Verifiable Random Functions from Identity-Based Key Encapsulation*

Michel Abdalla[1], Dario Catalano[2,**], and Dario Fiore[2]

[1] CNRS–LIENS, Ecole Normale Supérieure, Paris, France
michel.abdalla@ens.fr
[2] Dipartimento di Matematica e Informatica, Università di Catania, Italy
{catalano,fiore}@dmi.unict.it

Abstract. We propose a methodology to construct verifiable random functions from a class of identity based key encapsulation mechanisms (IB-KEM) that we call VRF suitable. Informally, an IB-KEM is VRF suitable if it provides what we call *unique decryption* (i.e. given a ciphertext C produced with respect to an identity ID, all the secret keys corresponding to identity ID', decrypt to the same value, even if $ID \neq ID'$) and it satisfies an additional property that we call pseudorandom decapsulation. In a nutshell, pseudorandom decapsulation means that if one decrypts a ciphertext C, produced with respect to an identity ID, using the decryption key corresponding to any other identity ID' the resulting value looks random to a polynomially bounded observer. Interestingly, we show that most known IB-KEMs already achieve pseudorandom decapsulation. Our construction is of interest both from a theoretical and a practical perspective. Indeed, apart from establishing a connection between two seemingly unrelated primitives, our methodology is *direct* in the sense that, in contrast to most previous constructions, it avoids the inefficient Goldreich-Levin hardcore bit transformation.

1 Introduction

Verifiable Random Functions (VRFs for short) were introduced by Micali, Rabin and Vadhan [21]. Informally a VRF is something that behaves like a random function but also allows for efficient verification. More precisely, this means that associated with a secret key sk (the seed), there is a public key pk and a function F such that the following properties are satisfied. First, the function is efficiently computable, given sk, on any input. Second, having only pk and oracle access to the function, the value $F_{pk}(x) = y$ looks random to any polynomially bounded observer who did not query $F_{pk}(x)$ explicitly. Third, a proof $\pi_{pk}(x)$ that $F_{pk}(x) = y$ is efficiently computable knowing sk and efficiently verifiable knowing only pk.

VRFs turn out to be very useful in a variety of applications essentially because they can be seen as a compact commitment to an exponential number of

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-3-642-01001-9_35

* The full version of this paper is available at http://www.dmi.unict.it/~fiore

** Work partially done while visiting the computer science department at Ecole Normale Supérieure.

A. Joux (Ed.): EUROCRYPT 2009, LNCS 5479, pp. 554–571, 2009.

(pseudo)random bits. To give a few examples, Micali and Reyzin [22] show how to use VRFs to reduce to 3 the number of rounds of resettable zero knowledge proofs in the bare model. Micali and Rivest [23] described a very simple non interactive lottery system used in micropayment schemes, based on VRFs. Jarecki and Shmatikov [17] employed VRFs to build a verifiable transaction escrow scheme that preserves users anonymity while enabling automatic de-escrow. Liskov [18] used VRFs to construct updatable Zero Knowledge databases. In spite of their popularity VRFs are not very well understood objects. In fact, as of today, only four constructions are known, in the standard model [21,20,11,13]. The schemes given in [21,20] build VRFs in two steps. First they focus on constructing a *Verifiable Unpredictable Function* (VUF). Informally a VUF is a function that is hard to compute but whose produced outputs do not necessarily look random. Next they show how to convert a VUF into a VRF using the Goldreich-Levin [15] theorem to "extract" random bits. Unfortunately, however, the VRF resulting from this transformation is very inefficient and, furthermore, it looses a quite large factor in its exact security reduction. This is because, the transformation involves several steps, all rather inefficient. First one uses the Goldreich Levin theorem [15] to construct a VRF with very small (i.e. slightly super polynomial in the security parameter) input space and output size 1. Next, one iterates the previous step in order to amplify the output size to (roughly) that of the input. Then, using a tree based construction, one iterates the resulting function in order to get a VRF with unrestricted input size and finally one evaluates the so obtained VRF several times in order to get an output size of the required length.

The constructions given in [11,13], on the other hand, are direct, meaning with this that they manage to construct VRF without having to resort to the Goldreich Levin transform. The VRF presented in [11] is based on a "DDH-like" assumption that the author calls *sum-free decisional Diffie-Hellman* (sf-DDH). This assumption is similar to that one employed by Naor-Reingold [24] to construct PRFs, with the difference that it applies an error correcting code C to the input elements in order to compute the function. The specific properties of the employed encoding allow for producing additional values that can be used as proofs. This construction is more efficient than [21,20] in the sense that it does not need the expensive Goldreich Levin transform. Still it has some efficiency issues as the size of the produced proofs and keys is linear in the input size. Dodis [11] also adapts this construction to provide a *distributed* VRF, that is a standard VRF which can be computed in a distributed manner.

The scheme proposed by Dodis and Yampolskiy [13], on the other hand, is more attractive, at least from a practical point of view, as it provides a simple implementation of VRFs with short (i.e. constant size) proofs and keys. It is interesting to note that, even though the latter construction is far more efficient than previous work, it builds upon a similar approach. Basically, the construction in [13] works in two steps. First they consider a simple VUF (which is basically Boneh Boyen [3] weakly secure signature scheme) that is secure for slightly superpolynomially sized input spaces. Next, rather than resorting to the Godreich Levin [15] hardcore bit theorem to convert it into a VRF, they show

how to modify the original VUF in order to make it a VRF, under an appropriate decisional assumption.

From the discussion above, it seems clear that, with the possible exception of [11], all known constructions of verifiable random functions, follow similar design criteria. First one builds a suitable VUF and then transforms it into a VRF by either using the Goldreich Levin transform, or via some direct, ad hoc, modifications of the original VUF. The main drawback of this approach is that, once a good enough VUF is found, one has to either be able to make it a VRF directly or accept the fact that the VRF obtained from the Goldreich Levin transform is not going to be a practical one. Thus it seems very natural to ask if there are alternative (and potentially more efficient) ways that allow to construct VRFs directly, without needing to resort to the two steps methodology sketched above.

Our Contribution. In this paper we show how to construct VRF from a class of identity based encryption (IBE) schemes [26] that we call *VRF suitable*. Roughly speaking an identity based encryption scheme, is an asymmetric encryption scheme where the public key can be an arbitrary string. Such schemes consists of four algorithms. A Setup algorithm, that generates the system common parameters as well as a master key msk; a key derivation algorithm that uses the master secret key to generate a private key d_{sk} corresponding to an arbitrary public key string ID (the identity); an encryption algorithm that encrypts messages using the public key ID and a decryption algorithm that decrypts ciphertexts using the corresponding private key.

Informally an IBE is said to be VRF suitable if the following conditions are met. First, the scheme has to provide *unique decryption*. This means that, given a ciphertext C produced with respect to some arbitrary identity ID, all the secret keys corresponding to any other identity ID' decrypt to the same value (i.e. even if $ID' \neq ID$). Second, the Key Encapsulation Mechanism (KEM) associated with the IBE (see below for a definition of key encapsulation mechanism) has to provide what we call *pseudorandom decapsulation*. Very informally, pseudorandom decapsulation means that if C is an encapsulation produced using some identity ID, the "decapsulated" key should look random even if the decapsulation algorithm is executed using the secret key corresponding to any other identity $ID^* \neq ID$. Having a scheme that achieves pseudorandom decapsulation may seem a strong requirement at first. We argue that it is not, as basically all currently known secure (in the standard model) IBE schemes *already* provide pseudorandom decapsulation.

Our result is of interest both from a theoretical and a practical point of view. Indeed, apart from establishing a connection between two seemingly unrelated primitives, our method is direct, in the sense that it allows to build a VRF from a VRF suitable IBE without having to resort to the inefficient Goldreich Levin transform. Moreover, the reduction is tight. This means that, once an efficient VRF suitable IBE is available, this leads to an equally efficient VRF, with no additional security loss. Furthermore, our constructions immediately allow

for efficient distributed VRFs as long as a distributed version of the underlying encryption scheme is available (which is the case for most schemes used in practice).

As a second contribution of this paper, we investigate on the possibility of implementing VRF suitable IBEs. Toward this goal, we first show how to construct a VRF suitable IB KEM from the Sakai-Kasahara IB KEM [25]. Interestingly, the resulting VRF turns out to be very similar to the Dodis-Yampolskiy VRF [13], thus showing that the latter construction can actually be seen as a special case of our general methodology. Next, we propose a new implementation of VRF suitable IB KEM inspired (but more efficient) by Lysyanskaya's VRF [20] (which in turn builds from the Naor Reingold's PRF [24]). The proposed scheme can be proved secure under the assumed intractability, in bilinear groups, of the decisional ℓ-th weak Bilinear Diffie Hellman Inversion problem (decisional ℓ-wBDHI* for short) introduced by Boneh, Boyen and Goh [4]. Interestingly, even though the decisional ℓ-wBDHI* assumption is asymptotic in nature, the ℓ parameter does not need to be too large in order for our security proof to go through. This is because it directly affects only the size of the space of valid identities *but not* the number of adversarial queries allowed in the security reduction[1] (as opposed to most known proofs using asymptotic assumptions). This means that in practice it is enough to assume the decisional ℓ-wBDHI* assumption to hold only for reasonably small values of ℓ (such as $\ell = 160$ or $\ell = 256$).

IBEs and Digital Signatures. Naor pointed out (see [5]) that a fully secure identity based encryption scheme can be transformed into a secure signature scheme as follows. One sets the message space as the set I of valid identities of the IBE. To sign $m \in I$ one executes the key derivation algorithm on input m, and outputs d_{sk} as the signature. A signature on m is verified by encrypting a random M with respect to the identity m, and then by checking that decrypting the resulting ciphertext one gets back M. Thus if one considers an IBE with unique key derivation (i.e. where for each identity one single corresponding decryption key can be computed) the methodology sketched above leads to a secure *unique* digital signature scheme (i.e. a digital signature scheme for which each message admits one single valid signature). Since secure unique signatures are, by definition, verifiable unpredictable functions, at first glance our construction might seem to (somewhat) follow from Naor's remark. We argue that this does not seem to be the case for two reasons. First, our construction does not require the underlying IB-KEM to have unique key derivation, but only to provide unique decryption. Clearly the former property implies the latter, but there is no reason to exclude the possibility of constructing a scheme realizing unique decryption using a randomized key derivation procedure. Second, a crucial requirement for Naor's transformation to work is that the original IBE is actually fully secure. A VRF-suitable IBE, on the other hand, is required to be secure only in a much weaker sense (that we call *weak selective* ID security).

[1] Here by not affecting the number of adversarial queries we mean that ℓ grows linearly with respect to the identity space but only logarithmically with respect to the number of adversarial queries.

Other related Work. As pointed out above the notion of VRF is related to the notion of unique signatures. Unique signatures were introduced by Goldwasser and Ostrovsky [16] (they called them invariant signatures). The only known constructions of unique signatures in the plain model (i.e. without common parameters or random oracles) are due to Micali, Rabin and Vadhan [21], to Lysyanskaya [20] and to Boneh and Boyen [3]. In the common string model, Goldwasser and Ostrovsky [16] also showed that unique signatures require the same kind of assumptions needed to construct non interactive zero knowledge.

Dodis and Puniya in [12] address the problem of constructing Verifiable Random Permutations from Verifiable Random Functions. They define VRPs as the verifiable analogous of pseudorandom permutations. In particular they point out that the technique of Luby-Rackoff [19] (for constructing PRPs from PRFs) cannot be applied in this case. This is due to the fact that VRP proofs must reveal the VRF outputs and proofs of the intermediate rounds. In their paper they show a construction in which a super-logarithmic number of executions of the Feistel transformation suffices to build a VRP.

More recently Chase and Lysyanskaya [8] introduced the notion of simulatable VRF. Informally a simulatable VRF is a VRF with the additional property that proofs can be simulated, meaning with this that a simulator can fake proofs showing that the value of $F_{sk}(x)$ is y for any y of its choice. Simulatable VRFs can be used to provide a direct transformation from single theorem non interactive zero knowledge to multi theorem NIZK and work in the common reference string model.

2 Preliminaries

Before presenting our results we briefly recall some basic definitions. In what follows we will denote with k a security parameter. The participants to our protocols are modeled as probabilistic Turing machines whose running time is bounded by some polynomial in k. Denote with $\mathbb{N}$ the set of natural numbers and with $\mathbb{R}^+$ the set of positive real numbers. We say that a function $\epsilon : \mathbb{N} \to \mathbb{R}^+$ is negligible if and only if for every polynomial $P(k)$ there exists an $k_0 \in \mathbb{N}$ such that for all $k > k_0$ $\epsilon(k) < 1/P(k)$. If A is a set, then $a \stackrel{\$}{\leftarrow} A$ indicates the process of selecting a at random and uniformly over A (which in particular assumes that A can be sampled efficiently).

Verifiable Random Functions Verifiable Random Functions (VRFs for short) were introduced by Micali, Rabin and Vadhan [21]. Intuitively, a VRF is something that behaves like a pseudorandom function, but also allows for a proof of its output correctness. More formally, a VRF is a triplet of algorithms $\mathsf{VRF} = (\mathsf{Gen}, \mathsf{Func}, \mathsf{V})$ providing the following functionalities. The key generation algorithm Gen is a probabilistic algorithm that takes as input the security parameter and produces a couple of matching public and private keys (vpk, vsk). The deterministic algorithm Func, on input the secret key vsk and the input x to the VRF, computes $(\mathsf{F}_{vsk}(x), \mathsf{Prove}_{vsk}(x))$, where $\mathsf{F}_{vsk}(x)$ is the value of the VRF and $\mathsf{Prove}_{vsk}(x)$ its proof of correctness. The verification algorithm V takes

as input (vpk, x, v, π) and outputs a bit indicating whether or not π is a valid proof that $\mathsf{F}_{vsk}(x) = v$.

Let $a : \mathbb{N} \to \mathbb{N} \cup \{*\}$ and $b : \mathbb{N} \to \mathbb{N}$ be functions computable in polynomial time (in k). Moreover we assume that $a(k)$ and $b(k)$ are bounded by a polynomial in k, except if a takes the value $*$ (in this case we simply assume that the VRF can take inputs of arbitrary length). Formally, we say that $\mathsf{VRF} = (\mathsf{Gen}, \mathsf{Func}, \mathsf{V})$ is a VRF of input length $a(k)$ and output length $b(k)$, if the following conditions are met.

Domain Range Correctness. For all $x \in \{0,1\}^{a(k)}$ it has to be the case that $\mathsf{F}_{vsk}(x) \in \{0,1\}^{b(k)}$. We require this condition to hold with overwhelming probability (over the choices of (vpk, vsk)).

Provability. For all $x \in \{0,1\}^{a(k)}$ if $\mathsf{Prove}_{vsk}(x) = \pi$ and $\mathsf{F}_{vsk}(x) = v$ then $\mathsf{V}(vpk, x, v, \pi) = 1$. We require this condition to hold with overwhelming probability (over the choices of (vpk, vsk) and the coin tosses of V).

Uniqueness. No values $(vpk, x, v_1, v_2, \pi_1, \pi_2)$ can satisfy (unless with negligible probability over the coin tosses of V) $\mathsf{V}(vpk, x, v_1, \pi_1) = \mathsf{V}(vpk, x, v_2, \pi_2) = 1$, when $v_1 \neq v_2$.

Pseudorandomness. For all probabilistic polynomial time adversaries $A = (A_1, A_2)$ we require that

$$\Pr\left[b' = b \;\middle|\; \begin{array}{l} (vpk, vsk) \xleftarrow{\$} \mathsf{Gen}(1^k);\ (x, \omega) \leftarrow A_1^{\mathsf{Func}(\cdot)}(vpk) \\ b \xleftarrow{\$} \{0,1\};\ v_0 \leftarrow \mathsf{F}_{vsk}(x);\ v_1 \xleftarrow{\$} \{0,1\}^{b(k)} \\ b' \leftarrow A_2^{\mathsf{Func}(\cdot)}(\omega, v_b) \end{array} \right] \leq \frac{1}{2} + \epsilon(k)$$

where the notation $A^{\mathsf{Func}(\cdot)}$ indicates that A has oracle access to the algorithm Func. In order to make this definition sensible, we impose that A cannot query the oracle on input x.

Remark 1. One might consider a relaxation of the pseudorandomness property in which the adversary is required to commit ahead of time (i.e. before seeing the public key) to the input value it intends to attack. We call *selective-VRF* a VRF that satisfies this weaker pseudorandomness[2].

ID BASED ENCRYPTION An identity based encryption scheme is a tuple of algorithms $\mathsf{IBE} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Enc}, \mathsf{Dec})$ providing the following functionality. The trusted authority runs Setup, on input 1^k, to generate a master key pair (mpk, msk). Without loss of generality we assume that the public key mpk specifies a message space $\mathcal{M}$ and a value n (polynomial in the security parameter) indicating the length of each identity. It publishes the master public key mpk and keeps the master secret key msk private. When a user with identity ID wishes to become part of the system, the trusted authority distributor generates a user decryption key $d_{ID} \xleftarrow{\$} \mathsf{KeyDer}(msk, ID)$, and sends this key over a secure and authenticated channel to the user. To send an encrypted message m to the user

[2] For lack of space we defer a more formal definition of this notion to the full version of this paper.

with identity ID, the sender computes the ciphertext $C \xleftarrow{\$} \mathsf{Enc}(mpk, ID, m)$, which can be decrypted by the user as $m \leftarrow \mathsf{Dec}(d_{ID}, C)$.

Boneh and Franklin [5] formally defined the notion of security for identity based encryption schemes. In particular they defined chosen plaintext security against adaptive chosen identity attack. Intuitively, such a notion, captures the requirement that security should be preserved even when facing an adversary who is allowed to choose the identity it wishes to attack. Later, Canetti, Halevi, and Katz [7] introduced a weaker notion of security in which the adversary is required to commit ahead of time (i.e. before the parameters of the scheme are made public) to the identity it intends to attack. A scheme meeting such a weaker security requirement is said selective ID, chosen plaintext secure IBE (IND-sID-CPA).

In this paper we introduce a new notion of security for IBE schemes that we call *weak selective ID security*. More precisely, we define weak selective ID security as the full fledged selective case with the exception that here the challenge identity is chosen by the challenger and given in input to the adversary. Clearly, this notion is weaker with respect to selective ID security as it is easy to see that the latter implies the former.

Identity Based Key Encapsulation. An identity-based key encapsulation mechanism (IB-KEM) scheme allows a sender and a receiver to agree on a random session key K in such a way that the sender can create K from public parameters and receiver identity and the receiver can recover K using his secret key. This notion, in the context of identity-based encryption, was first formalized by Bentahar et al. [1].

An IB-KEM scheme is defined by four algorithms:

- $\mathsf{Setup}(1^k)$ is a probabilistic algorithm that takes in input a security parameter k and outputs a master public key mpk and a master secret key msk.
- $\mathsf{KeyDer}(msk, ID)$ The key derivation algorithm uses the master secret key to compute a secret key sk_{ID} for identity ID.
- $\mathsf{Encap}(mpk, ID)$ The encapsulation algorithm computes a random session key K and a corresponding ciphertext C encrypted under the identity ID.
- $\mathsf{Decap}(C, sk_{ID})$ allows the possessor of a secret key sk_{ID} to decapsulate C to get back a session key K. We denote by $\mathcal{K}$ the session key space.

For correctness it is required that $\forall k \in \mathbb{N}, ID \in \mathcal{ID}, (C, K) \xleftarrow{\$} \mathsf{Encap}(mpk, ID)$ the following probability holds for all possible $(mpk, msk) \xleftarrow{\$} \mathsf{Setup}(1^k)$:

$$Pr[\mathsf{Decap}(C, \mathsf{KeyDer}(msk, ID)) = K] = 1$$

Here we define the notion of *weak selective ID security* for IB-KEM schemes. Let $\mathcal{IBKEM}$ be a IBE scheme with key encapsulation mechanism. Then $\mathcal{IBKEM}$ is *weakly selective ID secure against adaptively-chosen plaintext attacks* (wsIB-KEM-CPA) if there exists no polynomially bounded adversary $\mathcal{A}$ with non negligible advantage against the Challenger in the following game:

Setup. In this phase the challenger selects a challenge identity ID^* (according to an arbitrary distribution) and runs $(mpk, msk) \leftarrow \mathsf{Setup}(1^k)$. Then it

computes $(C^*, K^*) = \mathsf{Encap}(mpk, ID^*)$ and flips a binary coin $b \xleftarrow{\$} \{0,1\}$. Then it sets $\bar{K} = K^*$ if $b = 0$, otherwise it picks a random key $\bar{K} \xleftarrow{\$} \mathcal{K}$. Finally it runs $\mathcal{A}$ on input $(mpk, ID^*, C^*, \bar{K})$ and keeps msk for itself.

Key derivation queries. The adversary is allowed to ask key derivation queries for an arbitrary (but polynomial) number of adaptively chosen identities different from ID^*.

Guess. In the end of this game $\mathcal{A}$ outputs b' as its guess for b.

The adversary wins if $b' = b$. We formally define the advantage of $\mathcal{A}$ against $\mathcal{IBKEM}$ in the above game as

$$\mathbf{Adv}^{\text{wsIB-KEM-CPA}}_{\mathcal{IBKEM},\mathcal{A}}(k) = \left| Pr[b = b'] - \frac{1}{2} \right|$$

where the probability is taken over the coin tosses of the challenger and the adversary.

VRF-suitable IB-KEMs. Our VRF construction relies on a special class of identity based key encapsulation mechanisms that we call *VRF suitable*. In particular, a VRF suitable IB-KEM is defined by the following algorithms

- $\mathsf{Setup}(1^k)$ is a probabilistic algorithm that takes in input a security parameter k and outputs a master public key mpk and a master secret key msk.
- $\mathsf{KeyDer}(msk, ID)$ The key derivation algorithm uses the master secret key to compute a secret key sk_{ID} for identity ID and some auxiliary information aux_{ID} needed to correctly encapsulate and decapsulate the key.
- $\mathsf{Encap}(mpk, ID, aux_{ID})$ The encapsulation algorithm computes a random session key K, using (mpk, ID, aux_{ID}). Moreover it uses (mpk, ID) to computes a ciphertext C encrypted under the identity ID. Notice that aux_{ID} is required to compute K but not to compute C.
- $\mathsf{Decap}(C, sk_{ID}, aux_{ID})$ allows the possessor of sk_{ID} and aux_{ID} to decapsulate C to get back a session key K. We denote by $\mathcal{K}$ the session key space.

Remark 2. Notice that the description above differs from the one given for basic IB-KEM in that here we require the encapsulation and decapsulation mechanism to use some auxiliary information aux_{ID}, produced by KeyDer, to work correctly. Clearly if one sets $aux_{ID} = \perp$ one goes back to the original description. Thus the new paradigm is slightly more general as it allows to consider encapsulation mechanism where everybody can compute the ciphertext but only those knowing the aux_{ID} information can compute the key. Notice however that aux_{ID} does not allow, by itself, to decapsulate. In some sense, this auxiliary information should be seen as a value that completes the public key (rather than something that completes the secret key) [3]. Even though such a syntax may look totally meaningless in the standard public key scenario, it turns out to be extremely useful (see below) in our context.

[3] In fact this auxiliary information is not required to be kept secret in our constructions since the adversary can in principle obtain its value for any identity of its choice including the challenge identity (see definition of pseudorandom decapsulation).

Moreover, the IB-KEM has to satisfy the following properties:

1. **Unique decryption.** Let ID_0 be any valid identity and C a ciphertext encrypted under ID_0. We require that no valid identity ID can satisfy (unless with negligible probability) $\mathsf{Decap}(C, sk'_{ID}, aux_{ID}{}') \neq \mathsf{Decap}(C, sk''_{ID}, aux_{ID}{}'')$, where $(sk'_{ID}, aux_{ID}{}') \leftarrow \mathsf{KeyDer}(msk, ID)$, $(sk''_{ID}, aux_{ID}{}'') \leftarrow \mathsf{KeyDer}(msk, ID)$
2. **Pseudorandom decapsulation.** Let C be an encapsulation produced using identity ID_0, we require the session key to look random even if the decapsulation algorithm is executed using the secret key corresponding to any other $\overline{ID}$. More formally, we define the following experiment, for a polynomially bounded adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$.

Experiment $\mathbf{Exp}^{\text{IB-KEM-RDECAP}}_{\mathcal{IBKEM},\mathcal{A}}(k)$
$(mpk, msk) \xleftarrow{\$} \mathsf{Setup}(1^k)$
Choose $ID_0 \in \mathcal{ID}$ (according to any arbitrary distribution)
$C^* \xleftarrow{\$} \mathsf{Encap}(mpk, ID_0)$
$(\overline{ID}, st) \xleftarrow{\$} \mathcal{A}_1^{\mathsf{KeyDer}(\cdot)}(mpk, C^*, ID_0)$
$(aux_{\overline{ID}}, sk_{\overline{ID}}) \xleftarrow{\$} \mathsf{KeyDer}(msk, \overline{ID})$
$b \xleftarrow{\$} \{0,1\}$; $K_0 \xleftarrow{\$} \mathsf{Decap}(C^*, sk_{\overline{ID}}, aux_{\overline{ID}})$; $K_1 \xleftarrow{\$} \mathcal{K}$
$b' \leftarrow \mathcal{A}_2^{\mathsf{KeyDer}(\cdot)}(st, K_b, aux_{\overline{ID}})$
If $b' = b$ then return 1, else return 0

With $\mathcal{A}^{\mathsf{KeyDer}(\cdot)}$ we denote that an algorithm $\mathcal{A}$ has oracle access to the key derivation algorithm. Let $\mathcal{ID}$ denote identity space, i.e. the space from which the adversary (and everybody else) is allowed to choose the identities. In the experiment $\mathbf{Exp}^{\text{IB-KEM-RDECAP}}_{\mathcal{IBKEM},\mathcal{A}}$ we need the following restrictions:

- the identity $\overline{ID}$ output by $\mathcal{A}_1$ should not be asked before;
- $\mathcal{A}_2$ is not allowed to query the oracle on $\overline{ID}$.

We define the advantage of $\mathcal{A}$ in the IB-KEM-RDECAP experiment as

$$\mathbf{Adv}^{\text{IB-KEM-RDECAP}}_{\mathcal{IBKEM},\mathcal{A}}(k) = \left| Pr\left[\mathbf{Exp}^{\text{IB-KEM-RDECAP}}_{\mathcal{IBKEM},\mathcal{A}}(k) = 1\right] - \frac{1}{2}\right|.$$

$\mathcal{IBKEM}$ has pseudorandom decapsulation if for any polynomially bounded adversary $\mathcal{A}$ the advantage $\mathbf{Adv}^{\text{IB-KEM-RDECAP}}_{\mathcal{IBKEM},\mathcal{A}}(k)$ is a negligible function in k.

Remark 3. Requiring that an IB-KEM provides pseudorandom decapsulation might seem a very strong requirement at first. We argue that it is not, at least if the known constructions of IB-KEMs are considered. Indeed, all currently known schemes which are IND-CPA secure (but not IND-CCA secure) in the standard model *already* have this property (see the full version of the paper for details).

3 The Construction

In this section we show our construction of Verifiable Random Functions from a VRF-suitable IB-KEM $\mathcal{IBKEM} = (\mathsf{Setup}, \mathsf{KeyDer}, \mathsf{Encap}, \mathsf{Decap})$. Let $\mathcal{ID}$ be the identity space, $\mathcal{K}$ the session key space and $\mathcal{SK}$ the secret key space. Then we construct $\mathsf{VRF} = (\mathsf{Gen}, \mathsf{Func}, \mathsf{V})$ which models a function from input space $\mathcal{ID}$ to output space $\mathcal{K}$.

$\mathsf{Gen}(1^k)$ runs $(mpk, msk) \leftarrow \mathsf{Setup}(1^k)$, chooses an arbitrary identity $ID_0 \in \mathcal{ID}$ and computes $C_0 \leftarrow \mathsf{Encap}(mpk, ID_0)$. Then it sets $vpk = (mpk, C_0)$ and $vsk = msk$.

$\mathsf{Func}_{vsk}(x)$ computes $\pi_x = (sk_x, aux_x) = \mathsf{KeyDer}(msk, x)$ and $y = \mathsf{Decap}(C_0, \pi_x)$. It returns (y, π_x) where y is the output of the function and π_x is the proof.

$\mathsf{V}(vpk, x, y, \pi_x)$ first checks if π_x is a valid proof for x in the following way. It computes $(C, K) = \mathsf{Encap}(mpk, x, aux_x)$ and checks if $K = \mathsf{Decap}(C, \pi_x)$. Then it checks the validity of y by testing if $\mathsf{Decap}(C_0, \pi_x) = y$. If both the tests are true, then the algorithm returns 1, otherwise it returns 0.

Now we prove that the proposed construction actually realizes a secure VRF.

Theorem 1. *Assume $\mathcal{IBKEM}$ is a VRF Suitable IB-KEM scheme, as described in section 2 then the construction given above is a verifiable random function.*

Proof. According to the definition given in section 2, we prove that $\mathsf{VRF} = (\mathsf{Gen}, \mathsf{Func}, \mathsf{V})$ is a verifiable random function by showing that it satisfies all the properties. Domain range correctness and provability trivially follow from the IB-KEM scheme correctness. Since $\mathcal{IBKEM}$ has unique decryption the uniqueness property is satisfied for construction of VRF. To prove the residual pseudorandomness we assume there exists an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that breaks the residual pseudorandomness of VRF with non-negligible probability $\frac{1}{2} + \epsilon(k)$. Then we can build an adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ which has non-negligible advantage $\epsilon(k)$ in the IB-KEM-RDECAP game.

$\mathcal{B}_1$ receives in input from its Challenger the public key mpk and a ciphertext C_0^*. It sets $vpk = (mpk, C_0^*)$ and runs $\mathcal{A}_1(vpk)$. The adversary $\mathcal{A}$ is allowed to make queries to the function oracle $\mathsf{Func}(\cdot)$. $\mathcal{B}$ simulates this oracle in the following way. Given input $x \in \mathcal{ID}$, it queries the key derivation oracle on x. It obtains sk_x and returns $(\mathsf{Decap}(C_0^*, sk_x), sk_x)$ to the adversary. When $\mathcal{A}_1$ outputs an element $\bar{x}$, $\mathcal{B}_1$ gives the same element to its challenger. Thus the challenger produces K^*, which is either the decapsulation of C_0^* with $sk_{\bar{x}}$ or a random element of $\mathcal{K}$, and gives it to $\mathcal{B}_2$. Then $\mathcal{B}_2$ runs $b' \leftarrow \mathcal{A}_2(st, K^*)$ and outputs b' to the Challenger.

Since the simulation is perfect, if $\mathcal{A}$ outputs $b' = b$ with probability $\frac{1}{2} + \epsilon(k)$, then $\mathcal{B}$'s advantage is exactly $\epsilon(k)$.

Notice that, when describing the notion of VRF suitable IB-KEM, we did not expect the underlying scheme to meet any additional security requirement. With the following theorem (whose proof is deferred to the full version of this paper) we show that, indeed, a necessary condition, in order for an IB-KEM to be VRF suitable, is that it is secure only in a weak selective sense.

Theorem 2. *Let $\mathcal{IBKEM}$ be a VRF Suitable IB-KEM, then it is also a weakly selective secure IB-KEM (in the sense of the definition given in section 2).*

4 VRF Suitable IBEs

In this section we describe our constructions of Verifiable Random functions from VRF suitable encryption schemes. In particular, in light of the results presented in section 3, we focus on constructing VRF suitable IB-KEM schemes.

We start by describing, in section 4.1, a VRF from the Sakai-Kasahara [25] IB-KEM. Interestingly, the proposed VRF is basically the same as the VRF proposed by Dodis and Yampolskiy [13], thus showing that their construction can be seen as a special case of our general paradigm.

Next we present, in section 4.2, a new construction of VRF suitable IB-KEM from an assumption related to the ℓ-Bilinear Diffie Hellman Inversion assumption (see [2]), that is known as the decisional ℓ-weak Bilinear Diffie Hellman Inversion assumption (decisional ℓ-wBDHI*, following the acronym used in [4]). The decisional ℓ-wBDHI* was introduced by Boneh, Boyen and Goh in [4] and it (informally) states that given $g^b, g^c, g^{b^2}, ..., g^{b^\ell}$, the quantity $e(g,g)^{b^{\ell+1}c}$ should remain indistinguishable from random to any polynomially bounded adversary. The assumption is related to the ℓ bilinear Diffie Hellman Inversion assumption (ℓ-BDHI), in the sense that the former is known to hold in all those groups where the latter holds, but the converse is not known to be true. Interestingly, in order for our construction to work, the ℓ parameter does not need to be too large. This is because it only limits to 2^ℓ the size of the space of valid identities but it does not affect in any other way the number of adversarial queries allowed in the security proof (as in most known proofs using q-type assumptions). Said in a different way, ℓ is required to grow only in a logarithmic way (rather than linearly) with respect to the number of adversarial queries allowed. This means that it is enough to assume that the ℓ-wBDHI* assumption holds only for rather small values of ℓ (i.e. $\ell = 160$ or $\ell = 256$).

As a final note we mention that, in principle, one could construct a VRF from Boneh Franklin's IBE. Indeed, in the full version of this paper, we show that the KEM version of the scheme is actually a VRF suitable IB-KEM, under the decisional Bilinear Diffie Hellman assumption. This construction, however, is of very limited interest, since the proof holds in the random oracle model.

4.1 Sakai-Kasahara VRF

We briefly recall the KEM version of the Sakai-Kasahara IBE scheme (`SK` for short) [25]. This scheme relies on the q-decisional Bilinear Diffie-Hellmann Inversion assumption (DBDHI for short), introduced by Boneh and Boyen in [2]. Informally, the DBDHI assumption in bilinear group G of prime order p states that, for every PPT algorithm $\mathcal{A}$ and for a parameter q, $\mathcal{A}$ has negligible probability into distinguishing $e(g,g)^{1/x} \in G_T$ from a random one after seeing $(g, g^x, g^{(x^2)}, \cdots, g^{(x^q)})$. If we suppose that $\mathcal{G}(1^k)$ is a bilinear group generator

which takes in input a security parameter k, then (asymptotically) the DBDHI assumption holds for $\mathcal{G}$ if $\mathcal{A}$'s probability of success is negligible in k, for any q polynomial in k.

- $\mathsf{Setup}(1^k)$ runs $\mathcal{G}(1^k)$ to obtain the description of the groups G, G_T and of a bilinear map $e : G \times G \to G_T$. The description of G contains a generator $g \in G$. Then the algorithm picks a random $s \xleftarrow{\$} \mathbb{Z}_p$ and sets $h = g^s$, $mpk = (g, h)$, $msk = s$.
- $\mathsf{KeyDer}(msk, ID)$ We assume $ID \in \mathbb{Z}_p$. The key derivation algorithm constructs the secret key $sk_{ID} = g^{\frac{1}{s+ID}}$.
- $\mathsf{Encap}(mpk, ID)$ The encapsulation algorithm picks a random $t \xleftarrow{\$} \mathbb{Z}_p^\star$ and computes a random session key $K = e(g, g)^t$ and a corresponding ciphertext $C = (g^s g^{ID})^t$.
- $\mathsf{Decap}(C, sk_{ID})$ the decapsulation algorithm uses the secret key sk_{ID} to compute a session key K from a ciphertext C as follows: $K = e(C, sk_{ID})$.

First notice that, assuming $aux_{ID} = \perp \ \forall ID$, the above description fits our syntax of VRF suitable IB-KEMs. Now we prove (for lack of space the actual proof appears in the full version of this paper) that the Sakai-Kasahara IB-KEM scheme can be used to construct a VRF (i.e. that it actually provides unique decryption and pseudorandom decapsulation). In particular, the resulting VRF can only support superpolynomially-sized (in the security parameter) input space. Notice that all known constructions of VRF made the same assumption.

Theorem 3. *Assuming that the DBDHI assumption holds in a bilinear group G, then the Sakai-Kasahara IBE scheme [25] is a VRF-suitable IBE.*

Similarity with the Dodis-Yampolskiy VRF. Now we show that the Dodis-Yampolskiy VRF [13] (that we briefly recall in appendix A) can be seen as a special instantiation of the construction given above. Indeed, theorem 3 leads to the following VRF.

- $\mathtt{Gen}(1^k)$ Runs $\mathcal{G}(1^k)$ to obtain the description of the groups G, G_T and of a bilinear map $e : G \times G \to G_T$. The description of G contains a generator $g \in G$. Then the algorithm picks random $s, t \xleftarrow{\$} \mathbb{Z}_p$ and sets $h = g^s$, $C_0 = h^t$, $vpk = (g, h, C_0)$, $vsk = s$.
- $\mathtt{Func}_{vsk}(x)$ Let $\mathtt{Func}_{vsk}(x)$=$(F_{vsk}(x), \pi_{vsk}(x))$. One sets $\mathtt{Func}_{vsk}(x)$= $e(C_0, sk_x)$ $= e(g, g)^{(st)/(s+x)}$ as the VRF output and $\pi_{vsk}(x) = \mathsf{KeyDer}(x) = g^{1/(s+x)}$ as the proof of correctness.
- $\mathtt{V}(vpk, x, y, \pi_x)$ To verify whether y was computed correctly, one starts by running the Encap algorithm on input (vpk, x). Encap chooses $\omega \xleftarrow{\$} \mathbb{Z}_p$ and then computes $K \leftarrow e(g, g)^\omega$ and $C = (hg^x)^\omega$. Then one checks that $K = \mathsf{Decap}(C, \pi_x) = e((g^x \cdot h)^\omega, \pi_x)$ and $y = \mathsf{Decap}(C_0, \pi_x) = e(h^t, \pi_x)$.

Thus by setting $t = s^{-1} \bmod p$ and $\omega = 1$, the construction above can be optimized to get exactly the Dodis Yampolskiy VRF.

4.2 The New Construction

In this section we propose a new construction of VRF suitable IB-KEM from the (conjectured) computational intractability of the decisional weak ℓ-Bilinear Diffie-Hellman Inversion problem (see below for a formal description). The new scheme is inspired from Lysyanskaya [20] VRF in that the validity of each new auxiliary information aux_{ID} (required to compute the session key) is verified by exploiting the DDH-CDH separation in bilinear groups. The new scheme however is more efficient as it leads to a VRF directly (i.e. rather than having to construct a unique signature scheme first) and does not require error correcting codes.

Decisional weak ℓ-Bilinear Diffie Hellman Inversion Problem [4]. The decisional ℓ-wBDHI* problem in G is defined as follows. Let G be a bilinear group of prime order p and g a generator of G. Given $g^b, g^c, g^{b^2}, ..., g^{b^\ell}$, we say that an algorithm $\mathcal{A}$ has advantage ϵ in solving decisional ℓ-wBDHI* in G if

$$\Pr[\mathcal{A}(g^c, g^b, g^{b^2}, ..., g^{b^\ell}, e(g,g)^{b^{\ell+1}c}) = 1] - \Pr[\mathcal{A}(g^c, g^b, g^{b^2}, ..., g^{b^\ell}, e(g,g)^z) = 1] \geq \epsilon$$

where the probability is over the random choices of $b, c, z \in \mathbb{Z}_p^*$

We say that the decisional ℓ-wBDHI* assumption holds in G if no polynomially bounded adversary has advantage better than negligible in solving the decisional ℓ-wBDHI* problem in G.

Remark 4. Cheon showed in [9] an attack against the Strong Diffie-Hellman Assumption and its related problems (among which the DBDHI used to prove the security of the Dodis-Yampolskiy VRF). This attack reduces the security of a factor $\sqrt{q}$ and applies to the ℓ-wBDHI* as well. However, as it is stated at the beginning of this section, in our construction it is enough to assume that the ℓ-wBDHI* assumption holds only for rather small values of ℓ (i.e. $\ell = 160$ or $\ell = 256$). Thus in our case the security loss is not significant as in Dodis-Yampolskiy's.

The proposed scheme follows

$\mathsf{Setup}(1^k)$ runs $\mathcal{G}(1^k)$ to obtain the description of the groups G, G_T and of a bilinear map $e : G \times G \to G_T$. The description of G contains a generator $g \in G$. Let $\{0,1\}^\ell$ be the space of valid identities. Then the algorithm picks (at random) $a, \alpha_1, \beta_1, \ldots, \alpha_\ell, \beta_\ell \stackrel{\$}{\leftarrow} \mathbb{Z}_p$, sets $g_1 = g^a$, and for $i = 1, \ldots, \ell$ sets $g_{0i} = g^{\beta_i}$ and $g_{1i} = g^{\alpha_i}$. The public parameters are

$$mpk = \left(g, g_1, \{g_{ij}\}_{i=0,1;j=1..\ell}\right)$$

The master secret key is $msk = (a, \{\alpha_i, \beta_i\}_{i=1,..,\ell})$

$\mathsf{KeyDer}(msk, ID)$ We assume $ID = ID_1 \cdots ID_\ell$ where each $ID_i \in \{0,1\}$. The key derivation algorithm constructs the secret key sk_{ID} and the auxiliary information aux_{ID} as follows. Let $h_0 = g$, for $i = 1$ to ℓ one computes

$$h_i = (h_{i-1})^{\alpha_i^{ID_i} \beta_i^{(1-ID_i)}}$$

and sets $aux_{ID} = (h_1, \ldots, h_\ell)$ and $sk_{ID} = h_\ell^a$.

$\mathsf{Encap}(mpk, ID, aux_{ID})$ Let $aux_{ID} = (h_1, \ldots, h_\ell)$ computed as above, the encapsulation algorithm picks a random $t \xleftarrow{\$} \mathbb{Z}_p^\star$ and computes a random session key $K = e(g_1, h_\ell)^t$ and a corresponding ciphertext $C = g^t$.

$\mathsf{Decap}(C, sk_{ID}, aux_{ID})$ the decapsulation algorithm uses the secret key sk_{ID} and the auxiliary information aux_{ID} to compute a session key K from a ciphertext C. This is done as follows. First, in order to guarantee the unique decryption property, a check on the validity of the auxiliary information has to be performed. This is done as follows, let $h_0 = g$, for $i = 1, \ldots, \ell$

$$\begin{aligned} \textbf{if } ID_i = 1 \textbf{ check } & e(g, h_i) \stackrel{?}{=} e(g_{1i}, h_{i-1}) \\ \textbf{else} \quad \textbf{check } & e(g, h_i) \stackrel{?}{=} e(g_{0i}, h_{i-1}) \end{aligned}$$

If any of the above checks fails output reject. Second, the key K is computed as $K = e(C, sk_{ID}) = e(g_1, h_\ell)^t$ Notice that, the validity of sk_{ID} can be verified by first encrypting some random message m with respect to the public key (g, g_1, h_ℓ) and then by checking if one can decrypt it correctly using sk_{ID}.

Now we prove the following result

Theorem 4. *Suppose the decisional ℓ-wBDHI* assumption holds in G, then the scheme given above is a secure VRF suitable IB-KEM scheme.*

Proof. Let $\mathcal{ID} = \{0,1\}^\ell$ the identity space. First notice that the scheme fits the syntax of VRF suitable IB-KEMs. We prove the theorem by showing that the scheme has the unique decryption property and meets the pseudorandom decapsulation requirement.

Unique Decryption. We prove this by showing that for a given identity ID the corresponding h_ℓ is uniquely determined as

$$h_\ell = g^{\prod_{i=1}^{\ell} \alpha_i^{ID_i} \beta_i^{1-ID_i}}$$

The proof is by induction on i. First notice that it must be the case $h_1 = g^{\alpha_1^{ID_1} \beta_1^{1-ID_1}}$, as otherwise the check $e(g, h_1) \stackrel{?}{=} e(g_{ID_1 1}, h_0) = e(g^{\alpha_1^{ID_1} \beta_1^{1-ID_1}}, g)$ would fail. Now assume that the statement holds true for any index $j - 1 < \ell$, i.e. that $h_{j-1} = g^{\prod_{i=1}^{j-1} \alpha_i^{ID_i} \beta_i^{1-ID_i}}$. We prove that the same holds for j.

$$h_j = h_{j-1}^{\alpha_j^{ID_j} \beta_j^{1-ID_j}} = \left(g^{\prod_{i=1}^{j-1} \alpha_i^{ID_i} \beta_i^{1-ID_i}} \right)^{\alpha_j^{ID_j} \beta_j^{1-ID_j}} = g^{\prod_{i=1}^{j} \alpha_i^{ID_i} \beta_i^{1-ID_i}}$$

Pseudorandom Decapsulation. Assume that there is an an adversary $\mathcal{A}$ that breaks the pseudorandom decapsulation of the proposed scheme with advantage ϵ, we show how to build an adversary $\mathcal{B}$ that solves the decisional ℓ-wBDHI* problem with advantage $\epsilon/2^\ell$ and runs in time comparable to that needed by $\mathcal{A}$. $\mathcal{B}$ starts by receiving, from some challenging oracle, the values

($C = g^c, B_1 = g^b, B_2 = g^{b^2}, \ldots B_\ell = g^{b^\ell}$ and a value T that can be either of the form $e(g,g)^{b^{\ell+1}c}$ or of the form $e(g,g)^z$, for random $z \in \mathbb{Z}_p^*$, depending on some random (and hidden) bit d that $\mathcal{B}$ is supposed to guess. First, notice that in the proposed scheme the ciphertext C is independent of specific identities, thus $\mathcal{B}$ can produce it without having to commit to any ID_0. $\mathcal{B}$ chooses $\overline{ID}$ at random as its guess for the challenge identity. Then it sets $g_1 = B_1^a$, for random $a \in \mathbb{Z}_p^*$, chooses at random $\alpha_i, \beta_i \xleftarrow{\$} \mathbb{Z}_p^*$, for $i = 1, \ldots, \ell$, and computes for $i = 1, \ldots, \ell$

$$g_{0i} = \begin{cases} B_1^{\beta_i} & if \;\; \overline{ID}_i = 0 \\ g^{\beta_i} & if \;\; \overline{ID}_i = 1 \end{cases} \qquad g_{1i} = \begin{cases} g^{\alpha_i} & if \;\; \overline{ID}_i = 0 \\ B_1^{\alpha_i} & if \;\; \overline{ID}_i = 1 \end{cases}$$

Notice that the public parameters $mpk = \left(g, g_1, \{g_{ij}\}_{i=0,1;j=1..\ell}\right)$ are distributed exactly as those produced by the setup algorithm. The master secret key is implicitly set to $msk = (ab, \{\alpha_i b^{\overline{ID}_i}, \beta_i b^{1-\overline{ID}_i}\}_{i=1,..,\ell})$. Next, $\mathcal{B}$ computes C^* as follows $C^* \leftarrow C = g^c$. Thus, C^* is also correctly distributed. Now $\mathcal{B}$ runs $\mathcal{A}$ on input (mpk, C^*, ID_0), for some randomly chosen identity ID_0. Notice that, from the received inputs, $\mathcal{A}$ gets no information at all about the $\overline{ID}$ chosen by $\mathcal{B}$, thus such a choice will be identical to the challenge identity with probability $1/2^\ell$.

Now we show how $\mathcal{B}$ can answer key derivation queries for identities $ID \neq \overline{ID}$. Since $ID \neq \overline{ID}$ there exists (at least) an index j such that $ID_j \neq \overline{ID}_j$. For such index we have that either $g_{0j} = g^{\beta_j}$ (if $ID_j = 0$) or $g_{1j} = g^{\alpha_j}$ (otherwise). This means that the h_ℓ corresponding to identity ID will contain the (unknown) b with exponent $\ell - 1$, at most. Let $n < \ell$ denote the number of positions i such that $ID_i = \overline{ID}_i$. $\mathcal{B}$ computes the h_i as follows.

$$h_1 = \begin{cases} g^{\alpha_1^{ID_1}\beta_1^{1-ID_1}} & if \;\; ID_1 \neq \overline{ID}_1 \\ B_1^{\alpha_1^{ID_1}\beta_1^{1-ID_1}} & if \;\; ID_1 = \overline{ID}_1 \end{cases}$$

$$h_2 = \begin{cases} h_1^{\alpha_2^{ID_2}\beta_2^{1-ID_2}} & if \;\; ID_2 \neq \overline{ID}_2 \\ B_1^{\alpha_2^{ID_2}\beta_2^{1-ID_2}} & if \;\; ID_2 = \overline{ID}_2 \wedge ID_1 \neq \overline{ID}_1 \\ B_2^{\alpha_2^{ID_2}\beta_2^{1-ID_2}} & if \;\; ID_2 = \overline{ID}_2 \wedge ID_1 = \overline{ID}_1 \end{cases} \quad \cdots$$

Finally, letting $\omega_{\mathcal{ID}} = \prod_{i=1}^{\ell} \alpha_i^{ID_i}\beta_i^{1-ID_i}$, h_ℓ is computed as $B_n^{\omega_{ID}}$.

The sk_{ID} is set to $B_{n+1}^{a\omega_{ID}}$. Recall that, since $n < \ell$, $\mathcal{B}$ can do this operation using the values received by the challenger. It is easy to check that both the $aux_{ID} = (h_1, \ldots, h_\ell)$ and sk_{ID} are distributed as in the real key derivation algorithm.

Once $\mathcal{A}$ is done with its first phase of key derivation queries it outputs a challenge identity ID^*. If $ID^* \neq \overline{ID}$, $\mathcal{B}$ outputs a random bit and aborts. Otherwise it constructs $K_{\overline{ID}}$ as $T^{a\omega_{\overline{ID}}}$, where $\omega_{\overline{ID}} = \prod_{i=1}^{\ell} \alpha_i^{\overline{ID}_i}\beta_i^{1-\overline{ID}_i}$ and $aux_{\overline{ID}}$ is computed as before. This time however h_ℓ is set to $B_\ell^{\omega_{\overline{ID}}}$, thus $\mathcal{B}$ will not be able to explicitly compute $sk_{\overline{ID}}$. However this is not a problem as $\mathcal{B}$ is not required to do so. Finally $\mathcal{B}$ hands $(K_{\overline{ID}}, sk_{\overline{ID}})$ to $\mathcal{A}$. $\mathcal{A}$ replies back with a guess d' ($d' = 0$

means real, $d' = 1$ means random). $\mathcal{B}$ outputs d' as well. Additional key derivation queries are dealt with as in the first phase. This completes the description of the simulator.

Now notice that if $T = e(g,g)^{b^{\ell+1}c}$, $K_{\overline{ID}}$ is a valid key for the identity $\overline{ID}$. This is because, $K_{\overline{ID}} = e(g_1, h_{\overline{ID}})^c$, where $h_{\overline{ID}}$ is the h_ℓ corresponding to identity $\overline{ID}$. Thus, $h_{\overline{ID}} = g^{b^\ell \omega_{\overline{ID}}}$

$$K_{\overline{ID}} = e(g_1, h_{\overline{ID}})^c = e(g^{ab}, g^{b^\ell \omega_{\overline{ID}}})^c = T^{a\omega_{\overline{ID}}}$$

If on the other hand T is a random value so is $K_{\overline{ID}}$. Thus, by standard calculations one gets that, if $\mathcal{A}$ has advantage ϵ in breaking the pseudorandom decapsulation property of the scheme, $\mathcal{B}$ breaks the decisional ℓ-wBDHI* with advantage $\epsilon/2^\ell$. □

Remark 5. It is interesting to note that if one is interested only into a selective-VRF, then the above construction leads directly to a scheme with large input space. This does not hold for the Dodis-Yampolskiy VRF because in its security proof the simulator has to guess all the queries that the adversary is going to ask even in the weaker selective case.

5 Conclusions

In this paper we introduced a new methodology to construct verifiable random functions (VRF) from a class of identity based key encapsulation schemes that we call VRF suitable. We showed the applicability of our methods by providing two concrete realizations of the new primitive. The first one leads to a VRF that is very similar to the Dodis-Yampolskiy construction, while the second leads to a new construction. A natural question left open by this research is to find new (potentially better) instantiations of the primitive, possibly ones supporting exponentially large (in the security parameter) identity spaces and provably secure under non interactive assumptions. This would solve the long standing open problem of realizing a secure VRF with unbounded input size.

Acknowledgements

We thank Gregory Neven for collaborating with us at an early stage of this research. We also thank Eike Kiltz and Jonathan Katz for helpful discussions. This work was supported in part by the European Commission through the IST Program under Contract ICT-2007-216646 ECRYPT II and in part by the French National Research Agency through the PACE project.

References

1. Bentahar, K., Farshim, P., Malone-Lee, J., Smart, N.P.: Generic constructions of identity-based and certificateless KEMs. Journal of Cryptology 21(2), 178–199 (2008)

2. 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)
3. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 56–73. Springer, Heidelberg (2004)
4. 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)
5. Boneh, D., Franklin, M.: Identity-based encryption from the weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
6. Boneh, D., Gentry, C., Hamburg, M.: Space-efficient identity based encryption-without pairings. In: 48th FOCS, Providence, USA, pp. 647–657. IEEE Computer Society Press, Los Alamitos (2007)
7. 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)
8. Chase, M., Lysyanskaya, A.: Simulatable vRFs with applications to multi-theorem NIZK. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 303–322. Springer, Heidelberg (2007)
9. Cheon, J.H.: Security analysis of the strong diffie-hellman problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 1–11. Springer, Heidelberg (2006)
10. 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)
11. Dodis, Y.: Efficient construction of (Distributed) verifiable random functions. In: Desmedt, Y.G. (ed.) PKC 2003. LNCS, vol. 2567, pp. 1–17. Springer, Heidelberg (2002)
12. Dodis, Y., Puniya, P.: Verifiable random permutations. Cryptology ePrint Archive, Report 2006/078 (2006), `http://eprint.iacr.org/`
13. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 416–431. Springer, Heidelberg (2005)
14. Gentry, C.: Practical identity-based encryption without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 445–464. Springer, Heidelberg (2006)
15. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: 21st ACM STOC, Seattle, Washington, USA, May 15–17, pp. 25–32. ACM Press, New York (1989)
16. Goldwasser, S., Ostrovsky, R.: Invariant signatures and non-interactive zero-knowledge proofs are equivalent. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 228–245. Springer, Heidelberg (1993)
17. Jarecki, S., Shmatikov, V.: Handcuffing big brother: an abuse-resilient transaction escrow scheme. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 590–608. Springer, Heidelberg (2004)
18. Liskov, M.: Updatable zero-knowledge databases. In: Roy, B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 174–198. Springer, Heidelberg (2005)
19. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM Journal on Computing 17(2) (1988)

20. Lysyanskaya, A.: Unique signatures and verifiable random functions from the DH-DDH separation. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 597–612. Springer, Heidelberg (2002)
21. Micali, S., Rabin, M.O., Vadhan, S.P.: Verifiable random functions. In: 40th FOCS, October 17–19, pp. 120–130. IEEE Computer Society Press, New York (1999)
22. Micali, S., Reyzin, L.: Soundness in the public-key model. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 542–565. Springer, Heidelberg (2001)
23. Micali, S., Rivest, R.L.: Micropayments revisited. In: Preneel, B. (ed.) CT-RSA 2002. LNCS, vol. 2271, pp. 149–163. Springer, Heidelberg (2002)
24. Naor, M., Reingold, O.: Number-theoretic constructions of efficient pseudo-random functions. In: 38th FOCS, Miami Beach, Florida, October 19–22, pp. 458–467. IEEE Computer Society Press, Los Alamitos (1997)
25. Sakai, R., Kasahara, M.: Id based cryptosystems with pairing on elliptic curve. In: 2003 Symposium on Cryptography and Information Security – SCIS 2003, Hamamatsu, Japan (2003), `http://eprint.iacr.org/2003/054`
26. 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)
27. Waters, B.: Efficient identity-based encryption without random oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (2005)

A The VRF by Dodis and Yampolskiy

In this section we describe the VRF by Dodis and Yampolskiy [13].

$\texttt{Gen}(1^k)$ Runs $\mathcal{G}(1^k)$ to obtain the description of the groups G, G_T and of a bilinear map $e : G \times G \to G_T$. The description of G contains a generator $g \in G$. Then the algorithm picks a random $s \xleftarrow{\$} \mathbb{Z}_p$ and sets $h = g^s$, $vpk = (g, h)$, $vsk = s$.

$\texttt{Func}_{vsk}(x)$Let $\texttt{Func}_{vsk}(x)$=$(F_{vsk}(x), \pi_{vsk}(x))$. One sets $\texttt{Func}_{vsk}(x)$ =$e(g, g)^{1/(s+x)}$ as the VRF output and $\pi_{vsk}(x) = g^{1/(s+x)}$ as the proof of correctness.

$\texttt{V}(vpk, x, y, \pi_x)$ To verify if y was computed correctly, one checks that $e(g^x \cdot h, \pi_x) = e(g, g)$ and $y = e(g, \pi_x)$.