

Replacing a Random Oracle: Full Domain Hash from Indistinguishability Obfuscation

Susan Hohenberger^{1,*}, Amit Sahai^{2,**}, and Brent Waters^{3,***}

¹ Johns Hopkins University, USA
susan@cs.jhu.edu

² UCLA, USA

sahai@cs.ucla.edu

³ University of Texas at Austin, USA
bwaters@cs.utexas.edu

Abstract. Our main result gives a way to instantiate the random oracle with a concrete hash function in “full domain hash” applications. The term full domain hash was first proposed by Bellare and Rogaway [BR93, BR96] and referred to a signature scheme from any trapdoor permutation that was part of their seminal work introducing the random oracle heuristic. Over time the term full domain hash has (informally) encompassed a broader range of notable cryptographic schemes including the Boneh-Franklin [BF01] IBE scheme and Boneh-Lynn-Shacham (BLS) [BLS01] signatures. All of the above described schemes required a hash function that had to be modeled as a random oracle to prove security. Our work utilizes recent advances in indistinguishability obfuscation to construct specific hash functions for use in these schemes. We then prove security of the *original* cryptosystems when instantiated with our specific hash function.

Of particular interest, our work evades the impossibility results of Dodis, Oliveira, and Pietrzak [DOP05], who showed that there can be no black-box construction of hash functions that allow Full-Domain Hash

* Susan Hohenberger is supported in part by NSF CNS-1154035 and CNS-1228443; DARPA and the Air Force Research Laboratory under contract FA8750-11-2-0211, DARPA N11AP20006, the Office of Naval Research under contract N00014-11-1-0470, and a Microsoft Faculty Fellowship. Applying to all authors, the views expressed are those of the authors and do not reflect the official policy or position of DARPA, the NSF, or the U.S. Government.

** Amit Sahai is supported in part from a DARPA/ONR PROCEED award, NSF grants 1228984, 1136174, 1118096, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0389.

*** Brent Waters is supported in part by NSF CNS-0915361 and CNS-0952692, CNS-1228599 DARPA through the U.S. Office of Naval Research under Contract N00014-11-1-0382, DARPA N11AP20006, Google Faculty Research award, the Alfred P. Sloan Fellowship, Microsoft Faculty Fellowship, and Packard Foundation Fellowship.

Signatures to be based on trapdoor permutations, and its extension by Dodis, Haitner, and Tentes [DHT12] to the RSA Full-Domain Hash Signatures. This indicates our techniques applying indistinguishability obfuscation may be useful for circumventing other black-box impossibility proofs.

1 Introduction

Since Bellare and Rogaway [BR93] introduced the Random Oracle Model, a major effort in cryptography has been to understand when and if random oracles can be instantiated with families of actual hash functions while maintaining security. Over the years, we have seen real progress in this effort: Firstly we have seen the discovery of alternative schemes that do not require random oracles but achieve the same security properties as earlier schemes that do require random oracles. For example, Cramer and Shoup [CS98] achieved efficient chosen ciphertext security from DDH hard groups. As another example Canetti, Halevi, and Katz [CHK07] achieved secure IBE without random oracles, following the seminal work of [BF01] giving IBE in the Random Oracle Model. More recently, we have seen the discovery of schemes that not only work in the standard model without random oracles, but work in a manner very similar to the original schemes that used random oracles (e.g. [HSW13, FHPS13] following schemes in the random oracle model [BF01, BLS01]). However, all of these schemes proven secure without random oracles required *changing the underlying cryptographic scheme* in addition to instantiating the random oracle with a concrete hash function. Thus, despite these advances, the following basic question has remained open:

Can we instantiate the random oracle with an actual family of hash functions for existing cryptographic schemes in the random oracle model, such as Full Domain Hash signatures?

In other words, can we achieve security *without changing the underlying cryptographic scheme* at all, but only by replacing the random oracle with a specific family of hash functions? In this work, we give the first positive answer to this question. We do this by leveraging the notion of indistinguishability obfuscation [BGI⁺01, BGI⁺12] that was recently achieved in the work of [GGH⁺13].

Our result is particularly interesting in light of negative results on the Random Oracle Model [CGH98, GK03, BBP04] which have called into question the secure applicability of the Random Oracle Model. Our work is the first to show natural examples of schemes that were originally invented with the Random Oracle Model in mind, that nevertheless remain secure when the random oracle is specifically instantiated.

In particular, our work evades the impossibility result of Dodis, Oliveira, and Pietrzak [DOP05], who showed that there can be no black-box construction of hash functions that allow Full-Domain Hash Signatures to be based on trapdoor permutations. Because we make use of obfuscation, our constructions are inherently non-black-box, and thus are not ruled out by this type of black-box impossibility result. This indicates that our techniques applying indistinguishability

obfuscation may be useful in the future for circumventing other such black-box impossibility proofs.

Our Result. Our main result gives a way to instantiate the random oracle with a concrete hash function in “full domain hash” (FDH) signatures. The FDH signature scheme was first proposed¹ in the original Bellare-Rogaway [BR93] paper as a way to build a signature scheme from any trapdoor permutation using the introduced random oracle heuristic. This work was very influential and formed the foundation for part of the PKCS#1 standard [KS98]. While the terminology of “full-domain hash” originally applied to the trapdoor permutation signature scheme of Bellare and Rogaway, over time it has (informally) encompassed a broader range of notable cryptographic schemes including the Boneh-Franklin [BF01] IBE scheme, the Cocks’s IBE scheme [Coc01], and Boneh-Lynn-Shacham (BLS) [BLS01] signatures. Although these schemes exist in different algebraic domains and have different aims, they share common construction and proof structures that uses random oracle programming in very similar ways.

Our work develops a methodology for replacing the programming of a random oracle in these construction using indistinguishable obfuscation in a novel manner. We begin by describing a scheme that replaces the RO hash function in the original Bellare-Rogaway trapdoor permutation (TDP) signature scheme. Our newly instantiated scheme is then proven to be selectively secure.

Let’s begin by informally recalling the Bellare-Rogaway TDP-based FDH scheme. The signature setup algorithm generates a trapdoor permutation pair of functions $g_{\text{PK}}, g_{\text{SK}}^{-1}$. It chooses a hash function $H(\cdot)$ that maps from the message space to the domain (and co-domain) of the permutation. The permutation g_{PK} and hash function are published as the verification key and the inverse g_{SK}^{-1} is kept secret. To sign a message m , the signer computes $g_{\text{SK}}^{-1}(H(m))$. To verify a signature σ on message m , the verifier simply checks whether $g_{\text{PK}}(\sigma) \stackrel{?}{=} H(m)$.

The proof of the Bellare-Rogaway FDH system uses the random oracle heuristic to model $H(\cdot)$ as a programmable random oracle. Suppose a poly-time attacker makes at most Q_H oracle queries. One can create a reduction algorithm to the security of the trapdoor permutation as follows. For all but one of the (unique) queries of a message m to the oracle, the reduction algorithm chooses a random value t from the domain and outputs $g_{\text{PK}}(t)$ as the result of the query. For any of these messages, the reduction algorithm can easily generate a signature by outputting t . However, at one query point m^* it programs the output of the random oracle to be $z^* = g_{\text{PK}}(t^*)$ where z^* was given from the trapdoor permutation challenger. If the attacker forges at this message, then the forgery will be t^* which is immediately the solution for the trapdoor permutation inversion.

Our first result is creating a *replacement* hash function for the oracle $H(\cdot)$ and developing a security proof without relying on the random oracle heuristic. To keep with our original goals, our only modifications will be to $H(\cdot)$ and we

¹ The terminology “full-domain hash” was actually introduced by Bellare-Rogaway in 1996 [BR96]. They applied this label to the noted signature scheme of their earlier work.

will use the signature system construction as is, with no changes to the underlying trapdoor permutation family. The two main tools we use to build $H(\cdot)$ are an indistinguishability obfuscator [BGI⁺01, GGH⁺13] and a recently introduced primitive called constrained PRFs [BW13, BGI14, KPTZ13]. In short, a constrained PRF key is a secret key K that allows the evaluator to evaluate the a PRF at a limited set of points, while the rest will appear pseudorandom to him. For our results, we only need a simple form of constrained PRFs called “punctured PRFs” [SW13]. In this setting a private key will be associated with a polynomial set S , where a key $K(S)$ can evaluate the PRF $F(K, x)$ at all x except when $x \in S$. For our proofs we only ever need S to be a singleton set.

We now overview the hash function construction and how we prove it to be selectively secure. (One could use the usual complexity leveraging arguments to claim adaptive security, but we will address adaptive security in a direct way shortly.) To create the hash function the reduction algorithm first chooses a puncturable PRF key K (note this “master key” can evaluate the PRF at all points). Next, the hash function itself will be an obfuscation of the program which on input m computes $g_{PK}(F(K, m))$. That is the program simply computes the PRF at point m and then applies the trapdoor permutation. We call this program Full Domain Hash. To prove security we will apply the “punctured programs” method of Sahai and Waters [SW13], where we surgically remove a key element of a program, but in a way that does not alter input/output functionality.

Our security proof is formed from a sequence of hybrids. In the first hybrid, we replace the obfuscation of the program Full Domain Hash with an obfuscation of an equivalent program called Full Domain Hash*. This program operates the same as the original except on input m^* , where m^* is the message the attacker selectively chose to attack (before seeing the verification key). At this point instead of computing $F(K, m^*)$ the program is simply hardwired to output a constant z^* to output where z^* is set to be $F(K, m^*)$. Since $z^* = F(K, m^*)$, the input/output behavior is identical. In addition, the program is not given the full PRF key K , but instead is given a punctured PRF key $K(\{m^*\})$. By the security of indistinguishable obfuscation the advantage of any poly-time attacker must be negligibly close between these hybrids. In the next hybrid experiment we replace z^* with a random value chosen from the domain/range of the permutation. The advantage between of this hybrid must also be close due to the constrained PRF security. Now we are finally in a position where we can reduce to the security of the trapdoor permutation. The reduction algorithm receives a TDP challenge z^* and hardcodes that in as the output of $H(m^*)$. It can use a signature on this to invert the challenge. At all other points it knows the punctured PRF key and can therefore compute valid signatures without knowing the inverse of the trapdoor permutation.

Our reduction actually shares some of the spirit of the original random oracle reduction, where a challenge is programmed in at one point and signatures are made by knowing the pre images at all others. A key aspect is that the obfuscation hides the fact that at a certain hybrid m^* is treated differently. If an attacker were able to see inside the obfuscation it could actually see the preimages and

break the scheme. Another interesting aspect is that our proof does not leverage the fact that the function $g_{PK}(\cdot)$ is a permutation. It would go through equally well if we only assumed that it was an injective trapdoor function.

In our construction and proof, there is a hash function created as part of each public key. Taking things further, we might want to have *one* hash function built as a common reference string that could serve as part of many public keys. Creating similar results in this setting will require further work.

There is also a connection between the proof techniques we use and the proof technique for the selectively secure signature scheme of Sahai and Waters [SW13]. In the Sahai-Waters work, the verification key is an obfuscated program that evaluates the punctured PRF on a message and then outputs the evaluation of a one-way function on that. A signer signs by evaluating the punctured PRF on the message. In comparison, in our case, the hash function is the obfuscated program that evaluates the punctured PRF on a message and then outputs the evaluation of a trapdoor permutation on that. Here in contrast, the signing is done by applying the inverse permutation and the signer isn't necessarily "aware" of the punctured PRF.

Overcoming the Black-Box Impossibility. We now see more precisely why our work evades the impossibility result of Dodis, Oliveira, and Pietrzak [DOP05] and Dodis, Haitner, and Tentes [DHT12]. Our hash function is obfuscation of code that runs the underlying permutation. The obfuscation will intuitively hide the evaluation of this code. In particular, no attacker can tell if the trapdoor permutation was actually computed on an input or whether it was a special point where the output was hardcoded in. In the DOP negative result, they build an attack oracle that specifically leverages the black box access to the TDP to watch whenever it is called. It is interesting to see this very strong correlation between the negative result and how non-black box access to a primitive and indistinguishability obfuscation can combine to circumvent it.

Getting Adaptive Security. For our next result we show how to get adaptive (or standard) signature security without complexity leveraging for the case where the trapdoor permutation is the RSA function. The use of RSA as a trapdoor permutation candidate was suggested in Bellare-Rogaway'93 [BR93] and explicitly given in Bellare-Rogaway'96 [BR96]. The public parameters in their scheme are an RSA modulus $N = pq$ for hidden primes p, q and an RSA exponent e chosen such that $\gcd(\phi(N), e) = 1$. The secret key is the integer d where $d \cdot e = 1 \pmod{\phi(N)}$. A signature on message m is of the form $H(m)^d \pmod{N}$ and one verifies a signature σ by checking if $H(m) \stackrel{?}{=} \sigma^e \pmod{N}$.

We develop a different set of techniques that can leverage the particular structure of the RSA function. The first new ingredient is use of admissible hash functions first introduced in the context of Identity-Based Encryption by Boneh-Boyen [BB04a]. We use a simplification due to Freire et. al. [FHPS13]. At a high level the system is a pair of a hash function $h : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{n(\lambda)}$ that hashes from the message space to n bit strings and an efficient randomized algorithm AdmSample . The sampling algorithm takes in the security parameter as well as

second parameter Q which intuitively corresponds to the number of signature queries an attacker makes. It outputs a string $u \in \{0, 1, \perp\}^n$. Informally, we say that the system is admissible if the following conditions hold. Consider any sequence of Q values x_1, \dots, x_Q and $x^* \neq x_i$. The event we consider is where the string $h(x_i)$ has a bit in common with u in at least one position, but $h(x^*)$ is different from u at all positions. (Note, if $u_j = \perp$ then it is different at position j from all bit strings.) If this event occurs with non-negligible probability, we say it is an admissible system. Intuitively, when used in a proof of a signature scheme, the admissible hash function is utilized to partition the message space into messages that can be signed in the query phase and those that can be used in the challenge phase. A sampled string u corresponds to a particular partition. When running a reduction, one hopes that the actual signature oracle queries and forgery message align with a partition, and the reduction aborts otherwise.

To build the hash function candidate, the setup first chooses a random $v \in \mathbb{Z}_N^*$ as well as exponents $a_{i,b}$ chosen randomly in $[0, \phi(N)]$, for all $i \in [1, n]$, $b \in \{0, 1\}$. Next, it builds the hash function as an obfuscation of the program RSA Hash. The program will first compute $m' = h(m)$. Then, it computes and outputs $v^{\prod_{i \in [n]} a_{i,m'}}$.

Our proof proceeds in a few hybrid steps. In the first hybrid experiment the challenger creates a partition internally by calling $\text{AdmSample}(1^\lambda, Q) \rightarrow u$ for an attacker that makes at most $Q = Q(\lambda)$ queries. The game aborts and declares the attacker unsuccessful if any of the query messages or forgery message violates the partition. The property of admissible hashes states any attacker with non-negligible advantage in the real game will also have non-negligible advantage here. In the next hybrid, we change the way we sample the exponents $a_{i,b}$. One first chooses random $y_{i,b} \in [1, N]$. Then for when $u_i = b$ we set $c_{i,b} = e \cdot y_{i,b}$. If $u_i \neq b$ we set $c_{i,b} = e \cdot y_{i,b} + 1$. Note in the first case $c_{i,b}$ is a multiple of e and in the second case $e \nmid c_{i,b}$. The values $a_{i,b} = c_{i,b} \bmod \phi(N)$. We show that this way of choosing a values is statistically close to the previous uniform way, because $\gcd(\phi(N), e) = 1$.

Next, we use an alternative program where we directly use the $c_{i,b}$ values in place of the $a_{i,b}$ values. Since the group \mathbb{Z}_N^* is of order $\phi(N)$ we have that $v^{\prod_{i \in [n]} a_{i,m'}} = v^{\prod_{i \in [n]} c_{i,m'}}$ for all m' . Therefore the input/output behavior is the same between the two programs and we can argue the advantage in the hybrids for poly-time attackers must be close by indistinguishability obfuscation. This is the critical hybrid experiment in that it most radically departs from previous such proofs, by leveraging indistinguishability obfuscation. Observe that this hybrid experiment eliminates the need for the reduction to know $\phi(N)$, which is crucial to the reduction, since it uses $c_{i,b}$ values instead of $a_{i,b}$ values. However, if the values $c_{i,b}$ were completely visible to an attacker, they would be trivially distinguishable from the “true” uniform $a_{i,b}$ values. However, indistinguishability obfuscation guarantees that these values are hidden from the attacker, and that indeed the attacker cannot distinguish this hybrid from the previous one.

Finally, we show that any attacker that is successful in the last hybrid can be used to break the RSA assumption. For any signature query message m that

respects the partition, the reduction will view $H(m)$ as v raised to some integer that is a multiple of e and taking the e -th root is then easy. Any forgery on m^* that respects the partition, the reduction will view $H(m^*)$ as v^z for some z where $\gcd(e, z) = 1$ and from this can derive $v^{1/e}$.

BLS Signatures and More. We extend our techniques to replacing the random oracle in the BLS [BLS01] signature scheme. In Section 5 we give a candidate that has a selective proof of security based on the computational Diffie-Hellman problem (along with indistinguishability obfuscation). In the full version [HSW14], we give an adaptive proof of security based on an assumption equivalent to the n -Diffie-Hellman inversion assumption. The high level structures of these are similar to the respective selective and adaptive construction and proof methods above. The lower level mechanisms are adapted to the context of bilinear groups. In Section 7, we sketch how the BLS ideas extend to the Boneh-Franklin IBE scheme.

1.1 Other Related Work

Early work on replacing random oracles for the problem of obfuscating point functions under entropy conditions began with the work of Canetti [Can97].

Recently, the work of Bellare, Hoang and Keelveedhi [BHK13] looked at a complementary question of identifying a definitional abstraction to replace the random oracle heuristic in several random oracle-based constructions. The abstraction is a notion of security called UCE (Universal Computational Extractor). The authors emphasize that a random oracle is *known* not to exist and “behaves like a random oracle” is not a rigorously defined property, whereas UCE is a well defined property of a hash function. They then show how several previous constructions proven secure in the random schemes can be proven secure if we assume the hash functions are UCE secure. One can then conjecture that standard cryptographic hash functions like SHA-256 may satisfy the UCE security notion. In contrast, our work is focused on providing *new* candidate constructions for hash functions, that allow for a security proof to work with the original constructions in the random oracle model. Interestingly, the work of [BHK13] does not encompass the case of Full Domain Hash signatures, arguably one of the most natural and well-studied constructions in the Random Oracle Model, that we address here.

Dodis, Haitner, and Tentes [DHT12] show how to give an FDH signature that is secure for at most q queries when the hash function grows with q .

2 Preliminaries

We define indistinguishability obfuscation, and variants of pseudo-random functions (PRFs) that we will make use of. All the variants of PRFs that we consider will be constructed from one-way functions.

2.1 Indistinguishability Obfuscation

The definition below is from [GGH⁺13]; there it is called a “family-indistinguishable obfuscator”. They show that this notion follows immediately from their standard definition of indistinguishability obfuscator using a non-uniform argument.

Definition 1 (Indistinguishability Obfuscator ($i\mathcal{O}$)). *A uniform PPT machine $i\mathcal{O}$ is called an indistinguishability obfuscator for a circuit class $\{\mathcal{C}_\lambda\}$ if the following conditions are satisfied:*

- For all security parameters $\lambda \in \mathbb{N}$, for all $C \in \mathcal{C}_\lambda$, for all inputs x , we have that

$$\Pr[C'(x) = C(x) : C' \leftarrow i\mathcal{O}(\lambda, C)] = 1$$

- For any (not necessarily uniform) PPT adversaries $Samp, D$, there exists a negligible function α such that the following holds: if $\Pr[\forall x, C_0(x) = C_1(x) : (C_0, C_1, \tau) \leftarrow Samp(1^\lambda)] > 1 - \alpha(\lambda)$, then we have:

$$\left| \Pr [D(\tau, i\mathcal{O}(\lambda, C_0)) = 1 : (C_0, C_1, \tau) \leftarrow Samp(1^\lambda)] - \Pr [D(\tau, i\mathcal{O}(\lambda, C_1)) = 1 : (C_0, C_1, \tau) \leftarrow Samp(1^\lambda)] \right| \leq \alpha(\lambda)$$

In this paper, we will make use of such indistinguishability obfuscators for all polynomial-size circuits:

Definition 2 (Indistinguishability Obfuscator for $P/poly$). *A uniform PPT machine $i\mathcal{O}$ is called an indistinguishability obfuscator for $P/poly$ if the following holds: Let \mathcal{C}_λ be the class of circuits of size at most λ . Then $i\mathcal{O}$ is an indistinguishability obfuscator for the class $\{\mathcal{C}_\lambda\}$.*

Such indistinguishability obfuscators for all polynomial-size circuits were constructed under novel algebraic hardness assumptions in [GGH⁺13].

2.2 Constrained PRFs

We first consider some simple types of constrained PRFs [BW13, BGI14, KPTZ13], where a PRF is only defined on a subset of the usual input space. We focus on *puncturable* PRFs, which are PRFs that can be defined on all bit strings of a certain length, except for any polynomial-size set of inputs:

Definition 3. *A puncturable family of PRFs F mapping is given by a triple of Turing Machines Key_F , $Puncture_F$, and $Eval_F$, and a pair of computable functions $n(\cdot)$ and $m(\cdot)$, satisfying the following conditions:*

- [**Functionality preserved under puncturing**] *For every PPT adversary A such that $A(1^\lambda)$ outputs a polynomial-size set $S \subseteq \{0, 1\}^{n(\lambda)}$, then for all $x \in \{0, 1\}^{n(\lambda)}$ where $x \notin S$, we have that:*

$$\Pr [Eval_F(K, x) = Eval_F(K_S, x) : K \leftarrow Key_F(1^\lambda), K_S = Puncture_F(K, S)] = 1$$

- [**Pseudorandom at punctured points**] For every PPT adversary (A_1, A_2) such that $A_1(1^\lambda)$ outputs a polynomial-size set $S \subseteq \{0, 1\}^{n(\lambda)}$ and state τ , consider an experiment where $K \leftarrow \text{Key}_F(1^\lambda)$ and $K_S = \text{Puncture}_F(K, S)$. Then we have

$$\left| \Pr [A_2(\tau, K_S, S, \text{Eval}_F(K, S)) = 1] - \Pr [A_2(\tau, K_S, S, U_{m(\lambda)-|S|}) = 1] \right| = \text{negl}(\lambda)$$

where $\text{Eval}_F(K, S)$ denotes the concatenation of $\text{Eval}_F(K, x_1), \dots, \text{Eval}_F(K, x_k)$ where $S = \{x_1, \dots, x_k\}$ is the enumeration of the elements of S in lexicographic order, $\text{negl}(\cdot)$ is a negligible function, and U_ℓ denotes the uniform distribution over ℓ bits.

For ease of notation, we write $F(K, x)$ to represent $\text{Eval}_F(K, x)$. We also represent the punctured key $\text{Puncture}_F(K, S)$ by $K(S)$.

The GGM tree-based construction of PRFs [GGM84] from one-way functions are easily seen to yield puncturable PRFs, as observed by [BW13, BGI14, KPTZ13]. Thus we have:

Theorem 1. [GGM84, BW13, BGI14, KPTZ13] *If one-way functions exist, then for all efficiently computable functions $n(\lambda)$ and $m(\lambda)$, there exists a puncturable PRF family that maps $n(\lambda)$ bits to $m(\lambda)$ bits.*

2.3 RSA Assumption and Shamir’s Lemma

We begin by recalling (one of the) standard versions of the RSA assumption [RSA78].

Assumption 1 (RSA). *Let λ be the security parameter. Let positive integer N be the product of two λ -bit, distinct odd primes p, q . Let e be a randomly chosen positive integer less than and relatively prime to $\phi(N) = (p-1)(q-1)$. Given (N, e) and a random $y \in \mathbb{Z}_N^*$, it is hard to compute x such that $x^e \equiv y \pmod{N}$.*

We also make use of the following lemma due to Shamir.

Lemma 1 (Shamir [Sha83]). *Given $x, y \in \mathbb{Z}_N$ together with $a, b \in \mathbb{Z}$ such that $x^a = y^b \pmod{N}$ and $\text{gcd}(a, b) = 1$, there is an efficient algorithm for computing $z \in \mathbb{Z}_N$ such that $z^a = y \pmod{N}$.*

2.4 Bilinear Groups and the CDH Assumption

Let \mathbb{G} and \mathbb{G}_T be groups of prime order p . A *bilinear map* is an efficient mapping $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ which is both: (*bilinear*) for all $g \in \mathbb{G}$ and $a, b \leftarrow \mathbb{Z}_p$, $e(g^a, g^b) = e(g, g)^{ab}$; and (*non-degenerate*) if g generates \mathbb{G} , then $e(g, g) \neq 1$.

Assumption 2 (Computational Diffie-Hellman). *Let g generate a group \mathbb{G} of prime order $p \in \Theta(2^\lambda)$. For all p.p.t. adversaries \mathcal{A} , the following probability is negligible in λ :*

$$\Pr[a, b \leftarrow \mathbb{Z}_p; z \leftarrow \mathcal{A}(g, g^a, g^b) : z = g^{ab}].$$

2.5 The n -Diffie-Hellman Inversion Assumption

Our full version [HSW14] contains a construction of adaptively secure BLS signatures that makes use of the n -Diffie-Hellman Inversion assumption [BB04b]. This is a parameterized family of assumptions, where the number of group elements involved increases with n . (For our application, n will be dependent only on the security parameter.)

Assumption 3 (n -Diffie-Hellman Inversion). *Let h generate a group \mathbb{G} of prime order $p \in \Theta(2^\lambda)$. For all $p.p.t.$ adversaries \mathcal{A} , the following probability is negligible in λ :*

$$\Pr[b \leftarrow \mathbb{Z}_p; z \leftarrow \mathcal{A}(h, h^b, h^{b^2}, \dots, h^{b^n}) : z = g^{1/b}].$$

3 Full-Domain Hash Signatures (Selectively Secure)

In this section, we revisit the Bellare–Rogaway Full-Domain Hash (FDH) signature scheme [BR93, BR96], and show how to make it selectively secure in the standard model by instantiating the random oracle in a specific way. We stress that we do not modify the Bellare–Rogaway FDH signature scheme in any way; the only new aspect of our construction is our instantiation of the random oracle with a specific function whose description becomes part of the public key.

Recall that the Bellare–Rogaway FDH signature scheme required a trapdoor permutation family. Our method, in fact, not only applies to trapdoor permutation families, but indeed to any *injective* trapdoor function family. We prove the selective security of the FDH signature scheme based on the security of the indistinguishability obfuscator, the security of a puncturable PRF family, and the security of an injective trapdoor function family.

For simplicity of exposition, we assume that there is a polynomial $\ell(\lambda)$ which denotes the length of messages to be signed; we denote this message space by $\mathcal{M} = \{0, 1\}^{\ell(\lambda)}$. More generally, a collision-resistant hash function may be used to hash messages to this size.

- **Setup(1^λ)** : The setup algorithm first runs $\text{TDFSetup}(1^\lambda)$ and that produces a public index PK along with a trapdoor SK , yielding the map $g_{\text{PK}} : \{0, 1\}^n \rightarrow \{0, 1\}^w$ together with its inverse. Next, the setup algorithm chooses a puncturable PRF key K for F where $F(K, \cdot) : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^n$. Then, it creates an obfuscation of the of the program Full Domain Hash Figure 1. The size of the program is padded to be the maximum of itself and the program Full Domain Hash* of Figure 2. We refer to the obfuscated program as the function $H : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^w$, which acts as the random oracle type hash function in the Bellare–Rogaway scheme.

The verification key VK consists of the trapdoor index PK as well as the hash function $H(\cdot)$. The secret key is the trapdoor SK as well as $H(\cdot)$.

- **Sign($\text{SK}, m \in \mathcal{M}$)** : The signature algorithm outputs $\sigma = g_{\text{SK}}^{-1}(H(m)) \in \{0, 1\}^n$.




Fig. 1. Full Domain Hash




Fig. 2. Full Domain Hash*

- **Verify**(VK, m, σ) The verification algorithm tests if $g_{\text{PK}}(\sigma) \stackrel{?}{=} H(m)$ and outputs accept if and only if this holds.

Theorem 2. *If our obfuscation scheme is indistinguishably secure, F is a secure punctured PRF, and the injective trapdoor function is secure, then the above signature scheme is selectively secure.*

We describe a proof as a sequence of hybrid experiments where the first hybrid corresponds to the original signature security game. We prove that a poly-time attacker's advantage must be negligibly close between each successive one. Then, we show that any poly-time attacker in the final experiment that succeeds in forging with non-negligible probability can be used to invert the injective trapdoor function.

- **Hyb₀** : In the first hybrid the following game is played:
 1. The attacker selectively gives the challenger the message m^* .
 2. The TDF index is chosen by the challenger running $\text{TDFSetup}(1^\lambda)$.
 3. K is chosen as a key for the puncturable PRF.
 4. The hash function $H(\cdot)$ is created as an obfuscation of the program Full Domain Hash.
 5. The attacker queries the sign oracle a polynomial number of times on messages $m \neq m^*$. It receives back $g_{\text{SK}}^{-1}(H(m)) = F(K, m)$. (Note the equality holds since the function g_{PK} is injective.)
 6. The attacker sends a forgery σ^* and wins if $\text{Verify}(\text{VK}, m^*, \sigma^*) = 1$.
- **Hyb₁** : Is the same as **Hyb₀** except we let $z^* = g_{\text{PK}}(F(K, m^*))$ and let VK be the obfuscation of the program **Verify Signature*** of Figure 2.
- **Hyb₂** : Is the same as **Hyb₁** except $z^* = g_{\text{PK}}(t)$ for t chosen uniformly at random in $\{0, 1\}^n$.

The following three lemmas together yield our result in Theorem 2 that the full domain hash signature scheme in Section 3 is selectively secure.

Lemma 2. *If our obfuscation scheme is indistinguishability secure, then the advantage of a poly-time attacker in Hyb_0 is negligibly close to the advantage in Hyb_1 .*

Proof. We prove this lemma by giving a reduction to the indistinguishability security of the obfuscator. To do so, we must build the two algorithms $Samp$ and D .

$Samp(1^\lambda)$ behaves as follows: It invokes the adversary to obtain m^* and the adversary's state τ' . It runs $\text{TDFSetup}(1^\lambda)$ to obtain PK and SK. It then chooses K as a key for the puncturable PRF. It sets $z^* = g_{\text{PK}}(F(K, m^*))$. It sets $\tau = (m^*, \text{PK}, \text{SK}, K, \tau')$ and builds C_1 as the program for Full Domain Hash, and C_2 as the program for Full Domain Hash*.

Before describing D , we observe that by construction and the functionality preservation property of puncturable PRFs, the circuits C_1 and C_2 always behave identically on every input. Because of padding, both C_1 and C_2 have the same size. Thus, $Samp$ satisfies the conditions needed for invoking the indistinguishability property of the obfuscator.

Now, we can describe the algorithm D , which takes as input τ as given above, and either the obfuscation of C_1 , which is the program Full Domain Hash, or C_2 , which is the program Full Domain Hash*. D creates the verification key for the signature scheme by combining PK with the obfuscated program as the hash function description. It then invokes the adversary on this verification key, and the adversary then makes requests for signatures on messages $m \neq m^*$. For each such message, D constructs the signatures $g_{\text{SK}}^{-1}(H(m)) = F(K, m)$, through its knowledge of K within τ . Finally, the attacker sends a forgery σ^* and wins if $\text{Verify}(m^*, \sigma^*) = 1$. If the attacker wins, D outputs 1.

By construction, if D receives an obfuscation of C_1 , then the probability that D outputs 1 is exactly the probability of the adversary winning in hybrid Hyb_0 . On the other hand, if D receives an obfuscation of C_2 , then the probability that D outputs 1 is the probability of the adversary winning in hybrid Hyb_1 .

The lemma follows.

Lemma 3. *If our confined PRF is secure, then the advantage of a poly-time attacker in Hyb_1 is negligibly close to the advantage in Hyb_2 .*

Proof. We prove this lemma by giving a reduction to the pseudorandomness property at punctured points for punctured PRFs. To do so, we must build the algorithms A_1 and A_2 .

$A_1(1^\lambda)$ simply invokes the adversary to obtain the challenge message m^* and state τ' , and outputs the singleton set $S = \{m^*\}$ and $\tau = (1^\lambda, \tau')$.

A_2 obtains as input τ , the punctured key K_S , the singleton set $S = \{m^*\}$, and either a value $t^* = F(K, m^*)$ or a uniformly random value t^* . Then, A_2 invokes $\text{TDFSetup}(1^\lambda)$ to obtain PK and SK. Now given t^* , it can compute $z^* = g_{\text{PK}}(t^*)$. Note that this yields either the z^* value computed in hybrid Hyb_1

or in hybrid Hyb_2 . Since it knows K_S , now A_2 can obfuscate the program Full Domain Hash*, and then execute the adversary and answer its signature queries using the punctured key K_S . Finally, A_2 outputs 1 if the adversary succeeds.

By construction, the pseudorandomness property for punctured PRFs implies the lemma.

Lemma 4. *If our injective trapdoor function is hard to invert, then the advantage of a poly-time attacker in Hyb_2 is negligible.*

Proof. We prove this lemma by giving a reduction to the one-wayness of the injective trapdoor function. To do so, we build an inverting algorithm Inv .

Inv takes as input a public index PK for an injective trapdoor function, and a target $z^* = g_{\text{PK}}(t^*)$ for some (as yet unknown) random value t^* . The algorithm Inv then invokes the adversary to obtain m^* , and chooses a PRF key K and builds the punctured key $K(S)$ where $S = \{m^*\}$. It uses this key, together with PK and z^* , to obfuscate the program Full Domain Hash*. It can then execute the adversary, and use its knowledge of $K(S)$ to answer all adversary signing queries. The adversary then terminates with an attempted forgery σ^* on message m^* . By the definition of the program Full Domain Hash*, this forgery can only be valid if $g_{\text{PK}}(\sigma^*) = z^*$, and because g_{PK} is injective, this can only happen if $\sigma^* = t^*$. Thus if the adversary is successful, Inv can output σ^* as a valid pre-image of z^* .

We observe that by construction of Inv , the probability of success of Inv is exactly the probability that the attacker succeeds in hybrid Hyb_2 . The lemma follows.

4 Adaptively Secure RSA Full Domain Hash Signatures

We first overview what advantage indistinguishability obfuscation gives us in this situation: In several previous constructions of adaptively secure schemes in the plain model starting with the adaptively secure IBE scheme of [BB04a], a special hash function was chosen that allowed for a “partitioning” proof of security. In essence, for this to work, the hash function should have two “modes”:

- In the “normal” mode, the hash function’s parameters are typically just chosen at random, and it behaves like an ordinary hash function.
- In the “partitioning” mode, the hash function parameters are chosen according to a special distribution. This special distribution allows for the efficient computation of the inverse of the hash value for a large fraction of points, but it has the property that computing the inverse of the hash value at any other point is computationally hard.

It is crucial that the input/output functionality of the hash function should be identical in the two modes, and we will also use this property. However, in previous proofs (like [BB04a]), it was also critical that the hash function parameters in “partitioning” mode be information theoretically indistinguishable

from the parameters in “normal” mode, and thus the partition should be hidden from the adversary even when given the hash function parameters. This restriction significantly limited the applicability of this technique, as it could only be applied with algebraic structures that allowed for such “pseudorandom” hash parameters. Thanks to indistinguishability obfuscation, however, we can avoid this restriction by obfuscating the hash function description. Thus, even if the natural hash function parameters in “partitioning” mode clearly reveal the partition and thus are distinguishable from normal parameters, because the resulting hash function is *functionally* identical to a hash function in “normal” mode, the obfuscated hash function *must* hide the partition, and this allows the proof of adaptive security to go through.

In describing our signature scheme, For simplicity of exposition, we assume that there is a polynomial $\ell(\lambda)$ which denotes the length of messages to be signed; we denote this message space by $\mathcal{M} = \{0, 1\}^{\ell(\lambda)}$. More generally, a collision-resistant hash function may be used to hash messages to this size. Below, for any polynomial in λ , after the first mention of this polynomial, we will often suppress the dependence on λ for ease of notation. Thus, below often we will simply refer to the size of messages to be signed by ℓ .

Before describing our construction, we first recall a (simplified) description of the notion of *admissible* hash functions due to [BB04a]. Our definition is a slight variation of the simplified definition due to [FHPS13].

Definition 4. *Let ℓ, n and θ be efficiently computable univariate polynomials. We say that an efficiently computable function $h : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{n(\lambda)}$, and an efficient randomized algorithm AdmSample , is θ -admissible if the following condition holds:*

For any $u \in (\{0, 1\} \cup \{\perp\})^n$, define $P_u : \{0, 1\}^\ell \rightarrow \{0, 1\}$ as follows: $P_u(x) = 0$ iff $\forall i : h(x)_i \neq u_i$, and otherwise (if $\exists i : h(x)_i = u_i$) we have $P_u(x) = 1$.

Then we require that for any efficiently computable polynomial $Q(\lambda)$, for all $x_1, \dots, x_Q, z \in \{0, 1\}^\ell$, where $z \notin \{x_i\}$, we have that

$$\Pr [P_u(x_1) = P_u(x_2) = \dots = P_u(x_Q) = 1 \wedge P_u(z) = 0] \geq 1/\theta(Q)$$

where the probability is taken only over $u \leftarrow \text{AdmSample}(1^\lambda, Q)$.

Theorem 3 (Admissible Function Families [BB04a], see also [FHPS13] for a simple proof). *For any efficiently computable polynomials ℓ, n , there exists an efficiently computable polynomial θ such that there exist θ -admissible function families mapping ℓ bits to n bits.*

We leverage the structure of the RSA trapdoor permutation to prove adaptive security. The use of RSA as a candidate for a trapdoor permutation was first discussed in the original Bellare-Rogaway [BR93] paper, however, it was in [BR96] that Bellare and Rogaway gave an explicit full domain hash RSA construction. This construction formed the basis for part of the standard PKCS#1 [KS98].

- **Setup(1^λ)** : The setup algorithm first runs an RSA type setup. It chooses random primes p, q of λ bits each. We define $N = p \cdot q$ and $\phi(N) =$

$(p-1)(q-1)$. We let e be a random chosen integer between 1 and $\phi(N)$ such that $\gcd(\phi(N), e) = 1$. Next, it chooses integers $(a_{1,0}, a_{1,1}), \dots, (a_{n,0}, a_{n,1})$ each uniformly at random from the range $[1, \phi(N) - 1]$. In addition, it chooses a group element $v \in \mathbb{Z}_N^*$. It then creates an obfuscation of the of the program RSA Hash of Figure 3. The size of the program is padded to be the maximum of itself and the program RSA Hash* of Figure 4. We refer to the obfuscated program as the function $H(\cdot)$. This function $H(\cdot)$ will replace the random oracle in the RSA FDH scheme, but no other part of the scheme is modified. The verification key VK is the integers N, e and the hash function $H : \{0, 1\}^{\ell(\lambda)} \rightarrow \mathbb{Z}_N^*$. The secret key is the integer d where $e \cdot d \equiv 1 \pmod{\phi(N)}$.

- **Sign**(SK, $m \in \mathcal{M}$) : The signature algorithm outputs $\sigma = H(M)^d \pmod N$.
- **Verify**(VK, m, σ) The verification algorithm tests if $\sigma^e \equiv H(m) \pmod N$ and outputs accept if and only if this holds.

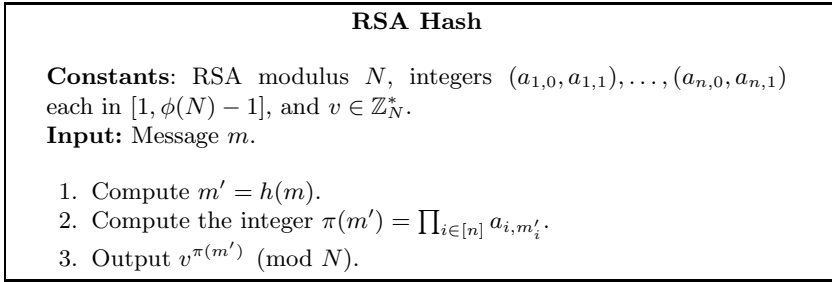


Fig. 3. RSA Hash

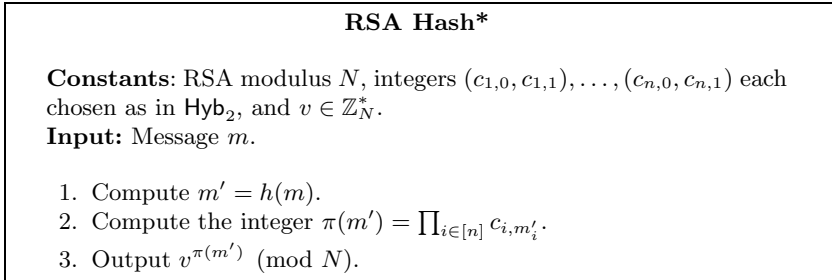


Fig. 4. RSA Hash*

Remark 1. For simplicity of exposition we describe computing the programs output by first computing a integer $\pi(m')$ as a product of n integers and then raising v to this mod N . In practice, it might be more efficient to incrementally raise an accumulated value to each a_{i, m'_i} .

Theorem 4. *If our obfuscation scheme is indistinguishably secure and the RSA assumption holds, the above signature scheme is existentially unforgeable against chosen message attacks.*

In the full version [HSW14], we describe a proof as a sequence of hybrid experiments where the first hybrid corresponds to the original signature security game. In the first hybrid step we do a “partitioning” of the message space. Consider a poly-time attacker that makes $Q = Q(\lambda)$ signature queries m_1, \dots, m_Q and attempts to forge on message $m^* \neq m_i$ for all i . Roughly, at the beginning of Hyb_1 the challenger will now (behind the scenes) partition the message space such that a large fraction of messages will fall into a “query” space and a much smaller, but still non-negligible fraction of messages will fall into the “challenge” space. Furthermore, in this new game the attacker is only considered to have won if he both forged a signature *and* all his signature queries m_1, \dots, m_n fall into the query space and m^* falls into the challenge space. We can show that if an attacker succeeds in the original security game (that does not have these additional restrictions on winning) with non-negligible advantage, then it will succeed in Hyb_1 with non-negligible advantage. Our system uses the Boneh-Boyen [BB04a] admissible hash function defined above, where if an attacker has advantage ϵ in Hyb_0 , he will have advantage $\epsilon/\theta(Q)$ in Hyb_1 . After the first proof step we prove that a poly-time attacker’s advantage must be negligibly close between each successive hybrid experiment. We finally show that any poly-time attacker in the final experiment that succeeds with non-negligible probability can be used to break the RSA assumption.

5 Selectively Secure BLS Signatures

We now give a concrete construction for the hash function modeled as a random oracle in the Boneh-Lynn-Shacham (BLS) signature scheme. BLS signatures fall into a broad interpretation (see e.g., [Boy08]) of the full domain hash paradigm of Bellare and Rogaway. Below we give the BLS signature scheme with a concrete hash function built from an indistinguishability obfuscator. We prove the signature scheme selectively secure based on the computational Diffie-Hellman problem in bilinear groups and a indistinguishability obfuscator.

On a technical level this selective proof of security follows a very similar structure to that of our selectively secure scheme from trapdoor functions from Section 3. The main difference is that here we deal with the mechanics of an algebraic bilinear group instead of a trapdoor function. We present the scheme for simplicity in terms of a symmetric bilinear group, however, moving to asymmetric groups is straightforward. As in Section 3, we assume that there is a polynomial $\ell(\lambda)$ which denotes the length of messages to be signed; we denote this message space by $\mathcal{M} = \{0, 1\}^{\ell(\lambda)}$. More generally, a collision-resistant hash function may be used to hash messages to this size.

- $\text{Setup}(1^\lambda)$: The setup algorithm first runs the group generator on input 1^λ to produce a description of groups \mathbb{G}, \mathbb{G}_T of prime order p along with generator $g \in \mathbb{G}$. These groups are related by a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. Next, it chooses a random exponent $a \in \mathbb{Z}_p$. Then, the setup algorithm chooses a puncturable PRF key K for F where $F(K, \cdot) : \{0, 1\}^{\ell(\lambda)} \rightarrow \mathbb{Z}_p$. Finally, it

creates an obfuscation of the program BLS Selective Hash of Figure 5. The size of the program is padded to be the maximum of itself and the program BLS Selective Hash* of Figure 6. We refer to the obfuscated program as the function $H : \{0, 1\}^\ell \rightarrow \mathbb{G}$, which acts as the random oracle type hash function in the BLS scheme.

The verification key VK consists of the group descriptions \mathbb{G}, \mathbb{G}_T , the order p , the generator g and $A = g^a$ as well as the hash function $H(\cdot)$. The secret key is $a \in \mathbb{Z}_p$ as well as $H(\cdot)$.

- **Sign**(SK, $m \in \mathcal{M}$) : The signature algorithm outputs $\sigma = H(m)^a \in \mathbb{G}$.
- **Verify**(VK, m, σ) The verification algorithm tests if $e(\sigma, g) \stackrel{?}{=} e(A, H(m))$ and outputs accept if and only if this holds.




Fig. 5. BLS Selective Hash




Fig. 6. BLS Selective Hash*

Remark 2. The confined PRFs from [BW13] use the GGM tree and get PRFs in range $\{0, 1\}^n$ for some n , whereas our PRFs need to hash to \mathbb{Z}_p . One can achieve a punctured PRF for the proper range by simply setting $n > 2 \lg(p)$ and taking interpreting the GGM output as an integer that is then mod by p . This is sufficient since sampling an integer in $[0, 2^n - 1]$ and then reducing it mod p is statistically close to choosing an integer in $[0, p - 1]$.

Theorem 5. *If our obfuscation scheme is indistinguishability secure, F is a secure punctured PRF, and the computational Diffie-Hellman problem holds in bilinear groups, then the above signature scheme is selectively secure.*

In the full version [HSW14], we describe a proof as a sequence of hybrid experiments where the first hybrid corresponds to the original signature security game. We prove that a poly-time attacker’s advantage must be negligibly close between each successive one. Then, we show that any poly-time attacker in the final experiment that succeeds in forging with non-negligible probability can be used to break the computational Diffie-Hellman assumption in bilinear groups.

6 Adaptively Secure BLS Signatures

In the full version [HSW14], we give a hash function for BLS signatures that can be used to prove adaptive (or standard) security. Our construction is identical to that given in Section 5 with the exception of how the setup creates the hash function. Our proof structure will follow in a similar path to that of our adaptively secure RSA full domain hash signatures in Section 4. In particular, we will again apply an admissible hash function to partition the message space in our proof. At the same time, there are important distinctions and corresponding challenges that arise in this setting as discussed in [HSW14]. Our proof of security relies on indistinguishability obfuscation and the Diffie-Hellman Inversion Assumption.

7 Extensions to Boneh-Franklin IBE and Aggregate Signatures

Boneh-Franklin IBE. We can adapt our techniques for proving security of BLS signatures to the Boneh-Franklin [BF01] Identity-Based Encryption system. BLS signatures directly correspond to IBE private keys in the BF scheme. The proof for the BF adapts with a few minor changes:

- For proving BF selectively secure we can use the Decisional Bilinear Diffie-Hellman assumption.
- The second random oracle in the BF IBE can be replaced with an extractor.
- For proving adaptive security we use the following assumption. Namely that given $g, g^s, g^a, g^{a^2}, \dots, g^{a^n}$ it is hard to distinguish $e(g, g)^{a^{n+1}s}$ from a random group element in \mathbb{G}_T . We note this assumption is weaker than the decision Bilinear Diffie-Hellman Exponent assumption [BGW05].

BGLS Aggregate Signatures. Boneh, Gentry, Lynn and Shacham [BGLS03] showed that the BLS signatures are aggregateable by reduction to the BDH assumption. Later Bellare, Namprempre and Neven [BNN07] showed how an aggregate signature scheme could be built directly from and reduced to the security of BLS signatures. Using their results we immediately get an aggregate signature scheme.

Acknowledgments. We thank Mihir Bellare for discussions relating to the origins and terminology of full domain hash signatures and other helpful discussions. We thank Dan Boneh for many helpful discussions and also for pointing out the equivalence of our assumption used for adaptive BLS case to the Diffie-Hellman

Inversion Assumption. We thank Dennis Hofheinz for clarifications on admissible hash functions and pointing us to the simplified version we used. Finally, we are grateful to the anonymous reviewers of Eurocrypt 2014 for their helpful comments.

References

- [BB04a] 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)
- [BB04b] 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)
- [BBP04] Bellare, M., Boldyreva, A., Palacio, A.: An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 171–188. Springer, Heidelberg (2004)
- [BF01] 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)
- [BGI⁺01] Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer, Heidelberg (2001)
- [BGI⁺12] Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. *J. ACM* 59(2), 6 (2012)
- [BGI14] Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer, Heidelberg (2014)
- [BGLS03] Boneh, D., Gentry, C., Lynn, B., Shacham, H.: Aggregate and verifiably encrypted signatures from bilinear maps. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 416–432. Springer, Heidelberg (2003)
- [BGW05] Boneh, D., Gentry, C., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 258–275. Springer, Heidelberg (2005)
- [BHK13] Bellare, M., Hoang, V.T., Keelveedhi, S.: Instantiating Random Oracles via UCes. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 398–415. Springer, Heidelberg (2013)
- [BLS01] Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 514–532. Springer, Heidelberg (2001)
- [BNN07] Bellare, M., Namprempre, C., Neven, G.: Unrestricted aggregate signatures. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 411–422. Springer, Heidelberg (2007)
- [Boy08] Boyen, X.: A tapestry of identity-based encryption: practical frameworks compared. *IJACT* 1(1), 3–21 (2008)
- [BR93] Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security, pp. 62–73 (1993)
- [BR96] Bellare, M., Rogaway, P.: The exact security of digital signatures - how to sign with RSA and Rabin. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 399–416. Springer, Heidelberg (1996)

- [BW13] Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (2013)
- [Can97] Canetti, R.: Towards realizing random oracles: Hash functions that hide all partial information. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 455–469. Springer, Heidelberg (1997)
- [CGH98] Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited (preliminary version). In: STOC, pp. 209–218 (1998)
- [CHK07] Canetti, R., Halevi, S., Katz, J.: A forward-secure public-key encryption scheme. *J. Cryptology* 20(3), 265–294 (2007)
- [Coc01] 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)
- [CS98] Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
- [DHT12] Dodis, Y., Haitner, I., Tentes, A.: On the instantiability of hash-and-sign RSA signatures. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 112–132. Springer, Heidelberg (2012)
- [DOP05] Dodis, Y., Oliveira, R., Pietrzak, K.: On the generic insecurity of the full domain hash. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 449–466. Springer, Heidelberg (2005)
- [FHPS13] Freire, E.S.V., Hofheinz, D., Paterson, K.G., Striecks, C.: Programmable hash functions in the multilinear setting. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 513–530. Springer, Heidelberg (2013)
- [GGH⁺13] Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: FOCS (2013)
- [GGM84] Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: FOCS, pp. 464–479 (1984)
- [GK03] Goldwasser, S., Kalai, Y.T.: On the (in)security of the Fiat-Shamir paradigm. In: FOCS, pp. 102–113 (2003)
- [HSW13] Hohenberger, S., Sahai, A., Waters, B.: Full domain hash from (leveled) multilinear maps and identity-based aggregate signatures. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 494–512. Springer, Heidelberg (2013)
- [HSW14] Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In: Eurocrypt (2014), Full version available at <http://eprint.iacr.org/2013/509>
- [KPTZ13] Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: ACM Conference on Computer and Communications Security, pp. 669–684 (2013)
- [KS98] Kaliski, B., Staddon, J.: PKCS #1: RSA Cryptography Specifications Version 2.0 (1998)
- [RSA78] Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21(2), 120–126 (1978)
- [Sha83] Shamir, A.: On the generation of cryptographically strong pseudorandom sequences. *ACM Trans. Comput. Syst.* 1(1), 38–44 (1983)
- [SW13] Sahai, A., Waters, B.: How to use indistinguishability obfuscation: Deniable encryption, and more. *Cryptology ePrint Archive*, Report 2013/454 (2013) (to appear in STOC, 2014), <http://eprint.iacr.org/>