

Optimal Security Proofs for Full Domain Hash, Revisited

Saqib A. Kakvi¹ and Eike Kiltz¹

Faculty of Mathematics, Horst Görtz Institute for IT Security
Ruhr University Bochum
{saqib.kakvi,eike.kiltz}@rub.de

Abstract. RSA Full Domain Hash (RSA-FDH) is a digital signature scheme, secure against chosen message attacks in the random oracle model. The best known security reduction from the RSA assumption is non-tight, i.e., it loses a factor of q_s , where q_s is the number of signature queries made by the adversary. It was furthermore proved by Coron (EUROCRYPT 2002) that a security loss of q_s is optimal and cannot possibly be improved. In this work we uncover a subtle flaw in Coron's impossibility result. Concretely, we show that it only holds if the underlying trapdoor permutation is *certified*. Since it is well known that the RSA trapdoor permutation is (for all practical parameters) not certified, this renders Coron's impossibility result moot for RSA-FDH. Motivated by this, we revisit the question whether there is a tight security proof for RSA-FDH. Concretely, we give a new tight security reduction from a stronger assumption, the Phi-Hiding assumption introduced by Cachin et al (EUROCRYPT 1999). This justifies the choice of smaller parameters in RSA-FDH, as it is commonly used in practice. All of our results (positive and negative) extend to the probabilistic signature scheme PSS.

1 Introduction

Among all digital signatures schemes based on the RSA problem, arguably among the most important ones is RSA Full Domain Hash (RSA-FDH) by Bellare and Rogaway [3]. It is extensively used in a wide variety of applications, and serves as the basis of several existing standards such as PKCS #1 [26]. It has been demonstrated by means of a security reduction that, in the random oracle model [2], breaking the security of RSA-FDH (in the sense of existential unforgeability against chosen message attacks) is asymptotically at least as hard as inverting the RSA function.

The seminal work by Bellare and Rogaway introduced the concept of concrete security [3] and highlights the importance of considering the tightness of a security reduction. A security reduction is *tight* if an adversary breaking the scheme yields another adversary breaking the underlying hardness assumption with roughly the same success probability and running time. The current state of RSA-FDH is as follows. Coron's reduction [11] (which improves on earlier results by Bellare and Rogaway [3]) bounds the probability ε of breaking RSA-FDH in time t by $\varepsilon' \cdot q_s$, where ε' is the probability of inverting RSA in time

$t' \approx t$ and q_s is the number of signature queries by the forger. In other words, the security reduction for RSA-FDH is loose (it loses a factor of q_s), which can have great negative impact on the practical parameter choices of the scheme. As a numerical example, for 80 bits of security and assuming that an adversary can make up to $q_s = 2^{30}$ signature queries [3], one should use a large enough RSA modulus N such that inverting the RSA function cannot be done in fewer than $2^{110} = 2^{30} \cdot 2^{80}$ operations. Concretely, using the recommended key sizes from [28], this leads to a modulus N of about 2432 bits, compared to 1248 bits if RSA-FDH had a tight reduction. We further refer to [9] for a recent discussion on the practical impact of non-tight security reductions in cryptography.

It is an interesting question of great practical impact whether or not there is a tight security reduction for general FDH signatures (based on any trapdoor permutation TDP) and, in particular, for RSA-FDH. Unfortunately, this question was already answered to the negative exactly 10 years ago by Coron [12, 13] who showed that the above non-tight security reduction is essentially *optimal*. That is, every security reduction from inverting the TDP (i.e., RSA in the case of RSA-FDH) to breaking FDH signatures will inevitably lose a q_s factor. Consequently, for RSA-FDH a large RSA modulus seems unavoidable to obtain a meaningful security proof.

1.1 An overview of our results

REVISITING CORON’S IMPOSSIBILITY RESULT. We uncover a gap in Coron’s result about the impossibility of a tight security reduction for FDH signatures [13]. As acknowledged by the author of [13], his impossibility result only holds if the underlying trapdoor permutation (i.e., RSA in the case of RSA-FDH) is a *certified trapdoor permutation*. A trapdoor permutation is certified [5, 22] if one can publicly verify that it actually defines a permutation. Unfortunately, the RSA trapdoor permutation is not known to be certified (unless the public exponent e is prime and larger than the modulus N) and therefore the impossibility result does not apply any longer to the case of RSA-FDH.

A TIGHT SECURITY REDUCTION FOR FDH SIGNATURES. In light of the above, we revisit the question whether there exists a tight security reduction for FDH signatures. Unfortunately, we are not able to give such a tight security reduction from the assumption that the TDP is one-way, but from a stronger (yet still non-interactive) assumption, namely that the TDP is lossy (in the sense of Peikert and Waters [25]). Our main result (Theorem 8) shows that there is a *tight* security reduction from the lossiness of the TDP to breaking security of FDH, in the random oracle model.

APPLICATIONS TO RSA-FDH. Recently, Kiltz et al. [20] showed that the RSA trapdoor permutation is lossy under the under the Φ -Hiding Assumption. The Φ -Hiding Assumption was introduced by Cachin, Micali, and Stadler in 1999 [8] and it states that, roughly, (N, e) with $\gcd(\varphi(N), e) = 1$ and $e < N^{1/4}$ is computationally indistinguishable from (N', e') with $e' \mid \varphi(N')$. (Here $\varphi(N)$ is

Euler’s totient function.) This give a tight security reduction for RSA-FDH from the Φ -Hiding Assumption. We remark that the Φ -Hiding Assumption (or, more generally, the assumption that RSA is lossy) is a stronger assumption than the assumption that RSA is one-way. However, it dates back to 1999 [8] and has ever since been used in a number of cryptographic applications (e.g., [20, 7, 14, 16, 23, 18]). It has been cryptanalyzed (e.g. [8, 7, 27]) and for the parameters of interest there is no known algorithm that breaks it without first factoring the modulus $N = pq$. The common interpretation is that the Φ -Hiding Assumption can in practice be viewed as *as hard as factoring* and hence gives a theoretical justification as to why RSA-FDH with a small modulus N is secure in practice.

1.2 Full Domain Hash and Coron’s Impossibility Result

Recall that FDH signatures on a message m is $\sigma = f^{-1}(H(m))$, where f is the public description of the TDP and H is a hash function modelled as a random oracle. A reduction \mathcal{R} that reduces inverting the TDP to breaking FDH inputs a challenge instance $(f, y = f(x))$ of the TDP and generates a public-key for FDH that is passed to a forger \mathcal{F} attacking FDH signatures. Next, \mathcal{F} makes a number of signature queries (which are answered by \mathcal{R}) and finally outputs a forgery. Finally, \mathcal{R} uses the gathered information to invert the TDP, i.e., to compute $x = f^{-1}(y)$. Reduction \mathcal{R} is *tight* if the success probability of \mathcal{R} is roughly the same as the one of \mathcal{F} .

Coron’s impossibility result shows that any reduction \mathcal{R} from inverting the TDP f to breaking FDH which is tight (i.e., does not lose more than a factor q_s) can be turned into an efficient inverting algorithm \mathcal{I} for the TDP f (that works without forger \mathcal{F}). In a nutshell, the argument is as follows. Given an instance of the TDP, the inverter \mathcal{I} runs reduction \mathcal{R} providing it with a simulated forger \mathcal{F} by making a number of hash queries and then signature queries to \mathcal{R} . Next, \mathcal{I} rewinds reduction \mathcal{R} to an earlier state (after the hash queries) and uses one of the signed messages/signature pairs (say (m^*, σ^*)) obtained before the rewind as its forgery. To \mathcal{R} , this counts as a valid forgery since after the rewind, \mathcal{I} did not make a signing query on m^* . The central argument is as follows: consider a real forger that is provided with the view as the simulated forger who outputs a forgery σ' on the same message m^* . FDH has *unique signatures*¹ and hence we can argue that σ^* (provided by \mathcal{R} before the rewind) equals σ' (provided by a real forger). Hence \mathcal{R} is convinced it interacts with a real forger and outputs a solution to the TDP instance. Consequently, from \mathcal{R} we were able to construct an algorithm \mathcal{I} that inverts the TDP without using any forger. It is shown by a combinatorial argument that the success probability of \mathcal{I} is non-negative as long as the reduction \mathcal{R} does not loose more than a factor of q_s , the number of signature queries.

THE GAP IN THE PROOF. During the proof of [12, Th. 5] it is silently assumed that the public-key pk generated by reduction \mathcal{R} is a *real public-key*, honestly

¹ A signature scheme has unique signatures if for each message there exists exactly one signature that verifies w.r.t. a given (honestly generated) public-key.

generated by the key-generation algorithm of FDH, i.e., it contains f which described a permutation.² However, that does not necessarily hold, the public-key generated by \mathcal{R} could be anything. In fact, it is possible that the public-key generated by the reduction \mathcal{R} is *fake* in the sense that the FDH signatures are no longer unique relative to this fake pk . Once signatures are no longer unique (with respect to the fake pk), it is possible that a *real forger* outputs a forgery σ' on m^* which is different from σ^* , the one provided by reduction \mathcal{R} before the rewind. In fact, it could be possible that $\sigma^* \neq \sigma'$ is no longer useful for \mathcal{R} in order to solve the RSA instance after the rewind and hence the impossibility result breaks down. In Section 3 we restate (and prove) a corrected version of Coron’s impossibility result. Fortunately, it turns out that Coron’s argument can be salvaged by requiring the trapdoor permutation in FDH to be certified. Note that in case of a certified trapdoor permutation it is not longer possible for the reduction \mathcal{R} to generate a fake public-key and hence signatures are guaranteed to be unique.

1.3 A tight security reduction for FDH signatures

It is precisely the non-uniqueness of FDH signatures with respect to a fake public-key that will allow us to prove a tight security from the lossiness from the lossiness of the TDP (i.e., the Φ -Hiding Assumption in the case of RSA-FDH). Our proof is surprisingly simple and is sketched as follows. In a first step we substitute the trapdoor permutation in public key with a lossy one. We use the programmability of the random oracle to show that this remains unnoticed by the adversary assuming lossiness of the TDP. Note that once the TDP is lossy, FDH signatures (i.e., σ with $f(\sigma) = H(m)$) are not longer unique since, by the definition of lossiness, each $H(m)$ has many pre-images under a lossy f . In the second step we show that any successful forger will be able to find a collision in the TDP, i.e., two values $x \neq \hat{x}$ with $f(x) = f(\hat{x})$, which is again hard assuming lossiness. The full proof is given in Section 3.

For the important case of RSA-FDH this gives a tight security reduction from the Φ -Hiding Assumption, in the random oracle model. The Φ -Hiding Assumption is believed to be true for sufficiently small public RSA exponents $e < N^{1/4-\epsilon}$ [8]. This in particular includes the important low-exponent cases of $e = 3$ and $e = 2^{16} + 1$ since they allow efficient verification of RSA-FDH signatures.³

It is interesting to remark, that, at a conceptual level FDH is the first signature scheme with *unique signatures* and a *tight* security reduction (from a non-interactive assumption).⁴ Previously, only tight security reductions for *randomized* signatures were known (e.g., [3, 17, 6, 15]).

² Such restricted reductions were called *key-preserving reductions* in [24].

³ We stress that our tight proof technically does not give a counter-example to Coron’s impossibility result since our reduction is from the Φ -Hiding Assumption, not the RSA Assumption. However, as corollary the impossibility result would exclude any (even non-tight) equivalence between the Φ -Hiding and the RSA assumption.

⁴ Here we do not count tight security proofs from “tautological assumptions” which are essentially assuming that the signature scheme is secure.

1.4 Extensions

Our observations can also be applied to the probabilistic signature scheme (PSS) [3] which is contained in IEEE P1363a [19], ISO/IEC 9796-2, and PKCS#1 v2.1 [26]. Coron proved that, if $\log_2(q_s)$ bits of random salt is used in PSS, then there is a tight security reduction from the one-wayness of the TDP [12, 13]. Furthermore, Coron also proved that $\log_2(q_s)$ bits of random salt are essentially optimal for a tight security reduction. Our results for PSS are similar to the ones for FDH. We note that Coron’s impossibility proof for PSS contains the same gap as the one in FDH, i.e., it is only correct if the underlying trapdoor permutation is certified. However, since PSS (with random salt of arbitrary length) is at least as secure as FDH, we obtain as a corollary from Section 3 a tight security proof from lossiness to the security of PSS, with random salt of arbitrary (possibly zero) length.

1.5 Related Work

There is a lot of work on FDH and tightly secure signature schemes, we try to summarize part of it relevant to this work.

TIGHT SECURITY REDUCTION FOR RSA-FDH FROM AN INTERACTIVE ASSUMPTION. Kobitz and Menezes [21, Sec. 3] show a tight reduction from an *interactive assumption* they call the *RSA1 assumption* (which is related to the one-more-RSA assumption RSA-CTI [1]): Given N , e , and a set of $q_s + q_h$ values y_i chosen uniformly from \mathbb{Z}_N , the adversary is permitted adaptively to select up to q_s of those y_i for which he is given solutions x_i to $x_i^e = y_i \bmod N$. The adversary wins if he produces a solution $x_i^e = y_i \bmod N$ for one of the remaining y_i . Even though the RSA1 assumption looks plausible, it is an interactive assumption and almost a tautology for expressing that RSA-FDH signatures are secure in the random oracle model. In fact, our tight security proof for RSA-FDH also serves to show a tight reduction from Φ -Hiding to RSA1.

NON-UNIQUE SIGNATURES WITH TIGHT REDUCTIONS. There exists several previous works that build digital signature schemes with a tight security reduction. We stress that all of them have, in contrast to FDH, a randomized signing algorithm, i.e., signatures are not unique. Goh et al. [17] show that adding one single bit of random salt to the hash function of FDH allows to prove a tight security reduction from the RSA assumption. Bernstein [6] shows a tight security reduction for (a certain randomized variant of) Rabin-Williams signature scheme from the factoring assumption. More generally, Gentry et al. [15] introduce the concept of preimage samplable trapdoor functions which are non-injective trapdoor functions with an efficient pre-image sampling algorithm. They further propose a probabilistic variant of FDH and prove it tightly secure. In fact, their proof technique is reminiscent to the second step in our proof of FDH from the lossiness but FDH can not be viewed as an instance of their probabilistic FDH variant.

RSA-OAEP. Recently, [20] used the Φ -Hiding Assumption to show that the RSA function is lossy and used this fact to prove positive instantiability results of RSA-OAEP in the standard model.

1.6 Open problems

On the one hand the Φ -Hiding Assumption is believed to be true for public exponents $e \leq N^{1/4-\epsilon}$ and hence for these values we get a tight security reduction for RSA-FDH. On the other hand, Coron’s impossibility results holds for prime e with $e > N$. This leaves the interesting open problem whether for public exponents $N^{1/4} \leq e \leq N$ there exists a tight security reduction for RSA-FDH (under a reasonable assumption).

2 Definitions

2.1 Notations and conventions

We denote our security parameter as k . For all $n \in \mathbb{N}$, we denote by 1^n the n -bit string of all ones. For any element x in a set S , we use $x \in_R S$ to indicate that we choose x uniformly random in S . All algorithms may be randomized. For any algorithm A , we define $x \leftarrow_{\$} A(a_1, \dots, a_n)$ as the execution of A with inputs a_1, \dots, a_n and fresh randomness and then assigning the output to x . We denote the set of prime numbers by \mathbb{P} and we denote the subset of k -bit primes as \mathbb{P}_k . Similarly, we have the integers denoted by \mathbb{Z} and \mathbb{Z}_k . We denote by \mathbb{Z}_N^* the multiplicative group modulo $N \in \mathbb{Z}$.

2.2 Games

A game (such as in Figure 2) is defined as a collection of procedures, as per the model of [4]. There is an **Initialize** procedure and a **Finalize** procedure, as well a procedure for each separate oracle. Executing a game G with an adversary \mathcal{A} means running the adversary and using the procedures to answer any oracle queries. The adversary must first make one query to **Initialize**. Then it may query the oracles as many times as allowed by the definition of the game. After this, the adversary must then make 1 query to **Finalize**, which is the final procedure call of the game. The output of **Finalize** is denoted by $G^{\mathcal{A}}$. Where the Finalize procedure simply returns the output of the adversary, we omit the Finalize procedure. We use a strongly typed pseudo-code with implicit initialization. Which means all variables maintain their type throughout the execution of the games and they are all implicitly declared and initialized. Boolean flags are initialized to false, numerical types are initialized to 0, sets are initialized to \emptyset . We use the notation $y \leftarrow_{\$} \mathcal{A}(a_1, \dots, a_n)$ to denote invoking the probabilistic algorithm \mathcal{A} with inputs a_1, \dots, a_n and fresh randomness and assigning the output to y .

2.3 Signature schemes

A digital signature is a message-dependant bit string σ , which can only be generated by the signer, using a secret signing key sk and is transmitted with the message. The signature can then be verified by the receiver using a public verification key pk . A digital signature scheme is defined as a triple of probabilistic algorithms $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{Verify})$, which we describe below:

1. **KeyGen** takes as an input the unary representation of our security parameter (1^k) and outputs a signing key sk and verification key pk .
2. **Sign** takes as input a signing key sk , message m and outputs a signature σ .
3. **Verify** is a deterministic algorithm, which on input of a public key and a message-signature pair (m, σ) outputs 1 (accept) or 0 (reject).

We say that SIG is correct if for all public key and secret key pairs generated by **KeyGen**, we have:

$$\Pr[\text{Verify}(pk, m, \text{Sign}(sk, m)) = 1] = 1.$$

We now define **UF-CMA** (unforgeability under chosen message attacks) assuming the signature scheme SIG contains a hash function $h : \{0, 1\}^* \rightarrow \text{Dom}$ which is modeled as a random oracle.

procedure Initialize $(pk, sk) \leftarrow_s \text{KeyGen}(1^k)$ return pk	procedure Sign(m) Game UF-CMA $\mathcal{M} \leftarrow \mathcal{M} \cup (m)$ return $\sigma \leftarrow_s \text{Sign}(sk, y)$
procedure Hash(m) if $(m, \cdot) \in \mathcal{H}$ then fetch $(m, y) \in \mathcal{H}$; return y else $y \in_R \text{Dom}$; $\mathcal{H} \leftarrow \mathcal{H} \cup (m, y)$; return y	procedure Finalize(m^*, σ^*) if $\text{Verify}(pk, m^*, \sigma^*) = 1 \wedge m^* \notin \mathcal{M}$ then return 1 else return 0

Fig. 1. Game defining UF-CMA security in the random oracle model.

We say a signature scheme SIG is $(t, \varepsilon, q_h, q_s)$ -UF-CMA secure in the random oracle model, if for all adversaries \mathcal{A} running in time upto t , making at most q_h hashing and q_s signing oracle queries, they have an advantage of at most ε , where the advantage of \mathcal{A} is defined as:

$$\text{Adv}_{\text{SIG}}^{\text{UF-CMA}}(\mathcal{A}) = \Pr \left[\text{UF-CMA}^{\mathcal{A}} \Rightarrow 1 \right].$$

2.4 Trapdoor Permutations

We recall the definition of trapdoor permutation families.

Definition 1 A family of trapdoor permutations $\text{TDP} = (\text{Gen}, \text{Eval}, \text{Invert})$ consists of following three polynomial-time algorithms.

1. The probabilistic algorithm **Gen**, which on input 1^k outputs a public description pub (which includes an efficiently sampleable domain Dom_{pub}) and a trapdoor td .
2. The deterministic algorithm **Eval**, which on input pub and $x \in \text{Dom}_{pub}$, outputs $y \in \text{Dom}_{pub}$. We write $f(x) = \text{Eval}(pub, x)$.
3. The deterministic algorithm **Invert**, which on input td and $y \in \text{Dom}_{pub}$, outputs $x \in \text{Dom}_{pub}$. We write $f^{-1}(y) = \text{Invert}(pub, y)$.

We require that for all $k \in \mathbb{N}$ and all (pub, td) output by $\text{Gen}(1^k)$, $f(\cdot) = \text{Eval}(pub, \cdot)$ defines a permutation over Dom_{pub} and that for all $x \in \text{Dom}_{pub}$, $\text{Invert}(td, \text{Eval}(pub, x)) = x$.

We want to point out that $f_{pub}(\cdot) = \text{Eval}(pub, \cdot)$ is only required to be a permutation for correctly generated pub but not every bit-string pub necessarily yields a permutation. A family of trapdoor permutations **TDP** is said to be *certified* [5] if the fact that it is a permutation can be verified in polynomial time given pub .

Definition 2 A family of trapdoor permutations **TDP** is called certified if there exists a deterministic polynomial-time algorithm **Certify** that, on input of 1^k and an arbitrary (polynomially bounded) bit-string pub (potentially not generated by **Gen**), returns 1 iff $f(\cdot) = \text{Eval}(pub, \cdot)$ defines a permutation over Dom_{pub} .

We now recall security notion for trapdoor permutations. A trapdoor permutation **TDP** is hard to invert (one-way) if given pub and $f_{pub}(x)$ for uniform $x \in \text{Dom}_{pub}$, it is hard to compute x . More formally, it is (t, ε) -hard to invert if for all adversaries running in time t , $\Pr[\mathcal{A}(pub, \text{Eval}(pub, x)) = x] \leq \varepsilon$, where the probability is taken over $(pub, td) \leftarrow \text{Gen}(1^k)$, $x \in_R \text{Dom}_{pub}$ and the random coin tosses of \mathcal{A} . The following security notion, lossiness [25], is a stronger requirement than one-wayness.

Definition 3 Let $l \geq 2$. A trapdoor permutation **TDP** is a (l, t, ε) lossy trapdoor permutation if the following two conditions hold.⁵

1. There exists a probabilistic polynomial-time algorithm **LossyGen**, which on input 1^k outputs pub' such that the range of $f_{pub'}(\cdot) := \text{Eval}(pub', \cdot)$ under $\text{Dom}_{pub'}$ is at least a factor of l smaller than the domain $\text{Dom}_{pub'}$: $|\text{Dom}_{pub'}|/|f_{pub'}(\text{Dom}_{pub'})| \geq l$. (Note that we measure the lossiness in its absolute value l , i.e., the function has $\lceil \log_2 l \rceil$ bits of lossiness.)
2. All distinguishers \mathcal{D} running in time at most t have an advantage $\text{Adv}_{\text{TDP}}^l(\mathcal{D})$ of at most ε , where

$$\text{Adv}_{\text{TDP}}^l(\mathcal{D}) = \Pr[\mathbf{L}_1^{\mathcal{D}} \Rightarrow 1] - \Pr[\mathbf{L}_0^{\mathcal{D}} \Rightarrow 1].$$

⁵ We deviate in two ways from the original definition of lossy trapdoor functions Peikert and Waters [25]. First, we define the permutation over arbitrary domains Dom , rather than $\{0, 1\}^k$; second, we measure the absolute lossiness l , rather than the bits of lossiness $\ell = \log_2(l)$.

procedure Initialize Game L_0 $(pub, td) \leftarrow_s \text{Gen}(1^k)$ return pub	procedure Initialize Game L_1 $(pub', \perp) \leftarrow_s \text{LossyGen}(1^k)$ return pub'
--	--

Fig. 2. The Lossy Trapdoor Permutation Games.

We say TDP is *regular* (l, t, ε) *lossy* if TDP is (l, t, ε) lossy and all functions $f_{pub'}(\cdot) = \text{Eval}(pub', \cdot)$ generated by LossyGen are l -to-1 on $\text{Dom}_{pub'}$.

2.5 The RSA trapdoor permutation

We define the RSA trapdoor permutation $\text{RSA} = (\text{RSAGen}, \text{RSAEval}, \text{RSAInv})$ as follows. The RSA instance generator $\text{RSAGen}(1^k)$ outputs $pub = (N, e)$ and $td = d$, where $N = pq$ is the product of two $k/2$ -bit primes, $\text{gcd}(e, \varphi(N)) = 1$, and $d = e^{-1} \bmod \varphi(N)$. The domain is $\text{Dom}_{pub} = \mathbb{Z}_N^*$. The evaluation algorithm $\text{RSAEval}(pub, x)$ returns $f_{pub}(x) = x^e \bmod N$, the inversion algorithm $\text{RSAInv}(td, y)$ returns $f_{pub}^{-1}(y) = y^d \bmod N$. The standard assumption is that RSA is hard to invert. We will review the (regular) lossiness of RSA in Section 4.

3 Full Domain Hash Signatures

3.1 The Scheme

For a family of trapdoor permutations $\text{TDP} = (\text{Gen}, \text{Eval}, \text{Invert})$ we define the Full Domain Hash (TDP-FDH) signature scheme [3] in Figure 3.

procedure KeyGen $(pub, td) \leftarrow_s \text{Gen}(1^k)$ Pick a hash function $h : \{0, 1\}^* \rightarrow \text{Dom}_{pub}$ return $(pk = (h, pub), sk = td)$	TDP-FDH
procedure Sign (sk, m) return $\sigma = \text{Invert}(td, h(m))$	$// \sigma = f_{pub}^{-1}(h(m))$
procedure Verify (pk, m, σ) if $\text{Eval}(pub, \sigma) = h(m)$ then return 1 else return 0	$// f_{pub}(\sigma) \stackrel{?}{=} h(m)$

Fig. 3. The Full Domain Hash Signature Scheme TDP-FDH.

3.2 Classical Security Results of TDP-FDH

The original reduction by Bellare and Rogaway from one-wayness of TDP loses a factor of $(q_h + q_s)$ [3], which was later improved by Coron to a factor of q_s [11] for the case of the RSA trapdoor permutation.

Theorem 4 *Assume the trapdoor permutation RSA is (t', ε') -hard to invert. Then for any (q_h, q_s) , RSA-FDH is $(t, \varepsilon, q_h, q_s)$ -UF-CMA secure in the Random Oracle Model, where*

$$\begin{aligned}\varepsilon' &= \frac{\varepsilon}{q_s} \cdot \left(1 - \frac{1}{q_s + 1}\right)^{q_s+1} \approx \frac{\varepsilon}{q_s} \cdot \exp(-1) \\ t' &= t + (q_h + q_s + 1) \cdot \mathcal{O}(k^3).\end{aligned}$$

3.3 A corrected version of Coron's optimality result

Coron showed that a security loss of a factor q_s (times some constant) is essentially optimal for TDP-FDH [12, 13]. To state a corrected version of Coron's impossibility result, we first recall the following definitions [12].

Definition 5 We say a reduction \mathcal{R} $(t_{\mathcal{F}}, t_{\mathcal{R}}, q_h, q_s, \varepsilon_{\mathcal{F}}, \varepsilon_{\mathcal{R}})$ -reduces solving a hard problem to breaking $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ if after running a forger \mathcal{F} that $(t_{\mathcal{F}}, q_h, q_s, \varepsilon_{\mathcal{F}})$ -breaks SIG , the reduction outputs a solution of the problem with probability at least $\varepsilon_{\mathcal{R}}$, with running time at most $t_{\mathcal{R}}$.

Definition 6 A signature scheme $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ is said to be a unique signature scheme if for every public key pk output by KeyGen , for every message m there exists exactly one bit-string $\sigma \in \{0, 1\}^*$ such that $\text{Verify}(pk, m, \sigma) = 1$.

We now state the corrected version of Coron's impossibility result which we prove in the full version of this paper.

Theorem 7 *Suppose TDP is a certified trapdoor permutation. Let \mathcal{R} be a reduction that $(t_{\mathcal{F}}, t_{\mathcal{R}}, q_h, q_s, \varepsilon_{\mathcal{F}}, \varepsilon_{\mathcal{R}})$ -reduces breaking one-wayness of TDP to breaking UF-CMA security of TDP-FDH. If \mathcal{R} runs the forger only once, then we can build an inverter \mathcal{I} which $(t_{\mathcal{I}}, \varepsilon_{\mathcal{I}})$ -breaks one-wayness of TDP with:*

$$\begin{aligned}t_{\mathcal{I}} &\leq 2 \cdot t_{\mathcal{R}} \\ \varepsilon_{\mathcal{I}} &\geq \varepsilon_{\mathcal{R}} - \varepsilon_{\mathcal{F}} \cdot \frac{\exp(-1)}{q_s} \cdot \left(1 - \frac{q_s}{q_h}\right)^{-1}.\end{aligned}$$

Hence, from a security reduction from one-wayness to the security of TDP-FDH which loses less than a factor of q_s , one obtains an efficient inverter \mathcal{I} for TDP.

3.4 A Tight Security Proof for TDP-FDH

The impossibility result of Theorem 7 only holds for TDP-FDH if TDP is certified trapdoor permutation. However if TDP is not certified, this leaves room for a tight proof for TDP-FDH. We now state our main result, namely that TDP-FDH is tightly secure assuming TDP is regular lossy.

Theorem 8 Assume $TDP = (\text{Gen}, \text{Eval}, \text{Invert})$ is a regular (l, t', ε') -lossy trapdoor permutation for $l \geq 2$. Then, for any (q_h, q_s) , TDP-FDH is $(t, \varepsilon, q_h, q_s)$ -UF-CMA secure in the Random Oracle Model, where

$$\varepsilon = \left(\frac{2l-1}{l-1} \right) \cdot \varepsilon'$$

$$t = t' - q_h \cdot T_{TDP}$$

and T_{TDP} is the time to evaluate TDP.

Proof. Let \mathcal{A} be an adversary that runs in time t against TDP-FDH executed in the UF-CMA experiment described in G_0 in Figure 1 with $\varepsilon = \Pr[G_0^{\mathcal{A}} \Rightarrow 1]$. Here we assume wlog that \mathcal{A} always makes a query to $\text{Hash}(m)$ before calling $\text{Sign}(m)$ or $\text{Finalize}(m, \cdot)$.

procedure Initialize	Game G_0 = (UF-CMA)	procedure Initialize	Games G_1 - G_4
$(pub, td) \leftarrow_s \text{Gen}(1^k)$ Return $pk = pub$		$(pub, td) \leftarrow_s \text{Gen}(1^k)$ $(pub, \perp) \leftarrow_s \text{LossyGen}(1^k)$ Return $pk = pub$	// G_1, G_4 // G_2, G_3
procedure Hash (m) if $(m, \cdot) \in \mathcal{H}$ then fetch (m, y_m) ; return y_m else $y_m \in_R \text{Dom}_{pub}$ $\mathcal{H} \leftarrow \mathcal{H} \cup (m, y_m)$; return y_m		procedure Hash (m) if $m \in \mathcal{H}$ lookup $(m, y_m, \sigma_m) \in \mathcal{H}$ return y_m else $\sigma_m \in_R \text{Dom}_{pub}$ $y_m = \text{Eval}(pub, \sigma_m)$ $\mathcal{H} \leftarrow \mathcal{H} \cup (m, y_m, \sigma_m)$; return y_m	
procedure Sign (m) $\mathcal{M} \leftarrow \mathcal{M} \cup (m)$ return $\sigma_m = \text{Invert}(td, h(m))$		procedure Sign (m) $\mathcal{M} \leftarrow \mathcal{M} \cup (m)$ lookup $(m, y_m, \sigma_m) \in \mathcal{H}$, return σ_m	
Procedure Finalize (m^*, σ^*) if $(\text{Verify}(pub, m^*, \sigma^*) = 1) \wedge (m^* \notin \mathcal{M})$ return 1 else return 0		Procedure Finalize (m^*, σ^*) lookup $(m^*, y_{m^*}, \sigma_{m^*}) \in \mathcal{H}$ // G_3, G_4 if $\sigma_{m^*} = \sigma^*$ then BAD = true return 0 // G_3, G_4 if $\text{Verify}(pub, m^*, \sigma^*) = 1 \wedge (m^* \notin \mathcal{M})$ return 1 else return 0	

Table 1. Games for the proof of Theorem 8.

Lemma 9 $\Pr[G_0^{\mathcal{A}} \Rightarrow 1] = \Pr[G_1^{\mathcal{A}} \Rightarrow 1]$.

Proof. In G_0 , we modelled the hash function as a random oracle. In G_1 we modify the random oracle and the signing queries. On any m the random oracle now

works by evaluating the permutation on a random element $\sigma_m \in \text{Dom}_{pub}$. We then modify the signing oracle to return this element σ_m . Note that signing no longer requires the trapdoor td . It can be seen that all our signatures will verify due to the fact that $\text{Eval}(pub, \sigma_m) = y_m$ for all m . Thus our simulation of the signatures is correct. Since TDP is a permutation, the distribution of our hash queries in \mathbf{G}_1 is identical to the distribution in \mathbf{G}_0 . Thus we have $\Pr[\mathbf{G}_0^A \Rightarrow 1] = \Pr[\mathbf{G}_1^A \Rightarrow 1]$.

Lemma 10 *There exists a distinguisher \mathcal{D}_1 against the lossiness of TDP, which runs in time $t = t_{\mathcal{A}} + q_h \cdot T_{\text{TDP}}$ and that $\Pr[\mathbf{G}_1^A \Rightarrow 1] - \Pr[\mathbf{G}_2^A \Rightarrow 1] = \text{Adv}_{\text{TDP}}^L(\mathcal{D}_1)$.*

Proof. From \mathbf{G}_1 to \mathbf{G}_2 , we change the key generation from a normal permutation to a lossy permutation, however the oracles are identical in both games. We now build a distinguisher \mathcal{D}_1 against the lossiness of TDP, using these games. The distinguisher will run \mathcal{A} and simulates the oracles $\text{Sign}(\cdot), \text{Hash}(\cdot)$ as described in games \mathbf{G}_1 & \mathbf{G}_2 , for which it requires time $q_h \cdot T_{\text{TDP}}$. Note that \mathcal{D}_1 does not require the trapdoor td to simulate the oracles. After \mathcal{A} calls `Finalize`, \mathcal{D}_1 returns the inverse of `Finalize`. Thus we can see that $\Pr[\mathbf{L}_0^{\mathcal{D}_1} \Rightarrow 1] = 1 - \Pr[\mathbf{G}_1^A \Rightarrow 1]$. Similarly, we have $\Pr[\mathbf{L}_1^{\mathcal{D}_1} \Rightarrow 1] = 1 - \Pr[\mathbf{G}_2^A \Rightarrow 1]$. Hence we have $\Pr[\mathbf{G}_1^A \Rightarrow 1] - \Pr[\mathbf{G}_2^A \Rightarrow 1] = (1 - \Pr[\mathbf{L}_0^{\mathcal{D}_1} \Rightarrow 1]) - (1 - \Pr[\mathbf{L}_1^{\mathcal{D}_1} \Rightarrow 1]) = \Pr[\mathbf{L}_1^{\mathcal{D}_1} \Rightarrow 1] - \Pr[\mathbf{L}_0^{\mathcal{D}_1} \Rightarrow 1] = \text{Adv}_{\text{TDP}}^L(\mathcal{D}_1)$.

Lemma 11 $\Pr[\mathbf{G}_3^A \Rightarrow 1] = \left(\frac{l-1}{l}\right) \Pr[\mathbf{G}_2^A \Rightarrow 1]$.

Proof. In \mathbf{G}_3 , we introduce a new rule, which sets BAD to true if the forgery σ^* provided by \mathcal{A} is the same as the simulated signature σ_{m^*} for the target message m^* . If this is the case, the adversary loses the game, i.e., \mathbf{G}_3 outputs 0. σ_{m^*} is independent of \mathcal{A} 's view and is uniformly distributed in set of pre-images of y_{m^*} . Due to the l regular lossiness of TDP, the probability of a collision is equal to exactly $1/l$. Thus we see that the BAD rule reduces the probability of the adversary winning the game by $1/l$, hence $\Pr[\mathbf{G}_3^A \Rightarrow 1] = (1 - \frac{1}{l}) \Pr[\mathbf{G}_2^A \Rightarrow 1] = \left(\frac{l-1}{l}\right) \Pr[\mathbf{G}_2^A \Rightarrow 1]$.

Lemma 12 *There exists a distinguisher \mathcal{D}_2 against the lossiness of TDP, which runs in time $t = t_{\mathcal{A}} + q_h \cdot T_{\text{TDP}}$ and that $\Pr[\mathbf{G}_3^A \Rightarrow 1] - \Pr[\mathbf{G}_4^A \Rightarrow 1] = \text{Adv}_{\text{TDP}}^L(\mathcal{D}_2)$.*

Proof. From \mathbf{G}_3 to \mathbf{G}_4 , we change the key generation from a lossy permutation to a normal permutation, however the oracles are identical in both games. We now build a distinguisher \mathcal{D}_2 against the lossiness of TDP, using these games. The distinguisher will act as the challenger to \mathcal{A} . It will simulate the oracles as described in games \mathbf{G}_3 & \mathbf{G}_4 , for which it requires time $q_h \cdot T_{\text{TDP}}$. After \mathcal{A} calls `Finalize`, \mathcal{D}_2 returns the output of `Finalize`. We can see that $\Pr[\mathbf{G}_4^A \Rightarrow 1] = \Pr[\mathbf{L}_0^{\mathcal{D}_2} \Rightarrow 1]$. Similarly, we have $\Pr[\mathbf{G}_3^A \Rightarrow 1] = \Pr[\mathbf{L}_1^{\mathcal{D}_2} \Rightarrow 1]$. Hence we have $\Pr[\mathbf{G}_3^A \Rightarrow 1] - \Pr[\mathbf{G}_4^A \Rightarrow 1] = \Pr[\mathbf{L}_1^{\mathcal{D}_2} \Rightarrow 1] - \Pr[\mathbf{L}_0^{\mathcal{D}_2} \Rightarrow 1] = \text{Adv}_{\text{TDP}}^L(\mathcal{D}_2)$.

Lemma 13 $\Pr[\mathbf{G}_4^A \Rightarrow 1] = 0$.

Proof. In G_4 we again use the original KeyGen such that $\text{Eval}(\text{pub}, \cdot)$ defines a permutation. This means that our signing function is now a permutation, thus any forgery implies a collision. Therefore whenever the adversary is able to make a forgery, the game outputs 0 due to the BAD rule. Whenever they are unable to make a forgery, the game outputs 0. Thus we can see that in all cases, the game will output 0, hence $\Pr[\mathsf{G}_4^A \Rightarrow 1] = 0$.

We combine Lemmas 9 to 13 to get:

$$\Pr[\mathsf{G}_0^A \Rightarrow 1] = \mathbf{Adv}_{\text{TDP}}^L(\mathcal{D}_1) + \left(\frac{l}{l-1}\right)\mathbf{Adv}_{\text{TDP}}^L(\mathcal{D}_2).$$

where l is the lossiness of TDP. Because the distinguishers run in the same time, we know that both distinguishers can have at most an advantage of ε' , giving us:

$$\varepsilon \leq \frac{2l-1}{l-1} \cdot \varepsilon'.$$

This completes the proof.

4 Lossiness of RSA from the Φ -Hiding Assumption

4.1 Lossiness of RSA

The lossiness of RSA for a number of specific instance generators RSAGen was first considered in [20]. We now recall (and extend) some of the results from [20].

First, we recall some definitions from [20]. We denote by $\mathcal{RSA}_k := \{(N, p, q) \mid N = pq, p, q \in \mathbb{P}_{k/2}\}$ the set of all the tuples (N, p, q) such that $N = pq$ is the product of two distinct $k/2$ -bit primes. Such an N is called an RSA modulus. By $(N, p, q) \in_R \mathcal{RSA}_k$ we mean the (N, p, q) is sampled according to the uniform distribution on \mathcal{RSA}_k . Let R be some relation on p and q . By $\mathcal{RSA}_k[R]$, we denote the subset of \mathcal{RSA}_k such that the relation R holds on p and q . For example, let e be a prime. Then $\mathcal{RSA}_k[p = 1 \bmod e]$ is the set of all (N, p, q) , where where $N = pq$ is the product of two distinct $k/2$ -bit primes p, q and $p = 1 \bmod e$. That is, the relation $R(p, q)$ is true if $p = 1 \bmod e$ and q is arbitrary. By $(N, p, q) \in_R \mathcal{RSA}_k[R]$ we mean that (N, p, q) is sampled according to the uniform distribution on $\mathcal{RSA}_k[R]$.

α - Φ -HIDING ASSUMPTION. We recall a variant of the Φ -Hiding Assumption introduced by Cachin, Micali and Stadler [8], where we build on a formalization by Kiltz, O'Neil and Smith [20]. The main statement of the assumption is that given an k -bit RSA modulus $N = pq$ and a random $\alpha \cdot k$ -bit prime e (where $0 < \alpha < \frac{1}{4}$ is a public constant), it is difficult to decide if $e \mid \varphi(N)$ or if $\gcd(e, \varphi(N)) = 1$. We note that if $e \mid \varphi(N)$ with $e \geq N^{1/4}$, then N can be factored using Coppersmith's attacks [10], see [8] for details. Hence for the Φ -Hiding Assumption to hold, the bit-length of e must not exceed one-fourth of the bit length of N .

Consider a distinguisher \mathcal{D} which plays one of the games P_0 or P_1 defined in Table 2, The advantage of \mathcal{D} is defined as:

$$\mathbf{Adv}^{\Phi\text{H}}(\mathcal{D}) = \Pr[\mathsf{P}_1^{\mathcal{D}} \Rightarrow 1] - \Pr[\mathsf{P}_0^{\mathcal{D}} \Rightarrow 1].$$

procedure Initialize	Game P ₀	procedure Initialize	Game P ₁
$e \in_R \mathbb{P}_{\alpha k}$	$(N, p, q) \in_R \mathcal{RSA}_k[\gcd(e, \varphi(N)) = 1]$	$e \in_R \mathbb{P}_{\alpha k}$	$(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \pmod e]$
$(N, p, q) \in_R \mathcal{RSA}_k[\gcd(e, \varphi(N)) = 1]$		$(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \pmod e]$	
return (N, e)		return (N, e)	

Table 2. The α - Φ -Hiding Assumption Games.

We say that the α - Φ -Hiding Problem is (t, ϵ) -hard if for all distinguishers \mathcal{D} running in time at most t have an advantage of at most ϵ .

Define an RSA instance generator RSAGen as an algorithm that returns (N, e, p, q) sampled as $e \in_R \mathbb{P}_{\alpha k}$ and $(N, p, q) \in_R \mathcal{RSA}_k[\gcd(e, \varphi(N)) = 1]$. (See [20] for details on the sampling algorithm.)

Lemma 14 *If the α - Φ -Hiding Problem is (t, ϵ) -hard, then the RSA = (RSAGen, RSAEval, RSAInv) defines a regular $(2^\alpha, t, \epsilon)$ -lossy trapdoor permutation.*

Proof. If (N, e) is sampled using RSAGen , then $\gcd(e, \varphi(N)) = 1$ and (N, e) defines a permutation $\text{RSA}(x) = x^e \pmod N$ over \mathbb{Z}_N^* . We define LossyGen to be an algorithm that returns (N, e) sampled as $e \in_R \mathbb{P}_{\alpha k}$ and $(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \pmod e]$. If (N, e) is sampled using LossyGen then $e \mid \varphi(N)$ and hence the RSA function is e -to-1 on the domain $\text{Dom}_{\text{pub}} = \mathbb{Z}_N^*$. By definition, the outputs of RSAGen and LossyGen are indistinguishable if the α - Φ -Hiding Problem is hard.

FIXED-PRIME Φ -HIDING ASSUMPTION. In practice, e is chosen to be small and is generally fixed to some specific numbers, such as $e = 3$ or $e = 2^{16} + 1$, which allows for fast exponentiation. We now show a minor variant of the α - Φ -Hiding Assumption for *fixed primes* e , where our formalization relies on discussions from [8] and [20, Footnote 9].

First, we discuss the special case of $e = 3$. We define our RSA instance RSAGen_3 generator as an algorithm that samples (N, p, q) uniformly from $\mathcal{RSA}_k[p = 2 \pmod 3, q = 2 \pmod 3]$, which is equivalent to $\mathcal{RSA}_k[\gcd(3, \varphi(N)) = 1]$. We note that $N \pmod 3$ is always 1. This means that for the lossy case, we must also ensure the $N \pmod 3 = 1$, otherwise there would be a simple distinguisher. To ensure this is to have 3 divide both $p - 1$ and $q - 1$. Thus, our lossy keys are sampled from the $\mathcal{RSA}_k[p = 1 \pmod 3, q = 1 \pmod 3]$.

procedure Initialize	Game 3F ₀	procedure Initialize	Game 3F ₁
$(N, p, q) \in_R \mathcal{RSA}_k[\gcd(3, \varphi(N)) = 1]$	$(N, p, q) \in_R \mathcal{RSA}_k[\gcd(3, \varphi(N)) = 1]$	$(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \pmod 3, q = 1 \pmod 3]$	$(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \pmod 3, q = 1 \pmod 3]$
return $(N, e = 3)$		return $(N, e = 3)$	

Table 3. The Fixed-Prime Φ -Hiding Assumption Games

Consider a distinguisher \mathcal{D} which plays one of the games in Table 3. The advantage of \mathcal{D} is defined as

$$\mathbf{Adv}^{\mathbb{F}\Phi\mathbb{H}}(\mathcal{D}) = \Pr[3\mathbb{F}_1^{\mathcal{D}} \Rightarrow 1] - \Pr[3\mathbb{F}_0^{\mathcal{D}} \Rightarrow 1].$$

We say that the Fixed-Prime Φ -Hiding Problem, with $e = 3$, is (t, ϵ) -hard if all distinguishers running in time at most t have an advantage of at most ϵ .

Lemma 15 *If the Fixed-Prime Φ -Hiding Problem, with $e = 3$, is (t, ϵ) -hard, then the $\text{RSA}_3 = (\text{RSAGen}_3, \text{RSAEval}, \text{RSAINv})$ defines a regular $(9, t, \epsilon)$ -lossy trapdoor permutation.*

Proof. If $(N, p, q) \in \mathcal{RSA}_k[\text{gcd}(3, \varphi(N)) = 1]$ then $(N, 3)$ clearly makes the RSA function a permutation. If $(N, p, q) \in \mathcal{RSA}_k[p = 1 \pmod{3}, q = 1 \pmod{3}]$ then $3 \mid \varphi(N)$ and hence the RSA function is 9-to-1 on the domain $\text{Dom}_{\text{pub}} = \mathbb{Z}_N^*$.

We now consider the general case of fixed $e > 3$. For this case, we define our RSA instance generator RSAGen_e as an algorithm that samples (N, p, q) from $\mathcal{RSA}_k[\text{gcd}(e, \varphi(N)) = 1]$. We note that $N \pmod{e}$ will be some value between 1 and $e - 1$. This means that for the lossy case, we require e to divide $p - 1$ and not $q - 1$, otherwise we would have a simple distinguisher. Our lossy keys are sampled from $\mathcal{RSA}_k[p = 1 \pmod{e}, q \neq 1 \pmod{e}]$. Consider a distinguisher \mathcal{D}

procedure Initialize	Game \mathbb{F}_0	procedure Initialize	Game \mathbb{F}_1
$(N, p, q) \in_R \mathcal{RSA}_k[\text{gcd}(e, \varphi(N)) = 1]$		$(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \pmod{e}, q \neq 1 \pmod{e}]$	
return (N, e)		return (N, e)	

Table 4. The Fixed-Prime Φ -Hiding Assumption Games

which plays one of the games in Table 4. The advantage of \mathcal{D} is defined as

$$\mathbf{Adv}^{\mathbb{F}\Phi\mathbb{H}}(\mathcal{D}) = \Pr[\mathbb{F}_1^{\mathcal{D}} \Rightarrow 1] - \Pr[\mathbb{F}_0^{\mathcal{D}} \Rightarrow 1].$$

We say that the Fixed-Prime Φ -Hiding Problem, with $e > 3$, is (t, ϵ) -hard if for all distinguishers running in time at most t have an advantage of at most ϵ .

Lemma 16 *If the Fixed-Prime Φ -Hiding Problem, with $e > 3$, is (t, ϵ) -hard, then $\text{RSA}_e = (\text{RSAGen}_e, \text{RSAEval}, \text{RSAINv})$ defines a regular (e, t, ϵ) -lossy trapdoor permutation.*

Proof. If $(N, p, q) \in \mathcal{RSA}_k[\text{gcd}(e, \varphi(N)) = 1]$ then (N, e) clearly defines a permutation. If $(N, p, q) \in \mathcal{RSA}_k[p = 1 \pmod{e}, q \neq 1 \pmod{e}]$ then $e \mid \varphi(N)$ and hence the RSA function is e -to-1 on the domain $\text{Dom}_{\text{pub}} = \mathbb{Z}_N^*$.

Acknowledgements

We thank Mihir Bellare and Dennis Hofheinz for valuable comments on an earlier draft.

References

1. M. Bellare, C. Namprempre, D. Pointcheval, and M. Semanko. The one-more-RSA-inversion problems and the security of Chaum's blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003.
2. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73. ACM Press, Nov. 1993.
3. M. Bellare and P. Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In U. M. Maurer, editor, *Advances in Cryptology – EUROCRYPT'96*, volume 1070 of *Lecture Notes in Computer Science*, pages 399–416. Springer, May 1996.
4. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, May / June 2006.
5. M. Bellare and M. Yung. Certifying permutations: Noninteractive zero-knowledge based on any trapdoor permutation. *Journal of Cryptology*, 9(3):149–166, 1996.
6. D. J. Bernstein. Proving tight security for Rabin-Williams signatures. In N. P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 70–87. Springer, Apr. 2008.
7. C. Cachin. Efficient private bidding and auctions with an oblivious third party. In *ACM CCS 99: 6th Conference on Computer and Communications Security*, pages 120–127. ACM Press, Nov. 1999.
8. C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. In J. Stern, editor, *Advances in Cryptology – EUROCRYPT'99*, volume 1592 of *Lecture Notes in Computer Science*, pages 402–414. Springer, May 1999.
9. S. Chatterjee, A. Menezes, and P. Sarkar. Another look at tightness. *Cryptology ePrint Archive*, Report 2011/442, 2011. <http://eprint.iacr.org/>.
10. D. Coppersmith. Finding a small root of a univariate modular equation. In U. M. Maurer, editor, *Advances in Cryptology – EUROCRYPT'96*, volume 1070 of *Lecture Notes in Computer Science*, pages 155–165. Springer, May 1996.
11. J.-S. Coron. On the exact security of full domain hash. In M. Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235. Springer, Aug. 2000.
12. J.-S. Coron. Optimal security proofs for pss and other signature schemes. *Cryptology ePrint Archive*, Report 2001/062, 2001. <http://eprint.iacr.org/>.
13. J.-S. Coron. Optimal security proofs for PSS and other signature schemes. In L. R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 272–287. Springer, Apr. / May 2002.
14. C. Gentry, P. D. Mackenzie, and Z. Ramzan. Password authenticated key exchange using hidden smooth subgroups. In V. Atluri, C. Meadows, and A. Juels, editors, *ACM CCS 05: 12th Conference on Computer and Communications Security*, pages 299–309. ACM Press, Nov. 2005.
15. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In R. E. Ladner and C. Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206. ACM Press, May 2008.

16. C. Gentry and Z. Ramzan. Single-database private information retrieval with constant communication rate. In L. Caires, G. F. Italiano, L. Monteiro, C. Palamidessi, and M. Yung, editors, *ICALP 2005: 32nd International Colloquium on Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 803–815. Springer, July 2005.
17. E.-J. Goh, S. Jarecki, J. Katz, and N. Wang. Efficient signature schemes with tight reductions to the Diffie-Hellman problems. *Journal of Cryptology*, 20(4):493–514, Oct. 2007.
18. B. Hemenway and R. Ostrovsky. Public-key locally-decodable codes. In D. Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 126–143. Springer, Aug. 2008.
19. IEEE P1363a Committee. IEEE P1363a / D9 — standard specifications for public key cryptography: Additional techniques. <http://grouper.ieee.org/groups/1363/index.html/>, June 2001. Draft Version 9.
20. E. Kiltz, A. O’Neill, and A. Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In T. Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 295–313. Springer, Aug. 2010.
21. N. Kobitz and A. J. Menezes. Another look at “provable security”. *Journal of Cryptology*, 20(1):3–37, Jan. 2007.
22. A. Lysyanskaya, S. Micali, L. Reyzin, and H. Shacham. Sequential aggregate signatures from trapdoor permutations. In C. Cachin and J. Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 74–90. Springer, May 2004.
23. S. Micali. Computationally sound proofs. *SIAM J. Comput.*, 30:1253–1298, October 2000.
24. P. Paillier and J. L. Villar. Trading one-wayness against chosen-ciphertext security in factoring-based encryption. In X. Lai and K. Chen, editors, *Advances in Cryptology – ASIACRYPT 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 252–266. Springer, Dec. 2006.
25. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In R. E. Ladner and C. Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 187–196. ACM Press, May 2008.
26. PKCS #1: RSA cryptography standard. RSA Data Security, Inc., Sept. 1998. Version 2.0.
27. C. Schridde and B. Freisleben. On the validity of the phi-hiding assumption in cryptographic protocols. In J. Pieprzyk, editor, *Advances in Cryptology – ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 344–354. Springer, Dec. 2008.
28. N. Smart. Ecrypt ii yearly report on algorithms and key sizes (2009-2010). *Framework*, page 116, March 2010.