive concept, it was only in 2001 that the first practical IBE construction was proposed based on elliptic curve pairings [11]. Later that year, Cocks proposed an alternative IBE construction based on the quadratic residuosity problem [19].

⋆ Work done while at ENS.

M. Fischlin, J. Buchmann, and M. Manulis (Eds.): PKC 2012, LNCS 7293, pp. 316–333, 2012.

The now-standard definition of security of IBE schemes, first suggested by Boneh and Franklin [11], is indistinguishability under adaptive chosen-identity attacks (we refer to it as *full security*). In this security model, the adversary is allowed to obtain secret keys for adaptively chosen identities before deciding the identity upon which it wishes to be challenged. By allowing these queries, this notion implicitly captures resistance against collusion attacks as different users should be unable to combine their keys in an attempt to decrypt ciphertexts intended to another user.

In 2002, Horwitz and Lynn introduced the notion of hierarchical identity-based encryption (HIBE), which allows intermediate nodes to act as private key generators. They also provided a two-level HIBE construction based on the Boneh-Franklin IBE scheme, but their scheme could provide full collusion resistance only in the upper level. The first HIBE scheme to provide full collusion resistance in all levels is due to Gentry and Silverberg [22]. Like the Horwitz-Lynn HIBE scheme, the Gentry-Silverberg HIBE scheme was also based on the Boneh-Franklin IBE scheme and proven secure in the random-oracle model [6].

The first HIBE to be proven secure in the standard model is due to Canetti, Halevi, and Katz [16], but in a weaker security model, called the *selective-identity* model. Unlike the security definitions used in previous constructions of (H)IBE schemes, the selective-identity model requires the adversary to commit to the challenge identity before obtaining the public parameters of the scheme. Despite providing weaker security guarantees, Canetti, Halevi, and Katz showed that the selective-identity model is sufficient for building forward-secure encryption schemes, which was the main motivation of their paper.

Although the selective-identity model has been considered in many works, and is interesting in its own right (e.g., it implies forward-secure public key encryption), if we focus solely on the (H)IBE application, then the selective notion is clearly unrealistic because it does not model the real capabilities of an adversary attacking a (H)IBE scheme. So while the design of selective-identity secure schemes seems to be an easier task, the quest for fully secure solutions is always considered the main goal for (H)IBE construction.

It is therefore a very interesting problem to investigate whether there are ways to efficiently convert a selective secure scheme into a fully secure one. In the random oracle model, this question has been resolved by Boneh, Boyen and Goh [9], who provided a very efficient black-box transformation. In the standard model, however, no such conversion is known[1], and all fully-secure (H)IBE schemes (e.g., [8], [30], [18]) had to be constructed and proved secure essentially from scratch.

Our Results. In this paper, we explore the relationship between selective-identity and fully secure (H)IBE schemes in the standard model.

[1] It was shown by Boneh and Boyen in [7] that any selective secure IBE scheme is already fully secure, but the concrete security degrades by a factor $1/|\mathcal{ID}|$, where $\mathcal{ID}$ is the scheme's identity space. Since $\mathcal{ID}$ is usually of exponential size, this conversion is too expensive in terms of efficiency to be considered practical.

From Selective-Secure WIBE to Fully-Secure HIBE. Our first main contribution is a generic construction of *fully-secure* HIBE schemes from *selective-pattern-secure* wildcarded identity-based encryption (WIBE) schemes. The notion of a WIBE, introduced by Abdalla *et al.* [1], is very similar to the notion of a HIBE except that the sender can encrypt messages not only to a specific identity, but to a whole range of receivers whose identities match a certain pattern defined through a sequence of fixed strings and a special wildcard symbol ($*$). The security notion, called selective-pattern security, requires the adversary to commit ahead of time to the pattern P^* that he intends to attack. He can then ask for the secret keys of any identity not matching P^*, and for the challenge ciphertext on any pattern P matching P^*. This notion of security is slightly more general than that given in [1]. Yet, as noted in Remark 1 at the end of Section 2, it is satisfied by all known WIBE constructions.

Our transformation from *any* selective-pattern-secure WIBE to a fully-secure HIBE is generic and relies on the notion of admissible hash functions (whose existence is implied by collision-resistant hash functions) introduced by Boneh and Boyen in [8]. Since building selective-pattern-secure WIBE schemes seems to be much easier than directly building a fully secure HIBE scheme, this transformation already significantly simplifies the latter task. In fact, it is worth noticing that the selective-pattern security of all currently-known instantiations of WIBE schemes (see [1]) follows from the selective-identity security of their respective underlying HIBE schemes.

One direct consequence of our construction is that several existing fully secure (H)IBE schemes can be seen as a particular case of our transformation. For instance, the fully secure IBE scheme of Boneh and Boyen in [8] turns out to be a particular case of our generic construction when instantiated with the selective-pattern-secure Boneh-Boyen WIBE scheme given in [1]. Likewise, the fully secure HIBE by Cash, Hofheinz, Kiltz, and Peikert [18] can be seen as the result of our generic transformation when applied to our new WIBE scheme in Section 5. Another consequence of our transformation is that one can obtain new constructions of fully secure HIBE schemes by applying our methodology to existing selective-pattern-secure WIBE schemes, such as the Boneh-Boyen-Goh WIBE in [1]. Interestingly, the result obtained from this instantiation closely resembles the Waters (H)IBE scheme [30].

An important point about our transformation from WIBE to (H)IBE is that it also works in the Continual Memory Leakage (CML) model [15,20]. In particular, we show how to modify the IBE scheme in [15] into a WIBE scheme and prove it selective-pattern-secure in the CML model under the same assumption. Then, by applying our transformation to this newly-constructed WIBE, we obtain a (CML) fully-secure version of the IBE in [15]. For lack of space we fully describe these extensions in the full version of our work.

The Role of WIBE in our Transformation. Somewhat surprisingly, our transformation seems to imply that the WIBE notion is of central importance when going from selective to full security in (H)IBE. To see why, one has to take a look at our proof strategy and at the notion of Admissible hash functions

(AHF). AHFs are a tool which allows to partition the identity space into two subsets, $\mathtt{B}$ and $\mathtt{R}$ (both of which are of exponential size) so that in the security proof the identities of secret key queries fall in $\mathtt{B}$ while the challenge identity falls in $\mathtt{R}$. In particular, by carefully selecting the AHFs parameters (as described in [8], for instance) one can make sure that the above (good) event occurs with non-negligible probability. In our proof from selective-secure WIBE to fully-secure HIBE, the simulator first uses AHFs to partition the identity space into $\mathtt{B}$ and $\mathtt{R}$. Next, it declares to the WIBE challenger a challenge pattern which corresponds to $\mathtt{R}$, by expressing $\mathtt{R}$ in the form of a pattern. By the property of AHFs, if the good event occurs (for all key derivation queries and the challenge identity chosen by the adversary), then the simulator can easily forward all queries to the WIBE challenger. In particular, it is guaranteed that the challenge identity falls in $\mathtt{R}$. When that happens, the simulator can output the challenge identity chosen by the adversary as its own challenge.

We remark that the proof strategy described above does not work if one starts from a selective-secure HIBE instead of a WIBE. Unlike the selective-WIBE simulator, the simulator against the selective security of a HIBE should commit to the challenge identity ID^* at the very beginning. And even if the simulator chooses the AHFs parameters so that all secret key queries fall in $\mathtt{B}$ and the challenge identity falls in $\mathtt{R}$, it still needs to guess ID^* in $\mathtt{R}$ at the very beginning. But the probability that the challenge identity chosen by the adversary matches such ID^* is $1/|\mathtt{R}|$, which is negligible (recall that both $\mathtt{B}$ and $\mathtt{R}$ are of exponential size).

Selective WIBE from selective HIBE. The second contribution of this paper is to identify conditions under which we can generically transform a selective-identity-secure HIBE scheme into a selective-pattern-secure WIBE scheme. Towards this goal, we introduce a new notion of security for HIBE schemes, called *security under correlated randomness*, which allows us to transform a given HIBE into a WIBE by simply re-encrypting the same message to a particular set of identities by reusing the same randomness. Informally speaking, in order for a HIBE scheme to be secure under correlated randomness, it must satisfy the following two properties. First, when given an encryption of the same message under the same randomness for two identity vectors $ID_0 = (ID_{0,1}, \ldots, ID_{0,j}, \ldots, ID_{0,\lambda})$ and $ID_1 = (ID_{1,1}, \ldots, ID_{1,j}, \ldots, ID_{1,\lambda})$ differing in exactly one position (say j), one can easily generate a ciphertext for any identity vector matching the pattern $ID = (ID_{1,1}, \ldots, *, \ldots, ID_{1,\lambda})$. Secondly, when given these two ciphertexts, the adversary should not be able to generate an encryption of the same message under the same randomness for any identity vector that does not match the pattern. In Section 4 we show that selective-correlated-randomness-secure HIBE schemes can be converted to selective-pattern-secure WIBEs. Moreover, in the full version, we show that several existing HIBE schemes already satisfy this slightly stronger notion of security, e.g., [7,9,30], and in particular we show that their security under correlated randomness black-box reduces to their selective-identity security.

Hence, if we combine our first generic transformation from selective-pattern-secure WIBE to fully-secure (H)IBE, together with our second result described above, we obtain a compiler that allows us to construct a fully secure (H)IBE starting from a selective-secure (H)IBE. In particular, the resulting transformation works in the standard model and is semi-generic because the second part assumes a specific property of the underlying scheme (i.e., security under correlated randomness). Nevertheless, by reducing the task of building fully secure HIBE schemes to that of building a selective-pattern-secure WIBE scheme, we believe that our result makes the former task significantly easier to achieve.

New WIBE Schemes. One final contribution of this paper are two constructions of selective-pattern-secure WIBE schemes. The first one, whose description is given in the full version of this paper, is obtained by modifying the IBE in [15]. It is based on pairings and is secure under the Decision Linear assumption in the CML model. Such modification essentially follows the correlated-randomness paradigm. Since for some technical reasons (related to the specific scheme) the selective-pattern security of this WIBE cannot be black-box reduced to the selective-identity security of the related IBE (like we do for other pairing-based WIBEs), we give a direct proof under the Decision Linear assumption. However, we notice that such proof closely follows the one in [15]. The second WIBE is based on lattices and its security follows from the selective-identity secure HIBE construction from [18]. Even though the Cash-Hofheinz-Kiltz-Peikert HIBE scheme does not meet the notion of security under correlated randomness introduced in Section 4 (because the scheme is not secure when the same randomness is reused for encryption), we show in Section 5 that one can easily modify it to obtain a selective-pattern-secure WIBE scheme. Similarly to the case of pairing-based WIBE schemes, the selective-pattern security of the new WIBE can be reduced directly to the selective-identity security of the original Cash-Hofheinz-Kiltz-Peikert HIBE scheme. However, in this case, it turns out to be even simpler to prove the selective-pattern security of our scheme directly from the decisional Learning With Errors Problem (`LWE`) [27,26].

Discussion. In this paper, we concentrate on building HIBE schemes that are adaptive-identity-secure against chosen-plaintext attacks. As shown by Boneh, Canetti, Halevi, and Katz [17,13,10], such schemes can easily be made chosen-ciphertext-secure with the help of one-time signature schemes or message authentication codes. Similarly to the (H)IBE schemes by Boneh and Boyen [8], by Waters [30], and by Cash, Hofheinz, Kiltz, and Peikert [18], the schemes obtained via our transformation are only provably secure when the maximum hierarchy's depth L is some fixed constant due to the loss of a factor which is exponential in L. While for lattice-based HIBE schemes [18,3,4], this seems to be the state of the art, the same is not true for pairing-based HIBE schemes. More precisely, there have been several proposals in recent years (e.g., [21,29,25,24]), which are fully secure even when the HIBE scheme has polynomially many levels. Most of these schemes use a new proof methodology, known as dual system encryption [29].

Organization. The paper is organized as follows. In Section 2, we start by recalling some standard definitions and notations used throughout the paper. Next, in Section 3, we present our first main contribution, which is a generic construction which can transform any selective-pattern-secure WIBE into a fully secure HIBE scheme. Then, in Section 4, we introduce the notion of security under correlated randomness for HIBE schemes and show how such schemes can be used to build selective-pattern-secure WIBEs. In Section 5, we show a selective-pattern-secure WIBE scheme that is obtained by transforming the Cash-Hofheinz-Kiltz-Peikert HIBE. Finally, in Section 6, we summarize some future directions left open by our work.

2 Basic Definitions

(Hierarchical) Identity Based Encryption. A *hierarchical identity-based encryption* scheme (HIBE) is defined by a tuple of algorithms $\mathcal{HIBE} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Enc}, \mathsf{Dec})$, a message space $\mathcal{M}$, and an identity space $\mathcal{ID}$. The algorithm Setup is run by a trusted authority to generate a pair of keys (mpk, msk) such that mpk is made public, whereas msk is kept private. The users are hierarchically organized in a tree of depth L whose root is the trusted authority. The identity of a user at level $1 \leq \ell \leq L$ is represented by a vector $\overrightarrow{ID} = (ID_1, \ldots, ID_\ell) \in \mathcal{ID}^\ell$. A user at level ℓ with identity $\overrightarrow{ID} = (ID_1, \ldots, ID_\ell)$ can use the key derivation algorithm $\mathsf{KeyDer}(sk_{\overrightarrow{ID}}, \overrightarrow{ID}')$ to generate a secret key for any of its children $\overrightarrow{ID}' = (ID_1, \ldots, ID_\ell, ID_{\ell+1})$ at level $\ell + 1$. Since this process can be iterated, every user can generate keys for all its descendants. Then, every user holding the master public key mpk, can encrypt a message $m \in \mathcal{M}$ for the identity $\overrightarrow{ID}$ by running $C \xleftarrow{\$} \mathsf{Enc}(mpk, \overrightarrow{ID}, m)$. Finally, the ciphertext C can be decrypted by running the deterministic decryption algorithm, $m \leftarrow \mathsf{Dec}(sk_{\overrightarrow{ID}'}, C)$. For correctness, it is required that for all honestly generated master keys $(mpk, msk) \xleftarrow{\$} \mathsf{Setup}$, for all messages $m \in \mathcal{M}$, all identities $\overrightarrow{ID} \in \mathcal{ID}^\ell$ and all $\overrightarrow{ID}'$ ancestors of $\overrightarrow{ID}$, $m \leftarrow \mathsf{Dec}(\mathsf{KeyDer}(msk, \overrightarrow{ID}'), \mathsf{Enc}(mpk, \overrightarrow{ID}, m))$ holds with overwhelming probability. An IBE is defined as an HIBE with a hierarchy of depth 1.

The security of a HIBE scheme is captured by the standard notion of indistinguishability under chosen-plaintext attacks. Informally, this is captured by the following game. The adversary $\mathcal{A}$ receives as input the master public key and it can ask for the secret key of any identities of its choice. Then it chooses a challenge identity $\overrightarrow{ID}^*$ and two messages m_0 and m_1, and it is given the encryption of m_β under $\overrightarrow{ID}^*$ for a random β. The goal of the adversary is to guess β under the restriction that $\mathcal{A}$ never asks for the secret key of $\overrightarrow{ID}^*$.

In the context of hierarchical identity-based encryption a lot of works in the literature also considered a weaker notion of security, called *selective-identity* indistinguishability under chosen-plaintext attacks (IND-sHID-CPA). The main difference with the standard IND-HID-CPA notion is that here the adversary

is required to commit ahead of time to the challenge identity $\overrightarrow{ID}^*$. The rest of the game is the same as IND-HID-CPA. Sometimes, in order to have a clear distinction with the standard notion of IND-HID-CPA, the latter is called "full security".

Identity Based Encryption with Wildcards. The notion of *Identity-Based Encryption with Wildcards* was introduced by Abdalla *et al.* in [1] as a generalization of the HIBE's notion. A WIBE scheme is defined by a tuple of algorithms $\mathcal{WIBE} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Enc}, \mathsf{Dec})$ that works exactly as a HIBE, except that here the encryption algorithm takes as input a value $P \in (\mathcal{ID} \cup *)^\ell$ (for $1 \leq \ell \leq L$), i.e., the pattern, instead of an identity vector. Such pattern may contain a special "don't care" symbol $*$, the wildcard, at some levels. An identity $\overrightarrow{ID} = (ID_1, \ldots, ID_\ell) \in \mathcal{ID}^\ell$ is said to *match* a pattern $P \in (\mathcal{ID} \cup *)^{\ell'}$, denoted as $\overrightarrow{ID} \in_* P$, if and only if $\ell \leq \ell'$ and $\forall i = 1, \ldots, \ell$: $ID_i = P_i$ or $P_i = *$. Note that under this definition, any ancestor of a matching identity is also a matching identity. This makes sense for the notion of WIBE, as any ancestor can derive the secret key of a matching descendant identity anyway. For any pattern $P \in (\mathcal{ID} \cup *)^\ell$, we denote with $\mathrm{W}(P)$ the set of indices $j \in [\ell]$ such that $P_j = *$. For correctness, it is required that for all honestly generated master keys $(mpk, msk) \xleftarrow{\$} \mathsf{Setup}$, for all messages $m \in \mathcal{M}$, all patterns $P \in (\mathcal{ID} \cup *)^{\ell'}$ and all identities $\overrightarrow{ID} \in \mathcal{ID}^\ell$ such that $\overrightarrow{ID} \in_* P$, $m \leftarrow \mathsf{Dec}(\mathsf{KeyDer}(msk, \overrightarrow{ID}), \mathsf{Enc}(mpk, P, m))$ holds with all but negligible probability.

Similarly to HIBEs, WIBE schemes allow for similar notions of security under chosen-plaintext attacks. In particular, in our work we consider only the notion of selective security that we call IND-sWID-CPA. Roughly speaking, it is similar to the IND-sHID-CPA notion for HIBE, except that here the adversary has to commit to a pattern P^* (instead of an identity $\overrightarrow{ID}^*$) at the beginning of the game. Next, when he has to choose the challenge pattern, he can provide any P that matches P^*, i.e., such that either P is an identity matching P^*, or P is a sub-pattern of P^*.

Remark 1. We notice that our notion of selective-security for WIBE schemes is slightly more general than the one that was originally proposed in [1]. The main difference is that in the original work of Abdalla *et al.* the notion is purely selective, meaning that the adversary declares the challenge pattern P^* at the beginning of the game, and later it receives an encryption of either m_0 or m_1 under P^*. Instead, our notion allows for more flexibility. Indeed, the adversary still declares P^* at the beginning of the game, but later it may ask the challenge ciphertext on a pattern P, possibly different from P^*, but such that P matches P^*. We stress that this property is not artificial for at least two reasons. First, it is more general than the previous one. Second, it is satisfied by all known WIBE schemes, and in particular we will show that it is satisfied by those schemes obtained through our transformation, from selective-secure HIBE to selective WIBE, that we describe in Section 4.

3 Fully-Secure HIBE from Selective-Secure WIBE

In this section we concentrate on the first part of our main result. We show how to construct a fully-secure HIBE scheme starting from any WIBE scheme that is secure only in a selective sense. Our transformation is black-box and makes use of admissible hash functions, a notion introduced by Boneh and Boyen in [8] that we recall below.

Admissible Hash Functions. Admissible hash functions were first introduced by Boneh and Boyen in [8] as a tool for proving the full security of their identity-based encryption scheme in the standard model. Such functions turn out to be particularly suitable for this purpose as they provide a way to implement the so-called "partitioning technique", a proof methodology that allows to secretly partition the identity space into two sets, the *blue* set and the *red* set, both of exponential size, so that there is a non-negligible probability that the adversary's secret key queries fall in the blue set and the challenge identity falls in the red set. This property has been shown useful to prove the full security of some identity-based encryption schemes (e.g., [8,30,18]). In particular, it fits those cases when, in the reduction, one can program the simulator so that it can answer secret key queries for all the blue identities, whereas it is prepared to generate a challenge ciphertext only for red identities.

In our work we employ admissible hash functions for a similar purpose, i.e., constructing a fully-secure HIBE from a selective-secure WIBE, and in particular we adopt a definition of admissible hash functions which follows the one used by Cash *et al.* in [18]. The formal definition follows.

Let $k \in \mathbb{N}$ be the security parameter, w and λ be two values that are at most polynomial in k, and Σ be an alphabet of size s. Let $\mathcal{H} = \{H : \{0,1\}^w \to \Sigma^\lambda\}$ be a family of functions. For $H \in \mathcal{H}$, $K \in (\Sigma \cup \{*\})^\lambda$ and any $x \in \{0,1\}^w$ we define the following function which colors strings in $\{0,1\}^w$ as follows:

$$F_{K,H}(x) = \begin{cases} \mathtt{R} \text{ if } \forall i \in \{1,\ldots,\lambda\} : H(x)_i = K_i \text{ or } K_i = * \\ \mathtt{B} \text{ if } \exists i \in \{1,\ldots,\lambda\} : H(x)_i \neq K_i \end{cases}$$

For any $\mu \in \{0,\ldots,\lambda\}$, we denote with $\mathcal{K}^{(\lambda,\mu)}$ the uniform distribution over $(\Sigma \cup \{*\})^\lambda$ such that exactly μ components are not $*$. Moreover, for every $H \in \mathcal{H}$, $K \in \mathcal{K}^{(\lambda,\mu)}$, and every vector $\boldsymbol{x} \in (\{0,1\}^w)^{Q+1}$ we define the function

$$\gamma(\boldsymbol{x}) = \Pr[F_{K,H}(x_0) = \mathtt{R} \wedge F_{K,H}(x_1) = \mathtt{B} \wedge F_{K,H}(x_2) = \mathtt{B} \wedge \cdots \wedge F_{K,H}(x_Q) = \mathtt{B}].$$

Definition 2. [Admissible Hash Functions] $\mathcal{H} = \{H : \{0,1\}^w \to \Sigma^\lambda\}$ is a family of (Q, δ_{min})-*admissible hash functions* if for every polynomial $Q = Q(k)$, there exists an efficiently computable function $\mu = \mu(k)$, efficiently recognizable sets $bad_H \subseteq (\{0,1\}^w)^*$ and an inverse of a polynomial $\delta_{min} = 1/\delta(k,Q)$ such that the following properties holds:

1. For every PPT algorithm $\mathcal{A}$ that, on input $H \in \mathcal{H}$, outputs $\boldsymbol{x} \in (\{0,1\}^w)^{Q+1}$, there exists a negligible function $\epsilon(k)$ such that:

$$\mathbf{Adv}^{adm}_{\mathcal{H}}(\mathcal{A}) = \Pr[\boldsymbol{x} \in bad_H : H \leftarrow \mathcal{H}, \boldsymbol{x} \leftarrow \mathcal{A}(H)] \leq \epsilon(k)$$

2. For every $H \in \mathcal{H}$, $K \xleftarrow{\$} \mathcal{K}^{(\lambda,\mu)}$, and every vector $\boldsymbol{x} \in (\{0,1\}^w)^{Q+1} \setminus bad_H$ such that $x_0 \notin \{x_1, \ldots, x_Q\}$ we have: $\gamma(\boldsymbol{x}) \geq \delta_{min}$.

Our Transformation. Let $\mathcal{WIBE}$ be a WIBE scheme with identity space $\mathcal{ID} = \Sigma$ of size s and depth $\leq \lambda \cdot L$, and $\mathcal{H} = \{H : \{0,1\}^w \to \Sigma^\lambda\}$ be a family of functions. Then we construct the following HIBE scheme that has identity space $\mathcal{ID}' = \{0,1\}^w$ and depth at most L:

$\mathcal{HIBE}.\mathsf{Setup}$**:** run $(mpk', msk') \xleftarrow{\$} \mathcal{WIBE}.\mathsf{Setup}$ and select $H_1, \ldots, H_L \xleftarrow{\$} \mathcal{H}$. Output $mpk = (mpk', H_1, \ldots, H_L)$ and $msk = msk'$.

$\mathcal{HIBE}.\mathsf{KeyDer}(msk, \overrightarrow{ID})$**:** let $\overrightarrow{ID} = (ID_1, \ldots, ID_\ell)$ and define $\boldsymbol{I} = (H_1(ID_1), \ldots, H_\ell(ID_\ell)) \in \Sigma^{\lambda \cdot \ell}$. Output $sk_{\overrightarrow{ID}} = \mathcal{WIBE}.\mathsf{KeyDer}(msk, \boldsymbol{I})$.

$\mathcal{HIBE}.\mathsf{Enc}(mpk, \overrightarrow{ID}, \mathfrak{m})$**:** let $\overrightarrow{ID} = (ID_1, \ldots, ID_\ell)$ and define $\boldsymbol{I} = (H_1(ID_1), \ldots, H_\ell(ID_\ell)) \in \Sigma^{\lambda \cdot \ell}$. Output $C = \mathcal{WIBE}.\mathsf{Enc}(mpk, \boldsymbol{I}, m)$.

$\mathcal{HIBE}.\mathsf{Dec}(sk_{\overrightarrow{ID}}, C)$**:** return $m = \mathcal{WIBE}.\mathsf{Dec}(sk_{\overrightarrow{ID}}, C)$.

Our scheme is very simple. Essentially, the HIBE algorithm uses the algorithms of the WIBE scheme in a black-box way, where each identity component ID_i is first hashed using a function $H_i \in \mathcal{H}$. Boneh and Boyen show how to construct admissible hash functions based on collision-resistance and error-correction, and propose some concrete parameters for their instantiation (which satisfy our definition). In particular, for convenience of their construction, they consider functions that map to strings in an alphabet Σ of size $s = 2$. Here we notice that if the given WIBE has an alphabet Σ' of size $s' > 2$, then one can simply choose two values $x_1, x_2 \in \Sigma'$, set $\Sigma = \{x_1, x_2\}$, and then consider the same WIBE restricted to these two identities.

The security of our scheme follows from the following theorem, whose proof is deferred to the full version.

Theorem 3. *If $\mathcal{H} = \{H : \{0,1\}^w \to \Sigma^\lambda\}$ is a family of (Q, δ_{min})-admissible hash functions, and $\mathcal{WIBE}$ is* IND-sWID-CPA*-secure, then the scheme $\mathcal{HIBE}$ given in Section 3 is* IND-HID-CPA*-secure, where the maximum hierarchy's depth L is some fixed constant.*

Intuitively speaking, the proof of Theorem 3 proceeds by showing an algorithm $\mathcal{B}$ that plays game IND-sWID-CPA against the scheme $\mathcal{WIBE}$ and simulates the game IND-HID-CPA to an adversary $\mathcal{A}$ against $\mathcal{HIBE}$. $\mathcal{B}$ first generates the parameters for the admissible hash functions, which define the partitions $\mathtt{B}$ and $\mathtt{R}$, and then it declares the set $\mathtt{R}$ as the challenge pattern (notice that by definition of $K \in \mathcal{K}^{(\lambda,\mu)}$, $\mathtt{R}$ can be described using a pattern). In this way, all secret key queries made by $\mathcal{A}$ for identities in $\mathtt{B}$ can be forwarded by $\mathcal{B}$ to its own challenger, and the same can be done if the challenge identity chosen by $\mathcal{A}$ falls in $\mathtt{R}$. In particular, by the properties of admissible hash functions, the event that the identities of secret key queries fall in $\mathtt{B}$ and the challenge identity falls in $\mathtt{R}$ occurs with non-negligible probability. However, things are not that simple, as there may be unlucky events in which $\mathcal{B}$ is unable to simulate the right game to $\mathcal{A}$ and thus it needs to abort. As it already occurred in other

works [30,18], these events may not be independent of the adversary's view, and one solution is to force the simulator to run an expensive artificial abort step. Our proof of Theorem 3 proceeds in this way, requiring $\mathcal{B}$ to (eventually) artificially abort at the end of the simulation. Alternatively, one can extend the techniques introduced by Bellare and Ristenpart in [5] to obtain a proof of Theorem 3 which avoids the need of artificial aborts. However, this requires a slightly different definition of admissible hash functions.

Remark 4. Even though our transformation requires a WIBE scheme with $\lambda \cdot L$ levels to get a HIBE with L levels, we observe that the HIBE key derivation algorithm will use the WIBE key derivation at most L times. The point is that while L is supposed to be a constant, λ can be instead non-constant, as it is the case for known constructions of admissible hash functions, whose output length depends on the number of secret key queries made by the adversary. This might have been a problem for those WIBE schemes that do not support key derivation (delegation) for a polynomial number of levels, such as the new lattice-based scheme described in the full version of this paper.

Extensions. Our transformation easily allows for two extensions. First, it can be used to build an IBE by using a WIBE without the delegation property. Second, we show that it works also in the Continual Memory Leakage model of [15,20]. We provide a complete description of these extensions in the full version of our work.

4 Selective WIBE Schemes from Selective HIBE

In this section we investigate methodologies that allow to build a selective-pattern secure WIBE scheme starting from a HIBE which is selective-identity secure. In particular, we identify conditions under which this transformation works, and then, in the full version we will show that such conditions are satisfied by many known schemes, e.g., [7,9,30]. Then, by combining this result, i.e., a transformation from selective-identity secure HIBE to selective-pattern secure WIBE, with the result of Section 3, i.e., a conversion from selective-pattern secure WIBE to fully-secure HIBE, we obtain a methodology which allows to turn a selective-secure HIBE into a fully-secure one.

Security under Correlated Randomness. Towards this goal, our first contribution is a notion of security for HIBE schemes, called *security under correlated randomness*.The main idea can be described as follows. Assume that one is given encryptions of the same message with the same randomness but for different identities $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$. Then there should be an efficient algorithm that allows to efficiently generate a new ciphertext encrypting the same message but intended to another identity $\overrightarrow{ID}' \in \mathcal{ID}' \subseteq \mathcal{ID}$. The first technical point is to delineate which is this subspace $\mathcal{ID}'$ of the identity space. So, our first contribution is to show that $\mathcal{ID}'$ follows from the differences between the identities

$\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$. More technically, we will show that starting from any set of identities $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$ one can define a matrix Δ whose column i contains the vector which is computed as the difference between $\overrightarrow{ID}^0$ and $\overrightarrow{ID}^i$ (i.e., $\Delta^{(i)} = \overrightarrow{ID}^0 - \overrightarrow{ID}^i$). Then the identity subspace $I\mathcal{D}'$ fixed by $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$ is the set of all identities that can be obtained by making affine operations over $\overrightarrow{ID}^0$ and Δ. (i.e., $\overrightarrow{ID}^0$ plus vectors obtained from integer linear combinations of vectors in Δ). Given this property, encrypting a message with the same randomness for $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$ is equivalent to encrypting for the entire $I\mathcal{D}'$, that we call $\mathsf{Span}(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$. As one may guess, this is already a first step towards building a WIBE, in which the set of recipients of an encryption is actually a subspace of $I\mathcal{D}$ described by the pattern P.

Given the intuitive notion of Span described above, we define below the property for HIBE schemes that we call *Ciphertext Conversion*.

Property 1 (Ciphertext Conversion). *A HIBE scheme satisfies Ciphertext Conversion if there exists an algorithm* $\mathsf{Convert}$ *that, on input* $n+1$ *ciphertexts* $(C_0, \ldots, C_n)$ *encrypting the same message with the same randomness* r*, under identities* $(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$ *respectively, can generate a new ciphertext (encrypting the same message) intended to any* $\overrightarrow{ID} \in \mathsf{Span}(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$.

For any HIBE satisfying Property 1, the notion of selective security under correlated randomness (IND-sCR-CPA) is defined by a game which is the same as the IND-sID-CPA one except that: at the beginning the adversary chooses $n+1$ identities $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$; it receives n+1 challenge ciphertexts generated using the same randomness under identities $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$ respectively; it cannot ask for the secret keys of identities in $\mathsf{Span}(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$. The IND-sCR-CPA notion is parametrized by a distribution $\mathcal{R}$ on the identities $\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n$ that can be chosen by the adversary.

We defer the interested reader to the full version of our work for more formal and precise definitions.

From HIBE Selective-Secure under Correlated Randomness to Selective-Secure WIBE. Now that we have defined the notion of selective security under correlated randomness (IND-sCR-CPA), we can show how to build a selective-pattern secure WIBE from an IND-sCR-CPA-secure HIBE. Towards this goal, let us first introduce some notation and basic definitions.

Let $I\mathcal{D} = \mathbb{Z}_q^\lambda$ be the identity space, for some $q \geq 2$ and $\lambda \geq 1$. For any pattern $P \in (I\mathcal{D} \cup \{*\})^\ell$ we define the function $(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n) \leftarrow F(P)$ as follows. Let $\{j_1, \ldots, j_{n'}\} = \mathrm{W}(P) \subseteq \{1, \ldots, \ell\}$ be the set of levels in which P contains $*$. Let $n = n' \cdot \lambda$, $(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$ is defined as:

$$ID_i^0 = \begin{cases} P_i & if \;\; P_i \neq * \\ 0^\lambda & if \;\; P_i = * \end{cases}$$

$$ID_{i,m}^{k+l-1} = \begin{cases} -1 & if \;\; i = j_k \wedge m = l \\ ID_{i,m}^0 & \text{otherwise} \end{cases} : \begin{array}{l} 1 \leq k \leq n',\ 1 \leq l \leq \lambda \\ 1 \leq i \leq \ell,\ 1 \leq m \leq \lambda \end{array}$$

Moreover, we let $B = \left[B^{(1)}||\cdots||B^{(\ell\lambda)}\right] \in \{0,1\}^{\ell\lambda\times\ell\lambda}$ be the canonical basis of $\mathbb{R}^{\ell\lambda}$.

The function $F(P)$ allows to specify a set of identities $(\overrightarrow{ID}^0,\ldots,\overrightarrow{ID}^n)$ such that the induced subspace $\mathsf{Span}(\overrightarrow{ID}^0,\ldots,\overrightarrow{ID}^n)$ is exactly the same subspace described by the pattern P. Intuitively, this can be seen by looking at the way the identities are defined. $\overrightarrow{ID}^0$ is equal to P on all the positions different from $*$ and 0 elsewhere. Instead each identity $\overrightarrow{ID}^i$ is such that its difference with $\overrightarrow{ID}^0$ leads to a 1 in the *single* position where they differ and 0 elsewhere. Basically, this means that the matrix Δ obtained from $F(P)$ contains a subset of vectors in B. In this way, adding linear combinations of these vectors to $\overrightarrow{ID}^0$ allows to reach identities $\overrightarrow{ID}$ such that $ID_i = P_i$ where $P_i \neq *$, while ID_i can take any value in $I\mathcal{D}$ in those positions i where $P_i = *$. Notice that the number n of such linearly independent vectors strictly depends on the number of $*$ in P. We formally show this property of $F(\cdot)$ by proving the following claim (the proof appears in the full version of our paper):

Claim 5. *For any $P \in (I\mathcal{D}\cup\{*\})^\ell$ and any $\overrightarrow{ID} \in I\mathcal{D}^\ell$ it holds $\overrightarrow{ID} \in \mathsf{Span}(F(P))$ iff $\overrightarrow{ID} \in_* P$.*

Our WIBE Scheme. Let $\mathcal{HIBE} = (\mathsf{Setup}', \mathsf{KeyDer}', \mathsf{Enc}', \mathsf{Dec}', \mathsf{Convert})$ be a HIBE scheme with identity space $I\mathcal{D} = \mathbb{Z}_q^\lambda$ (for $q \geq 2$ and $\lambda \geq 1$), and equipped with an efficient algorithm $\mathsf{Convert}$ satisfying Property 1. Then we construct the scheme $\mathcal{WIBE} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Enc}, \mathsf{Dec})$ as follows.

Setup: Return the output of Setup'.

$\mathsf{KeyDer}(sk_{\overrightarrow{ID}'}, \overrightarrow{ID})$: Run $sk_{\overrightarrow{ID}} \xleftarrow{\$} \mathsf{KeyDer}'(sk_{\overrightarrow{ID}'}, \overrightarrow{ID})$ and output $sk_{\overrightarrow{ID}}$.

$\mathsf{Enc}(mpk, P, m)$: Let $(\overrightarrow{ID}^0,\ldots,\overrightarrow{ID}^n) \leftarrow F(P)$. For all $i = 0$ to n, compute $C_i \xleftarrow{\$} \mathsf{Enc}'(mpk, \overrightarrow{ID}^k, m; r)$, where r is taken at random from the randomness space of $\mathcal{HIBE}.\mathsf{Enc}$. Finally, output $C = (C_0,\ldots,C_n)$.

$\mathsf{Dec}(sk_{\overrightarrow{ID}}, C, P)$: If $\overrightarrow{ID} \notin_* P$, then output $\perp$. Otherwise, compute $(\overrightarrow{ID}^0,\ldots,\overrightarrow{ID}^n) \leftarrow F(P)$, run $C' \leftarrow \mathsf{Convert}(mpk, C_0, \overrightarrow{ID}^0, \ldots, C_n, \overrightarrow{ID}^n, \overrightarrow{ID})$ and then output $m \leftarrow \mathsf{Dec}'(sk_{\overrightarrow{ID}}, C')$.

Remark 6. We notice that our transformation assumes a HIBE scheme that works with the identities returned by our function $F(\cdot)$. This function is defined so that it assigns to the identities values P_i, 0 or -1. However, it may be the case that 0 and/or 1 are not considered valid values in some specific identity space (e.g., assume $I\mathcal{D} = \mathbb{Z}_q \setminus \{0\}$). This issue can be overcome by observing that everything still works if one takes any two different (and valid) values of the identity space, instead of 0 and 1. All we want is that when we compute the matrix Δ, if two identity components are equal, then their difference becomes 0, otherwise they lead to some value c (not necessarily 1). To see that everything works even with any constant c, observe that it is possible to consider our operations over Δ/c.

Now, we state the security of our scheme via the following theorem, whose proof can be found in the full version.

Theorem 7. *If $\mathcal{HIBE}$ satisfies Property 1 and is* IND-sCR-CPA*-secure w.r.t. $\mathcal{R} = \mathcal{ID}^{\ell\times(n+1)}$, then the scheme $\mathcal{WIBE}$ described above is correct and* IND-sWID-CPA *secure.*

A Sufficient Distribution for Building a WIBE. In the previous section, we showed that an HIBE scheme satisfying Property 1 and the notion of selective-security under correlated randomness can be transformed into a WIBE. In particular, Theorem 7 considers the most general definition where the distribution $\mathcal{R}$ is arbitrary, i.e., $\mathcal{R} = \mathcal{ID}^{\ell\times(n+1)}$. However, we observe that in order for the transformation to work, it is sufficient to consider a more restricted distribution that we call $\mathcal{R}_{WIBE}$.

Let $B = \left[B^{(1)}||\cdots||B^{(\ell\lambda)}\right] \in \{0,1\}^{\ell\lambda\times\ell\lambda}$ be the canonical basis. defined in the previous section. We define the distribution

$$\mathcal{R}_{WIBE} = \{(\overrightarrow{ID}^0,\ldots,\overrightarrow{ID}^n) : \overrightarrow{ID}^0 \in \mathbb{Z}_q^{\lambda\ell},\ \overrightarrow{ID}^i = \overrightarrow{ID}^0 + k_i \cdot B^{(j_i)},\ 1 \le i \le n,$$

$$j_i \in \{1,\ldots,\lambda\ell\}, \boldsymbol{k} \in \mathbb{Z}^n\}$$

It is interesting to observe that for any pattern P the identities obtained from $F(P)$ follow the distribution $\mathcal{R}_{WIBE}$. We show the following claim whose proof appears in the full version.

Claim 8. *For any pattern $P \in (\mathcal{ID} \cup \{*\})^\ell$ we have $F(P) \in \mathcal{R}_{WIBE}$.*

Hence, we can combine the results of Theorem 7 and Claim 8 to obtain the following Corollary.

Corollary 9. *If $\mathcal{HIBE}$ satisfies Property 1 and is secure under the* IND-sCR-CPA *notion w.r.t. $\mathcal{R}_{WIBE}$, then the scheme $\mathcal{WIBE}$ described above is correct and* IND-sWID-CPA*-secure.*

5 Lattice-Based WIBE

In this section, we give a construction of a lattice-based selectively-secure WIBE, based on the hardness of the LWE Problem [27], that very closely resembles the selectively-secure HIBE construction from [18]. In fact, the master/secret key generation and delegation procedures are exactly the same for the HIBE and the WIBE. The only difference lies in the encryption and decryption procedures; yet even there, the distinction is fairly minor. For the benefit of those readers familiar with the HIBE of [18], we present the constructions of the WIBE along with the construction of the HIBE and also try to use the same notational conventions.

Algorithms Used in Constructing the HIBE and WIBE. We now briefly describe the algorithms that were used in [18] to construct the HIBE, which we will be using in this section for constructing the WIBE.

1. $\texttt{GenBasis}(1^n, 1^m, q)$: This algorithm generates a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ (where $m = \Omega(n \log q)$) and a basis $\mathbf{S} \in \mathbb{Z}^{m \times m}$ of $\Lambda^{\perp}(\mathbf{A})$ such that the distribution of $\mathbf{A}$ is negligibly close to uniform over $\mathbb{Z}_q^{n \times m}$ and $\|\tilde{\mathbf{S}}\| = O(\sqrt{n \log q})$.
2. $\texttt{ExtBasis}(\mathbf{S}, \mathbf{A}' = \mathbf{A}||\bar{\mathbf{A}})$: This algorithm takes as input a matrix $\mathbf{A}' = \mathbf{A}||\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times (m+\bar{m})}$ and a matrix $\mathbf{S} \in \mathbb{Z}^{m \times m}$, which is basis of $\Lambda^{\perp}(\mathbf{A})$, and outputs a matrix $\mathbf{S}' \in \mathbb{Z}^{(m+\bar{m}) \times (m+\bar{m})}$ that is a basis of $\Lambda^{\perp}(\mathbf{A}')$ such that $\|\tilde{\mathbf{S}}\| = \|\widetilde{\mathbf{S}'}\|$.
3. $\texttt{SampleD}(\mathbf{S}, \mathbf{A}, \mathbf{y}, s)$: This algorithm takes as input a basis $\mathbf{S} \in \mathbb{Z}^{m \times m}$ of the lattice $\Lambda^{\perp}(\mathbf{A})$, a vector $\mathbf{y} \in \mathbb{Z}_q^n$, and a real number $s \geq \|\tilde{\mathbf{S}}\| \cdot \omega(\sqrt{\log n})$ and outputs a vector $\mathbf{z} \sim D_{\Lambda_{\mathbf{y}}^{\perp}(\mathbf{A}), s}$.
4. $\texttt{RandBasis}(\mathbf{S}, s)$: This algorithm takes as input an $m \times m$ lattice basis $\mathbf{S}$ and a real number $s \geq \|\tilde{\mathbf{S}}\| \cdot \omega(\sqrt{\log n})$, and outputs a basis $\mathbf{S}'$ of the same lattice such that $\|\mathbf{S}'\| \leq s\sqrt{m}$. Furthermore, if $\mathbf{S_0}, \mathbf{S_1}$ are bases of the same lattice and $s > \max\{\|\widetilde{\mathbf{S_0}}\|, \|\widetilde{\mathbf{S_1}}\|\}$, then the distributions of $\texttt{RandBasis}(\mathbf{S_0}, s)$ and $\texttt{RandBasis}(\mathbf{S_1}, s)$ are statistically close.

The Lattice-Based WIBE Scheme. We now describe the master key generation, key derivation, encryption and decryption algorithms of our WIBE scheme. For any distribution χ over $\mathbb{Z}$, and any vector $\mathbf{x} \in \mathbb{Z}_q^n$ let $\texttt{Noisy}_{\chi}(\mathbf{x})$ be the distribution obtained by first creating a vector $\mathbf{y} \in \mathbb{Z}^n$ each of whose coordinates is independently sampled according to χ, and then outputting $\mathbf{x} + \mathbf{y} \bmod q$.

Master Key Generation. We assume that the identities are of the form $\{0,1\}^t$, for all $t \leq L$. The generation of the master public and secret keys is done exactly in the same fashion in the HIBE and in the WIBE. The WIBE is parametrized by the integers n, m, q where n is the security parameter, m is an integer of size $\Omega(n \log q)$ and q is some prime whose size is related to the number of allowable key derivations, and is discussed in Section 5. We first run the $\texttt{GenBasis}(1^n, 1^m, q)$ procedure to obtain a matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{n \times m}$ and a basis $\mathbf{S}_0 \in \mathbb{Z}^{m \times m}$ of $\Lambda^{\perp}(\mathbf{A})$. Then for all $(i, j) \in \{0,1\} \times \{1, \ldots, L\}$, we generate a uniformly random matrix $\mathbf{A}_j^{(i)} \in \mathbb{Z}_q^{n \times m}$, and choose a uniformly-random $\mathbf{y} \in \mathbb{Z}_q^n$. The master public key is

$$\left[\mathbf{A}_0, \mathbf{A}_1^{(0)}, \mathbf{A}_1^{(1)}, \ldots, \mathbf{A}_L^{(0)}, \mathbf{A}_L^{(1)}, \mathbf{y}\right],$$

and the master secret key is $\mathbf{S}_0$.

Key Derivation. The key derivation procedure is again performed exactly the same for the HIBE and the WIBE. The public key of identity $id = (id_1, \ldots, id_t)$ is $(\mathbf{A}_{id}, \mathbf{y})$, where $\mathbf{A}_{id} = \mathbf{A}_0 \| \mathbf{A}_1^{(id_1)} \| \ldots \| \mathbf{A}_t^{(id_t)}$. The secret key of user id is $(\mathbf{S}_{id}, \mathbf{x}_{id})$ where $\mathbf{S}_{id}$ is a "short" basis of the lattice $\Lambda^{\perp}(\mathbf{A}_{id})$ and $\mathbf{x}_{id}$ is a "short" vector satisfying $\mathbf{A}_{id}^T \mathbf{x}_{id} = \mathbf{y}$. The matrix $\mathbf{S}_{id}$ will be used for delegation, while the vector $\mathbf{x}_{id}$ will be used for decryption.

If a user with $id = (id_1, \ldots, id_t)$ would like to generate a secret key for a user $id' = (id_1, \ldots, id_t, id_{t+1}, \ldots, id_{t'})$ whose public key is $(\mathbf{A}_{id'} = \mathbf{A}_{id} || \bar{\mathbf{A}}, \mathbf{y})$, where

$\bar{\mathbf{A}} = \mathbf{A}_{t+1}^{(id_{t+1})} \| \ldots \| \mathbf{A}_{t'}^{(id_{t'})}$, he computes the following:

$$\mathbf{S}_{id'} \leftarrow \texttt{RandBasis}(\texttt{ExtBasis}(\mathbf{S}_{id}, \mathbf{A}_{id'}), s)$$
$$\mathbf{x}_{id'} \leftarrow \texttt{SampleD}(\texttt{ExtBasis}(\mathbf{S}_{id}, \mathbf{A}_{id'}), \mathbf{A}_{id'}, \mathbf{y}, s)$$

where $s \geq \|\widetilde{\mathbf{S}_{id}}\| \cdot \omega(\sqrt{\log n})$. We point out that with every key derivation, the value of $\|\widetilde{\mathbf{S}_{id}}\|$ increases by a factor of $\tilde{O}(\sqrt{n})$. When the norm of the secret key gets too large, decryption becomes impossible, and so, just like in [18], it is important to adjust the ratio of the size of the secret key $\mathbf{S}_0$ and the prime q based on how many levels of delegations one wishes to have. With each level of delegation increasing the norm of the user id by a factor of $\tilde{O}(\sqrt{n})$, the ratio between $\|\widetilde{\mathbf{S}_0}\|$ and q should be on the order of $\sqrt{n}^d$, where d is the maximum allowable levels in the hierarchy. Since the $\texttt{LWE}_{n,q,\chi}$ problem becomes easier as q gets larger (and the distribution χ stays the same), there is a trade-off between security and the maximum number of delegation levels. We direct the reader to [18] for the precise parameters.

Encryption and Decryption. In the HIBE, encryption of a message $\kappa \in \{0,1\}$ is performed to identity $id = (id_1, \ldots, id_t)$ by picking a random $\mathbf{r} \in \mathbb{Z}_q^n$ and outputting the pair $(\mathbf{u}_{id}, v) \in \mathbb{Z}_q^{m(t+1)+1}$, where

$$(\mathbf{u}_{id}, v) = \left(\texttt{Noisy}_\chi\left(\mathbf{A}_{id}^T \mathbf{r}\right), \texttt{Noisy}_\chi\left(\mathbf{y}^T \mathbf{r} + \kappa \cdot \lfloor q/2 \rfloor\right)\right)$$

where

$$\mathbf{A}_{id} = \mathbf{A}_0 \| \mathbf{A}_1^{(id_1)} \| \ldots \| \mathbf{A}_t^{(id_t)} \tag{1}$$

and χ is some "narrow" distribution such that the $\texttt{LWE}_{n,q,\chi}$ problem is hard.

The decryption of the HIBE ciphertext by the identity $id = (id_1, \ldots, id_t)$ is performed as follows: for a ciphertext $(\mathbf{u}_{id}, v)$ and secret key $\mathbf{x}_{id}$, the algorithm computes $v - \mathbf{x}_{id}^T \mathbf{u}_{id} \bmod q$ and outputs 0 if the result is closer to 0 than to $q/2$, and outputs 1 otherwise.

In our WIBE, encryption is defined in essentially the same way as in the HIBE. To encrypt to a pattern $pat = (pat_1, \ldots, pat_t) \in \{0, 1, *\}^t$, we pick a random $\mathbf{r} \in \mathbb{Z}_q^n$, define

$$\mathbf{A}_{pat} = \mathbf{A}_0 \| \mathbf{A}_1^{(pat_1)} \| \ldots \| \mathbf{A}_t^{(pat_t)} \tag{2}$$

where $\mathbf{A}_i^* := \mathbf{A}_i^{(0)} \| \mathbf{A}_i^{(1)}$, and output the pair $(\mathbf{u}_{pat}, v) \in \mathbb{Z}_q^{m(t+t_*+1)+1}$ (where t_* is the number of $*$ in the pattern pat),

$$(\mathbf{u}_{pat}, v) = \left(\texttt{Noisy}_\chi(\mathbf{A}_{pat}^T \mathbf{r}), \texttt{Noisy}_\chi\left(\mathbf{y}^T \mathbf{r} + \kappa \cdot \lfloor q/2 \rfloor\right)\right).$$

Notice that instead of the matrix $\mathbf{A}_{pat}$ being $n \times mt$ as in the HIBE, it can be as large as $n \times 2mt$ because every position pat_i that contains the wildcard $*$ results in the concatenation of both $\mathbf{A}_i^{(0)}$ and $\mathbf{A}_i^{(1)}$ into the matrix $\mathbf{A}_{pat}$. Therefore the ciphertext of the WIBE could be twice as large as the HIBE ciphertext.

The decryption procedure of the WIBE is also very similar to that of the HIBE. For every $id = (id_1, \ldots, id_t) \in \{0,1\}^t$, the matrix $\mathbf{A}_{pat}$ contains the

matrix $\mathbf{A}_{id}$, where $\mathbf{A}_{id}$ is defined as in (1). Therefore, since we know $\mathbf{u}_{pat} = \mathtt{Noisy}_{\chi}(\mathbf{A}_{pat}^T\mathbf{r})$, we can retrieve from it $\mathbf{u}_{id} = \mathtt{Noisy}_{\chi}(\mathbf{A}_{id}^T\mathbf{r})$. And now, using the secret key $\mathbf{x}_{id}$, the user can decrypt the ciphertext $(\mathbf{u}_{id}, v)$ the same way as in the HIBE scheme by computing $v - \mathbf{x}_{id}^T\mathbf{u}_{id} \bmod q$ and outputting 0 if the result is closer to 0 than to $q/2$, and 1 otherwise.

Security. The security proof of our scheme, which can be found in the full version of this paper, is a simple adaptation of the HIBE security proof in [18].

Theorem 10. *Given an adversary $\mathcal{A}$ who breaks the WIBE with parameters n, m, q allowing d key derivations, there exists an algorithm $\mathcal{S}$ that solves the $\mathtt{LWE}_{n,q,\chi}$ problem where $q > \sigma \cdot n^{d/2} \cdot poly(n)$ where σ is the standard deviation of the distribution χ and $poly(n)$ is some fixed polynomial function in n.*

6 Future Directions

First, in its most general form (i.e., without restrictions on $\mathcal{R}$), our notion of security under correlated randomness gives a generic methodology for encrypting messages to sets S of recipients that are defined by $\mathsf{Span}(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$. In this sense, a WIBE can be seen as a special case of this notion in which the recipients' sets always have a fixed form specified by the pattern P, i.e., $S = \mathsf{Span}(F(P))$. However, one may think of a more general notion in which these sets can have a more "irregular" form that we can express using a set of identities $(\overrightarrow{ID}^0, \ldots, \overrightarrow{ID}^n)$ and its Span.

Since we were mostly interested in building WIBE schemes in this work, we considered security under correlated randomness w.r.t. the distribution $\mathcal{R}_{WIBE}$. However, as a future direction, it would be interesting to explore whether there exist HIBE schemes that are IND-sCR-CPA-secure according to the most generic notion, i.e., without any restriction on $\mathcal{R}$. Perhaps more interestingly, the resulting primitive could be seen as the dual version of the notion of Spatial Encryption proposed by Boneh and Hamburg in [12]. In Spatial Encryption, ciphertexts are associated to points in $\mathbb{Z}_p^\ell$, while secret keys correspond to affine subspaces of $\mathbb{Z}_p^\ell$. In this setting, a ciphertext for $x \in \mathbb{Z}_p^\ell$ can be decrypted by any secret key for $W \subseteq \mathbb{Z}_p^\ell$ as long as $x \in W$. In contrast, our new notion would consider ciphertexts that are associated to affine subspaces of $\mathcal{ID}^\ell$.

As a second direction, it would be interesting to investigate whether our techniques can be applied to other cryptographic primitives. Indeed, the problem of selective vs. full security has already been considered in the context of other cryptographic notions, such as attribute-based encryption or verifiable random functions (VRFs). In the particular case of VRFs, finding a fully secure scheme has been a long standing open problem until the very recent works by Hohenberger and Waters [23] and by Boneh *et al.* [14]. In fact, both of these works can be seen as obtaining a fully secure VRF from a selective secure one. While the work of Boneh *et al.* explicitly builds a selective-secure VRF and then turns it into a fully secure one, the work of Hohenberger and Waters can be interpreted as a fully secure version of the selective-secure VRF scheme of Abdalla *et al.* [2].

Acknowledgments. This work was supported in part by the European Research Council, in part by the European Commission through the ICT Program under Contract ICT-2007-216676 ECRYPT II, and in part by the French ANR-10-SEGI-15 PRINCE Project.

References

1. Abdalla, M., Catalano, D., Dent, A.W., Malone-Lee, J., Neven, G., Smart, N.P.: Identity-Based Encryption Gone Wild. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006, Part II. LNCS, vol. 4052, pp. 300–311. Springer, Heidelberg (2006)
2. Abdalla, M., Catalano, D., Fiore, D.: Verifiable Random Functions from Identity-Based Key Encapsulation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 554–571. Springer, Heidelberg (2009)
3. Agrawal, S., Boneh, D., Boyen, X.: Efficient Lattice (H)IBE in the Standard Model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg (2010)
4. Agrawal, S., Boneh, D., Boyen, X.: Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98–115. Springer, Heidelberg (2010)
5. Bellare, M., Ristenpart, T.: Simulation without the Artificial Abort: Simplified Proof and Improved Concrete Security for Waters' IBE Scheme. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 407–424. Springer, Heidelberg (2009)
6. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, pp. 62–73. ACM Press (November 1993)
7. Boneh, D., Boyen, X.: Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 223–238. Springer, Heidelberg (2004)
8. Boneh, D., Boyen, X.: Secure Identity Based Encryption Without Random Oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (2004)
9. Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical Identity Based Encryption with Constant Size Ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)
10. Boneh, D., Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. SIAM Journal on Computing 36(5), 1301–1328 (2007)
11. Boneh, D., Franklin, M.K.: Identity based encryption from the Weil pairing. SIAM Journal on Computing 32(3), 586–615 (2003)
12. Boneh, D., Hamburg, M.: Generalized Identity Based and Broadcast Encryption Schemes. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 455–470. Springer, Heidelberg (2008)
13. Boneh, D., Katz, J.: Improved Efficiency for CCA-Secure Cryptosystems Built Using Identity-Based Encryption. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 87–103. Springer, Heidelberg (2005)
14. Boneh, D., Montgomery, H.W., Raghunathan, A.: Algebraic pseudorandom functions with improved efficiency from the augmented cascade. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM CCS 2010, pp. 131–140. ACM Press (October 2010)

15. Brakerski, Z., Kalai, Y.T., Katz, J., Vaikuntanathan, V.: Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In: 51st FOCS, pp. 501–510. IEEE Computer Society Press (2010)
16. Canetti, R., Halevi, S., Katz, J.: A Forward-Secure Public-Key Encryption Scheme. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 255–271. Springer, Heidelberg (2003)
17. Canetti, R., Halevi, S., Katz, J.: Chosen-Ciphertext Security from Identity-Based Encryption. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222. Springer, Heidelberg (2004)
18. Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai Trees, or How to Delegate a Lattice Basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer, Heidelberg (2010)
19. Cocks, C.: An Identity Based Encryption Scheme Based on Quadratic Residues. In: Honary, B. (ed.) Cryptography and Coding 2001. LNCS, vol. 2260, pp. 360–363. Springer, Heidelberg (2001)
20. Dodis, Y., Haralambiev, K., López-Alt, A., Wichs, D.: Cryptography against continuous memory attacks. In: 51st FOCS, pp. 511–520. IEEE Computer Society Press (2010)
21. Gentry, C., Halevi, S.: Hierarchical Identity Based Encryption with Polynomially Many Levels. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 437–456. Springer, Heidelberg (2009)
22. Gentry, C., Silverberg, A.: Hierarchical ID-Based Cryptography. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 548–566. Springer, Heidelberg (2002)
23. Hohenberger, S., Waters, B.: Constructing Verifiable Random Functions with Large Input Spaces. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 656–672. Springer, Heidelberg (2010)
24. Lewko, A., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully Secure Functional Encryption: Attribute-Based Encryption and (Hierarchical) Inner Product Encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 62–91. Springer, Heidelberg (2010)
25. Lewko, A., Waters, B.: New Techniques for Dual System Encryption and Fully Secure HIBE with Short Ciphertexts. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 455–479. Springer, Heidelberg (2010)
26. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Mitzenmacher, M. (ed.) 41st ACM STOC, pp. 333–342. ACM Press (May/June 2009)
27. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC, pp. 84–93. ACM Press (May 2005)
28. Shamir, A.: Identity-Based Cryptosystems and Signature Schemes. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (1985)
29. Waters, B.: Dual System Encryption: Realizing Fully Secure IBE and HIBE under Simple Assumptions. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg (2009)
30. Waters, B.: Efficient Identity-Based Encryption Without Random Oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)