

Optimal Security Proofs for Full Domain Hash, Revisited

Saqib A. Kakvi¹ Eike Kiltz¹

¹ Faculty of Mathematics, Horst Görtz Institute for IT Security, Ruhr University Bochum
Universitätsstraße 150, 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 (with message recovery).

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 [BR96]. It is extensively used in a wide variety of applications, and serves as the basis of several existing standards such as PKCS #1 [PKC98]. It has been demonstrated by means of a security reduction that, in the random oracle model [BR93], 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 [BR96] 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 [Cor00] (which improves on earlier results by Bellare and Rogaway [BR96]) 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 [BR96], 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 [Sma10], 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 [CMS11] 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 [Cor01, Cor02] 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 [Cor02]. As acknowledged by the author of [Cor02], 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 [BY96, LMRS04] 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 no longer applies 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 [PW08]). Our main result (Theorem 3.5) shows that there is a *tight* security reduction from the lossiness of the TDP to breaking security of FDH, in the random oracle model. Our results also extend to the probabilistic signature scheme (PSS) (with message recovery) [BR96]. We obtain a tight reduction for RSA-PSS with arbitrary (possibly zero) size random seed from the assumption that the TDP is lossy.

APPLICATIONS TO RSA-FDH AND RSA-PSS. Recently, Kiltz et al. [KOS10] showed that the RSA trapdoor permutation is lossy under the Φ -Hiding Assumption. The Φ -Hiding Assumption was introduced by Cachin, Micali, and Stadler in 1999 [CMS99] 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 gives 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 [CMS99] and has ever since been used in a number of cryptographic applications (e.g., [KOS10, Cac99, GMR05, GR05, Mic00, HO08]). It has been cryptanalyzed (e.g. [CMS99, Cac99, SF08]) 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.

We also obtain a tight reduction for RSA-PSS (with message recovery) from the Φ -Hiding Assumption for arbitrary (possibly zero) size random seed. Assuming again that the Φ -Hiding Assumption is as hard as factoring we get a signature scheme (with message recovery) with only 160 bits of overhead for 80 bits security.

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 lose more than a factor of q_s , the number of signature queries.

THE GAP IN THE PROOF. During the proof of [Cor01, Th. 5] it is silently assumed that the public-key pk generated by reduction \mathcal{R} is a *real public-key*, honestly generated by the key-generation algorithm of FDH, i.e., it contains f which described a permutation.² However, that does not necessarily hold, and 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 $e < N^{1/4-\epsilon}$ [CMS99]. 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., [BR96, GJKW07, Ber08, GPV08]).

1.4 The Probabilistic Signature Scheme PSS

Our observations can also be applied to the probabilistic signature scheme (PSS) [BR96] which is contained in IEEE P1363a [IEE01], ISO/IEC 9796-2, and PKCS#1 v2.1 [PKC98]. Improving an earlier result by Bellare and Rogaway [BR96] Coron proved that, if at least $\log_2(q_s)$ bits of random salt is used in PSS,

¹A 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.

²Such restricted reductions were called *key-preserving reductions* in [PV06].

³We 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.

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

then there is a tight security reduction from the one-wayness of the TDP [Cor01, Cor02]. 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. We show a tight security proof from lossiness to the security of PSS, with random salt of arbitrary (possibly zero) length.

Our results also apply to PSS with message recovery (PSS-R), where the signature encodes (parts of) the message. We show that PSS-R with arbitrary-length salt is tightly secure assuming lossiness of the TDP. Concretely, our security reduction shows that PSS-R with zero-length salt has an overhead (signature length minus message length) of only $2k$ bits, where k is the security parameter. Interestingly, this matches the overhead of BLS short signatures [BLS01] over bilinear maps.

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. Kobiltz and Menezes [KM07, 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 [BNPS03]): 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. [GJKW07] 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 [Ber08] shows a tight security reduction for (a certain randomized variant of) Rabin-Williams signature scheme from the factoring assumption. More generally, Gentry et al. [GPV08] introduce the concept of pre-image 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, [KOS10] 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 $x \leftarrow_{\S} A(a_1, \dots, a_n)$ 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 [BR06]. There is an **Initialize** procedure and a **Finalize** procedure, as well a procedure for each separate oracle. Executing a game G with and 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 .

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.

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

Figure 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 2.1 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 $f^{-1}(y) = \text{Invert}(td, y)$.

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* [BY96] if the fact that it is a permutation can be verified in polynomial time given pub .

Definition 2.2 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 [PW08], is a stronger requirement than one-wayness.

Definition 2.3 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

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

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

Figure 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 $\text{Dom}_{pub'}$.

⁵We deviate in two ways from the original definition of lossy trapdoor functions Peikert and Waters [PW08]. 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 $\ell = \log_2(l)$.

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 $\text{pub} = (N, e)$ and $\text{td} = d$, where $N = pq$ is the product of two $k/2$ -bit primes, $\gcd(e, \varphi(N)) = 1$, and $d = e^{-1} \bmod \varphi(N)$. The domain is $\text{Dom}_{\text{pub}} = \mathbb{Z}_N^*$. The evaluation algorithm $\text{RSAEval}(\text{pub}, x)$ returns $f_{\text{pub}}(x) = x^e \bmod N$, the inversion algorithm $\text{RSAInv}(\text{td}, y)$ returns $f_{\text{pub}}^{-1}(y) = y^d \bmod N$. 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 [BR96] in Figure 3.

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

Figure 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 $(q_h + q_s)$ [BR96], which was later improved by Coron to a factor of q_s [Cor00] for the case of the RSA trapdoor permutation.

Theorem 3.1 (Coron [Cor00]) *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 [Cor01, Cor02]. To state a corrected version of Coron's impossibility result, we first recall the following definitions [Cor01].

Definition 3.2 We say a reduction \mathcal{R} $(t_{\mathcal{F}}, t_{\mathcal{R}}, q_h, q_s, \varepsilon_{\mathcal{F}}, \varepsilon_{\mathcal{R}})$ -reduces inverting a trapdoor permutation 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 an inverse with probability at least $\varepsilon_{\mathcal{R}}$, with running time at most $t_{\mathcal{R}}$.

Definition 3.3 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 Section 5.

Theorem 3.4 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 (with non-negative success probability $\varepsilon_{\mathcal{I}}$).

3.4 A Tight Security Proof for TDP-FDH

The impossibility result of Theorem 3.4 only holds for TDP-FDH if TDP is a 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 3.5 Assume $\text{TDP} = (\text{Gen}, \text{Eval}, \text{Invert})$ 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

$$\begin{aligned} \varepsilon &= \left(\frac{2l-1}{l-1}\right) \cdot \varepsilon' \\ t &= t' - (q_h + q_s) \cdot T_{\text{TDP}} \end{aligned}$$

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 $\varepsilon = \Pr[\mathsf{G}_0^{\mathcal{A}} \Rightarrow 1]$. Here we assume wlog that \mathcal{A} always makes a query to $\text{Hash}(m)$ before calling $\text{Sign}(m)$ or $\text{Finalize}(m, \cdot)$.

Lemma 3.6 $\Pr[\mathsf{G}_0^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathsf{G}_1^{\mathcal{A}} \Rightarrow 1]$.

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 $\sigma_m \in \text{Dom}_{\text{pub}}$. 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 $\text{Eval}(\text{pub}, \sigma_m) = y_m$ for all m . Thus our simulation of the signatures is correct. Since TDP is a permutation, the distribution of our hash queries in G_1 is identical to the distribution in G_0 . Thus we have $\Pr[\mathsf{G}_0^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathsf{G}_1^{\mathcal{A}} \Rightarrow 1]$. ■

Lemma 3.7 There exists a distinguisher \mathcal{D}_1 against the lossiness of TDP, which runs in time $t = t_{\mathcal{A}} + (q_h + q_s) \cdot T_{\text{TDP}}$ and that $\Pr[\mathsf{G}_1^{\mathcal{A}} \Rightarrow 1] - \Pr[\mathsf{G}_2^{\mathcal{A}} \Rightarrow 1] = \text{Adv}_{\text{TDP}}^{\text{L}}(\mathcal{D}_1)$.

Proof: From G_1 to 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

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

Table 1: Games for the proof of Theorem 3.5.

described in games $G_1 \& G_2$, for which it requires time $(q_h + q_s) \cdot T_{\text{TDP}}$. 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 $\Pr[\mathcal{L}_0^{\mathcal{D}_1} \Rightarrow 1] = 1 - \Pr[G_1^{\mathcal{A}} \Rightarrow 1]$. Similarly, we have $\Pr[\mathcal{L}_1^{\mathcal{D}_1} \Rightarrow 1] = 1 - \Pr[G_2^{\mathcal{A}} \Rightarrow 1]$. Hence we have $\Pr[G_1^{\mathcal{A}} \Rightarrow 1] - \Pr[G_2^{\mathcal{A}} \Rightarrow 1] = (1 - \Pr[\mathcal{L}_0^{\mathcal{D}_1} \Rightarrow 1]) - (1 - \Pr[\mathcal{L}_1^{\mathcal{D}_1} \Rightarrow 1]) = \Pr[\mathcal{L}_1^{\mathcal{D}_1} \Rightarrow 1] - \Pr[\mathcal{L}_0^{\mathcal{D}_1} \Rightarrow 1] = \text{Adv}_{\text{TDP}}^{\mathcal{L}}(\mathcal{D}_1)$. ■

Lemma 3.8 $\Pr[G_3^{\mathcal{A}} \Rightarrow 1] = (\frac{l-1}{l}) \Pr[G_2^{\mathcal{A}} \Rightarrow 1]$.

Proof: In 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., G_3 outputs 0. σ_{m^*} is independent of \mathcal{A} 's view and is uniformly distributed in the 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 $\Pr[G_3^{\mathcal{A}} \Rightarrow 1] = (1 - \frac{1}{l}) \Pr[G_2^{\mathcal{A}} \Rightarrow 1] = (\frac{l-1}{l}) \Pr[G_2^{\mathcal{A}} \Rightarrow 1]$. ■

Lemma 3.9 *There exists a distinguisher \mathcal{D}_2 against the lossiness of TDP, which runs in time $t = t_{\mathcal{A}} + (q_h + q_s) \cdot T_{\text{TDP}}$ and that $\Pr[G_3^{\mathcal{A}} \Rightarrow 1] - \Pr[G_4^{\mathcal{A}} \Rightarrow 1] = \text{Adv}_{\text{TDP}}^{\mathcal{L}}(\mathcal{D}_2)$.*

Proof: From G_3 to 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 $G_3 \& G_4$, for which it requires time $(q_h + q_s) \cdot T_{\text{TDP}}$. After \mathcal{A} calls **Finalize**, \mathcal{D}_2 returns the output of **Finalize**. We can see that $\Pr[G_4^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathcal{L}_0^{\mathcal{D}_2} \Rightarrow 1]$. Similarly, we have $\Pr[G_3^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathcal{L}_1^{\mathcal{D}_2} \Rightarrow 1]$. Hence we have $\Pr[G_3^{\mathcal{A}} \Rightarrow 1] - \Pr[G_4^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathcal{L}_1^{\mathcal{D}_2} \Rightarrow 1] - \Pr[\mathcal{L}_0^{\mathcal{D}_2} \Rightarrow 1] = \text{Adv}_{\text{TDP}}^{\mathcal{L}}(\mathcal{D}_2)$. ■

Lemma 3.10 $\Pr[G_4^{\mathcal{A}} \Rightarrow 1] = 0$.

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 $\Pr[\mathsf{G}_4^{\mathcal{A}} \Rightarrow 1] = 0$. ■

We combine Lemmas 3.6 to 3.10 to get:

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

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 [KOS10]. We now recall (and extend) some of the results from [KOS10].

First, we recall some definitions from [KOS10]. 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 [CMS99], where we build on a formalization by Kiltz, O'Neil and Smith [KOS10]. 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 [Cop96], see [CMS99] 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 .

procedure Initialize	Game P_0	procedure Initialize	Game P_1
$e \in_r \mathbb{P}_{\alpha k}$		$e \in_r \mathbb{P}_{\alpha k}$	
$(N, p, q) \in_r \mathcal{RSA}_k[\gcd(e, \varphi(N)) = 1]$		$(N, p, q) \in_r \mathcal{RSA}_k[p = 1 \bmod e, p \neq 1 \bmod e^2, q \neq 1 \bmod e]$	
return (N, e)		return (N, e)	

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

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:

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

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 $(N, p, q) \in_R \mathcal{RSA}_k[\text{gcd}(e, \varphi(N)) = 1]$. (See [KOS10] for details on the sampling algorithm.)

Lemma 4.1 *If the α - Φ -Hiding Problem is (t, ε) -hard, then the RSA = $(\text{RSAGen}, \text{RSAEval}, \text{RSAINV})$ defines a regular $(2^\alpha, t, \varepsilon)$ -lossy trapdoor permutation.*

Proof: If (N, e) is sampled using RSAGen , then $\text{gcd}(e, \varphi(N) = 1)$ and (N, e) defines a permutation $\text{RSA}(x) = x^e \bmod N$ 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 $(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \bmod e, p \neq 1 \bmod e^2, q \neq 1 \bmod e]$. 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 [CMS99] and [KOS10, 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 $\mathcal{RSA}_k[p = 2 \bmod 3, q = 2 \bmod 3]$, which is equivalent to $\mathcal{RSA}_k[\text{gcd}(3, \varphi(N)) = 1]$. We note that $N \bmod 3$ is always 1. This means that for the lossy case, we must also ensure the $N \bmod 3 = 1$, 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 $\mathcal{RSA}_k[p = q = 1 \bmod 3, p \neq 1 \bmod 9, q \neq 1 \bmod 9]$.

procedure Initialize Game $3F_0$ $(N, p, q) \in_R \mathcal{RSA}_k[\text{gcd}(3, \varphi(N)) = 1]$ return $(N, e = 3)$	procedure Initialize Game $3F_1$ $(N, p, q) \in_R \mathcal{RSA}_k[p = q = 1 \bmod 3, p \neq 1 \bmod 9, q \neq 1 \bmod 9]$ 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

$$\text{Adv}^{\text{F}\Phi\text{H}}(\mathcal{D}) = \Pr[3F_1^{\mathcal{D}} \Rightarrow 1] - \Pr[3F_0^{\mathcal{D}} \Rightarrow 1].$$

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 4.2 *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 $(9, t, \varepsilon)$ -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 $(N, p, q) \in \mathcal{RSA}_k[p = q = 1 \bmod 3, p \neq 1 \bmod 9, q \neq 1 \bmod 9]$ then $9 \mid \varphi(N)$ 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 $N \bmod e$ 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 $\mathcal{RSA}_k[p = 1 \bmod e, p \neq 1 \bmod e^2, q \neq 1 \bmod e]$. Consider a distinguisher \mathcal{D} which plays one of the games in Table

procedure Initialize Game F_0 $(N, p, q) \in_R \mathcal{RSA}_k[\text{gcd}(e, \varphi(N)) = 1]$ return (N, e)	procedure Initialize Game F_1 $(N, p, q) \in_R \mathcal{RSA}_k[p = 1 \bmod e, p \neq 1 \bmod e^2, q \neq 1 \bmod e]$ return (N, e)
---	---

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

4. The advantage of \mathcal{D} is defined as

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

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 4.3 *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 $(N, p, q) \in \text{RSA}_k[\text{gcd}(e, \varphi(N)) = 1]$ then (N, e) clearly defines a permutation. If $(N, p, q) \in \text{RSA}_k[p = 1 \pmod{e}, p \neq 1 \pmod{e^2}, q \neq 1 \pmod{e}]$ then $e \mid \varphi(N)$ and hence the RSA function is e -to-1 on the domain $\text{Dom}_{pub} = \mathbb{Z}_N^*$. \blacksquare

5 Proof of Coron's Impossibility Results (Corrected)

5.1 Proof of Theorem 3.4

Proof: Assume \mathcal{R} is a reduction that $(t_{\mathcal{F}}, t_{\mathcal{R}}, q_h, q_s, \varepsilon_{\mathcal{R}}, \varepsilon_{\mathcal{F}})$ -reduces inverting TDP to breaking TDP-FDH. We will now build an adversary \mathcal{I} against the one-wayness of TDP. \mathcal{I} receives pub and $y = \text{Eval}(pub, x) \in \text{Dom}_{pub}$ for unknown $x \in_R \text{Dom}_{pub}$. The goal of \mathcal{I} is to compute x .

1. Adversary \mathcal{I} runs the reduction \mathcal{R} providing it with (pub, y) and in reply receives a public key $pk = pub'$ for TDP-FDH. Note that pub' provided by \mathcal{R} may be different from pub . \mathcal{I} verifies that pub' defines a permutation on $\text{Dom}_{pub'}$ by running $\text{Certify}(1^k, pub')$. If not, then \mathcal{I} outputs \perp to \mathcal{R} . Otherwise, \mathcal{I} continues by simulating a forger \mathcal{F}_{sim} for \mathcal{R} . (We note that it may be the case that $pub' \neq pub$, in particular we do not know if pub' was correctly generated but this does not matter for what follows.)
2. Adversary \mathcal{I} picks q_h messages m_1, \dots, m_{q_h} from the message space $\text{Dom}_{pub'}$ (e.g., at random or the lexicographically first) and makes hash queries to \mathcal{R} on the messages from universe $\mathcal{U} = (m_1, \dots, m_{q_h})$.
3. \mathcal{I} picks $i \in_R \{1, \dots, q_s\}$, $\beta \in_R \{1, \dots, q_h\}$ and $(\alpha_1, \dots, \alpha_{q_s}) \in_R (\{1, \dots, q_h\} \setminus \{\beta\})^{q_s}$. This defines the following two sequences of integers:

$$\alpha^{\text{rw}} = (\alpha_1, \dots, \alpha_{q_s}), \quad \alpha^{\text{first}} = (\alpha_1, \dots, \alpha_{i-1}, \beta).$$

4. Adversary \mathcal{I} queries the signing oracle on the messages indexed by α^{first} (i.e., $(m_{\alpha_1}, \dots, m_{\alpha_{i-1}}, m_{\beta})$) and receives the response $\mathcal{S}^{\text{first}} = (\sigma_{\alpha_1}, \dots, \sigma_{\alpha_{i-1}}, \sigma^*)$, if the reduction does not abort.
5. Adversary \mathcal{R} is then rewound back to after it answered all the hash queries.⁶ Now \mathcal{I} queries the signing oracle on the messages indexed by α^{rw} (i.e., $(m_{\alpha_1}, \dots, m_{\alpha_{q_s}})$) and receives $\mathcal{S}^{\text{rw}} = (\sigma_1, \dots, \sigma_{q_s})$, if the reduction does not abort.
6. Adversary \mathcal{I} then tosses a biased coin τ with probability $\varepsilon_{\mathcal{F}}$ of returning 1, and if $\tau = 0$, then \mathcal{I} sends \perp to \mathcal{R} . If $\tau = 1$, then \mathcal{I} submits $(m^* = m_{\beta}, \sigma^*)$ as a forgery.
7. Because the reduction was rewound, this constitutes a valid forgery, as m^* was not queried to the signature oracle and σ^* is indeed a valid signature on m^* . \mathcal{R} will then return x (with probability $\varepsilon_{\mathcal{R}}$) which \mathcal{I} submits as its solution to one-way experiment.

We now analyze \mathcal{I} 's success probability in breaking one-wayness of TDP. To this end we define \mathcal{Q} as the set of all sequences of indices such that the corresponding signature queries are correctly answered by \mathcal{R} , after the hash queries (in time less than $t_{\mathcal{R}}$). If a sequence of signature queries is correctly answered by

⁶More formally, \mathcal{R} is run again with the same random tape, providing it with exactly the same queries until it has answered the hash queries. Hence upto this point \mathcal{R} behaves exactly the same as in its first execution.

\mathcal{R} , then also the same sequence of signature queries without the last message is correctly answered by \mathcal{R} , so for any sequence $(\alpha_1, \dots, \alpha_j) \in \mathcal{Q}$, we have $(\alpha_1, \dots, \alpha_{j-1}) \in \mathcal{Q}$.

Consider (as a thought-experiment) a (possibly non-efficient) real forger \mathcal{F}_{real} who inputs arbitrary pub' such that $\text{Certify}(pub', 1^k) = 1$, makes hash queries to the messages from \mathcal{U} (receiving answers \mathcal{H}), signature queries to the messages from sequence α^{rw} (receiving answers \mathcal{S}^{rw}) and outputs a valid forgery σ^* on the message $m^* = m_\beta$ with probability $\varepsilon_{\mathcal{F}_{real}}$. This is a valid forger so by the assertion of the theorem reduction \mathcal{R} (interacting with \mathcal{F}_{real}) outputs $x = \text{Invert}(pub, y)$ with probability at least $\varepsilon_{\mathcal{R}}$. After the rewind, \mathcal{R} (interacting with \mathcal{I}) sees exactly the same transcript as he would interact with \mathcal{F}_{real} , except if $\alpha^{\text{first}} \notin \mathcal{Q}$ (\mathcal{R} does not answer all the signature queries before the rewind) and $\alpha^{rw} \in \mathcal{Q}$ (\mathcal{R} answers all the signature queries after the rewind). In that case the forger \mathcal{F}_{real} would output a valid forgery (with probability $\varepsilon_{\mathcal{F}_{real}}$) but our simulated forger does not. Here we are relying on the fact that $\text{Eval}(pub', \cdot)$ is a permutation and hence the forgery σ^* on $m^* = m_\beta$ output by \mathcal{F}_{real} is the same as the simulated one. (Otherwise it could be the case that the simulated forgery is useless for the reduction.)

Let $\mathcal{R}^{\mathcal{F}_{real}}$ denote the execution of \mathcal{R} with the above real forger and let $\mathcal{R}^{\mathcal{F}_{sim}}$ denote the execution of \mathcal{R} with the forger \mathcal{F}_{sim} simulated by \mathcal{I} . The two games $\mathcal{R}^{\mathcal{F}_{real}}$ and $\mathcal{R}^{\mathcal{F}_{sim}}$ are identical until $\alpha^{rw} \in \mathcal{Q}$ and $\alpha^{\text{first}} \notin \mathcal{Q}$, and $\tau = 1$, where all three ‘‘bad events’’ are defined in the game involving inverter \mathcal{I} . By the above we get

$$|\Pr[\mathcal{R}^{\mathcal{F}_{sim}}(pub, y) = x] - \Pr[\mathcal{R}^{\mathcal{F}_{real}}(pub, y) = x]| \leq \varepsilon_{\mathcal{F}} \cdot \Pr[\alpha^{rw} \in \mathcal{Q} \wedge \alpha^{\text{first}} \notin \mathcal{Q}].$$

We need the following combinatorial lemma due to Coron [Cor01, Appendix D].

Lemma 5.1 *Let \mathcal{Q} be a set of sequences of at most q_s integers in $\{1, \dots, q_h\}$, such that for any sequence $(\alpha_1, \dots, \alpha_j) \in \mathcal{Q}$, we have $(\alpha_1, \dots, \alpha_{j-1}) \in \mathcal{Q}$. Then:*

$$\Pr_{\substack{i \in_R \{1, \dots, q_s\} \\ (\alpha_1, \dots, \alpha_{q_s}, \beta) \in_R \{1, \dots, q_h\}^{q_s+1}}} [(\alpha_1, \dots, \alpha_{q_s}) \in \mathcal{Q} \wedge (\alpha_1, \dots, \alpha_{i-1}, \beta) \notin \mathcal{Q}] \leq \frac{\exp(-1)}{q_s}.$$

By Lemma 5.1 we have

$$\Pr[\alpha^{rw} \in \mathcal{Q} \wedge \alpha^{\text{first}} \notin \mathcal{Q}] \leq \frac{\exp(-1)}{q_s} \left(1 - \frac{q_s}{q_h}\right)^{-1}$$

(We note that the $1 - q_s/q_h$ term is due to the fact that α_i are drawn from $\{1, \dots, q_h\} \setminus \{\beta\}$.) Overall, we obtain for the success probability $\varepsilon_{\mathcal{I}}$ of \mathcal{I} :

$$\begin{aligned} \varepsilon_{\mathcal{I}} &= \Pr[\mathcal{R}^{\mathcal{F}_{sim}}(pub, y) = x] \\ &\geq \Pr[\mathcal{R}^{\mathcal{F}_{real}}(pub, y) = x] - \varepsilon_{\mathcal{F}} \cdot \frac{\exp(-1)}{q_s} \left(1 - \frac{q_s}{q_h}\right)^{-1} \\ &= \varepsilon_{\mathcal{R}} - \varepsilon_{\mathcal{F}} \cdot \frac{\exp(-1)}{q_s} \left(1 - \frac{q_s}{q_h}\right)^{-1}. \end{aligned}$$

The time bound comes from running the reduction once, plus the rewind and the additional time for at most q_s additional signing queries and the rewind. This is essentially equal to running the reduction twice. \blacksquare

5.2 Impossibility for any certified unique signature scheme

We now extend Theorem 3.4 in two directions, simultaneously. Firstly, we extend it to reductions from any non-interactive hard problem Π to forging any certified unique signature scheme **SIG**. Secondly we allow the reduction to run the forger \mathcal{F} multiple times.

We first define a certified unique signature scheme.

Definition 5.2 A signature scheme SIG is said to be a certified unique signature scheme if it is a unique signature scheme and there exists a polynomial time algorithm Certify, that on input 1^k and a public key pk will output 1 iff pk defines a unique signature scheme.

We will now recall the definition of a non-interactive problem instance generator due to Abe et al. [AGO11]. This definition is very generic and encompasses both search and decisional problems.

Definition 5.3 A non-interactive problem instance generator consists of a triple of PPT algorithms $\Pi = (\text{Gen}, \text{Verify}, \text{U})$ such that:

- $\text{Gen}(1^k)$ outputs an instance I and a witness w .
- $\text{Verify}(I, S, w)$, where S is a candidate solution and outputs 1 or 0 that represents acceptance or rejection, respectively.
- $\text{U}(I)$ outputs a candidate solution S .

U can be seen as a trivial guessing algorithm against which the advantage of an adversary \mathcal{A} is measured. Concretely, we define the advantage of an adversary \mathcal{A} as:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{NIP}}(\mathcal{A}) &= \Pr[(I, w) \leftarrow_{\S} \text{Gen}(1^k), S \leftarrow_{\S} \mathcal{A}(I) : \text{Verify}(I, S, w) \rightarrow 1] \\ &\quad - \Pr[(I, w) \leftarrow_{\S} \text{Gen}(1^k), S \leftarrow_{\S} \text{U}(I) : \text{Verify}(I, S, w) \rightarrow 1]. \end{aligned}$$

We say that Π is (t, ε) -hard if for all PPT adversaries \mathcal{A} running in time at most t have and $\text{Adv}_{\Pi}^{\text{NIP}}(\mathcal{A}) \leq \varepsilon$.

EXAMPLES OF NON-INTERACTIVE PROBLEMS. We now give a the description of some known problems, both search and decisional, in terms of non-interactive problems. Firstly, we have the RSA problem which states that it is hard to invert the RSA trapdoor permutation. An instance of the RSA problem is given as $I = (N, e, y = x^e \pmod N)$, where (N, e) are taken from the output of RSAGen , as defined in section 4.1, $x \in_{\mathcal{R}} \mathbb{Z}_N$, and the secret information is $s = x$. If the RSA problem is (t, ε) -hard, then this defines a (t, ε) -hard non-interactive problem, where the trivial guessing algorithm U simply picks a random element in \mathbb{Z}_N^* . Secondly, we consider the Φ -Hiding Problem, as described in Section 4.1. An instance of this problem is given as $I = (N, e)$ and the secret information is $s = 1$ if $\gcd(e, \varphi(N)) = 1$ and $s = 0$ if $e \mid \varphi(N)$. The Verify algorithm takes as input (I, b, s) and outputs 1 if $b = s$, else 0. Hence, if the Φ -Hiding Problem is (t, ε) -hard, then this defines a (t, ε) -hard problem with U picking a random bit. Note that we can also define the Computational and Decisional Diffie-Hellman problems in terms of a hard non-interactive problems. Furthermore, this definition also covers problems with non-unique solutions, such as computing modular square roots or approximate-SVP. It is due to these types of problem that we require the Verify function. In these cases, we may not be able to store all the possible solutions, but we can verify a solution using the witness.

We are now ready to state our main impossibility result which generalizes Theorem 3.4.

Theorem 5.4 *Suppose SIG is a certified unique signature scheme, where the size of the message space is at least 2^ℓ . Let \mathcal{R} be a reduction that $(t_{\mathcal{F}}, t_{\mathcal{R}}, q_s, \varepsilon_{\mathcal{F}}, \varepsilon_{\mathcal{R}})$ -reduces breaking a non-interactive problem Π to breaking UF-CMA security of SIG, where \mathcal{R} can run the forger at most r times sequentially. Then we can build an adversary \mathcal{A} which $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ -breaks Π with*

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

Proof: The proof of this theorem is very similar to that of Theorem 3.4. For $1 \leq n \leq r$, we say that the reduction \mathcal{R} is in the n^{th} round when it has run the forger $n - 1$ times. We use this reduction to construct an adversary \mathcal{A} against the problem Π . The reduction is initialized by the adversary \mathcal{A} sending it an instance I of the problem Π . We consider the n^{th} round of the reduction below:

1. The reduction \mathcal{R} provides the adversary \mathcal{A} with a public key pk_n . Note that this public key may be different to the public keys for the previous rounds. \mathcal{A} verifies that pk_n defines a unique signature scheme by running $\text{SIG.Certify}(1^k, pk_n)$. If not, then \mathcal{A} outputs \perp to \mathcal{R} . Otherwise, \mathcal{A} continues.
2. Adversary \mathcal{A} picks 2^ℓ messages m_1, \dots, m_{2^ℓ} from the message space \mathcal{M} (e.g., at random or the lexicographically first) and defines the universe $\mathcal{U}_n = (m_{n,1}, \dots, m_{n,2^\ell})$.
3. \mathcal{A} picks $i_n \in_R \{1, \dots, q_s\}$, $\beta_n \in_R \{1, \dots, 2^\ell\}$ and $(\alpha_{n,1}, \dots, \alpha_{n,q_s}) \in_R (\{1, \dots, 2^\ell\} \setminus \{\beta_n\})^{q_s}$. This defines the following two sequences of integers:

$$\alpha_n^{\text{rw}} = (\alpha_{n,1}, \dots, \alpha_{n,q_s}), \quad \alpha_n^{\text{first}} = (\alpha_{n,1}, \dots, \alpha_{n,i-1}, \beta_n).$$

4. Adversary \mathcal{A} then queries the signing oracle on the messages indexed by α_n^{first} (i.e., $(m_{\alpha_{n,1}}, \dots, m_{\alpha_{n,i-1}}, m_{\beta_n})$) and receives the response $\mathcal{S}_n^{\text{first}}$ which are the signatures, with the last σ^* being the signature on the message indexed by β_n , if the reduction does not abort.
5. Adversary \mathcal{R} is then rewound back to before it answered the signatures queries. Now \mathcal{A} queries the signing oracle on the messages indexed by α_n^{rw} (i.e., $(m_{\alpha_{n,1}}, \dots, m_{\alpha_{n,q_s}})$) and receives \mathcal{S}^{rw} which contains the signatures of the messages indexed by α_n^{rw} , if the reduction does not abort.
6. Adversary \mathcal{A} then tosses a biased coin τ_n with probability $\varepsilon_{\mathcal{F}}$ of returning 1, and If $\tau_n = 0$, then \mathcal{I} sends \perp to \mathcal{R} . If $\tau = 1$, then \mathcal{I} submits $(m_n^* = m_{\beta_n}, \sigma_n^*)$ as a forgery.
7. Because the reduction was rewound, this constitutes a valid forgery, as m_n^* was not queried to the signature oracle and σ_n^* is indeed a valid signature on m_n^* .

This process is repeated for each round, with $n = 1$ to r . We note that in each round, we need to draw a new set of messages for the universe \mathcal{U}_n , as the message space may have changed, due to the public key being different. After at most r rounds, the reduction \mathcal{R} outputs a solution \widehat{S} with probability $\varepsilon_{\mathcal{R}}$, which \mathcal{A} will use as its solution to I .

We now analyze the success probability of \mathcal{A} in breaking the search problem Π . For the n^{th} round, we define \mathcal{Q}_n as the set of sequences of indices such that the corresponding signature queries are correctly answered by \mathcal{R} in the n^{th} round. We can see that if a sequence is correctly answered by \mathcal{R} in the n^{th} round, then the same sequence without the last query is also correctly answered by \mathcal{R} in the n^{th} round. That is to say, for any sequence $(\alpha_{n,1}, \dots, \alpha_{n,j}) \in \mathcal{Q}_n$, we have $(\alpha_{n,1}, \dots, \alpha_{n,j-1}) \in \mathcal{Q}_n$.

Consider now a real forger $\mathcal{F}_{\text{real}}$, which is run r times and in the n^{th} round it receives with public keys pk_n and makes signature queries corresponding to α_n , receives $\mathcal{S}_n^{\text{rw}}$ as an answer, and then submits $(m_{\beta_n}, \sigma_n^*)$ as forgery, with probability $\varepsilon_{\mathcal{F}_{\text{real}}}$. By definition, after at most the r^{th} round, the reduction will output a candidate solution S with probability $\varepsilon_{\mathcal{R}}$. After each n^{th} rewind, the reduction \mathcal{R} sees the same transcript, in the n^{th} round, when interacting with \mathcal{F}_{sim} as it does when it interacts with $\mathcal{F}_{\text{real}}$, unless $\alpha_n^{\text{first}} \notin \mathcal{Q}_n$ and $\alpha_n^{\text{rw}} \in \mathcal{Q}_n$. In this case, the real forger $\mathcal{F}_{\text{real}}$ will output a forgery with probability $\varepsilon_{\mathcal{F}_{\text{real}}}$, however the forger simulated by \mathcal{A} will be unable to output a forgery. This relies on the fact that pk_n defines a unique signature scheme and hence the forgery output by the forger $\mathcal{F}_{\text{real}}$ would be the same as the simulated one.

We denote the execution of the reduction with the real forger as $\mathcal{R}^{\mathcal{F}_{\text{real}}}$ and the execution of the reduction with the forger simulated by \mathcal{A} as $\mathcal{R}^{\mathcal{F}_{\text{sim}}}$. From above, we get:

$$\Pr[\mathcal{R}^{\mathcal{F}_{\text{real}}}(I) \rightarrow S : \text{Verify}(I, S, w) = 1] - \Pr[\mathcal{R}^{\mathcal{F}_{\text{sim}}}(I) \rightarrow S : \text{Verify}(I, S, w) = 1] \leq \sum_{n=1}^r \varepsilon_{\mathcal{F}} \cdot \Pr[\alpha_n^{\text{rw}} \in \mathcal{Q}_n \wedge \alpha_n^{\text{first}} \notin \mathcal{Q}_n]$$

By Lemma 5.1 we have

$$\Pr[\alpha_n \in \mathcal{Q}_n \wedge \alpha_n^{\text{first}} \notin \mathcal{Q}_n] \leq \frac{\exp(-1)}{q_s} \left(1 - \frac{q_s}{2^\ell}\right)^{-1}$$

Overall, we obtain for the success probability $\varepsilon_{\mathcal{A}}$ of \mathcal{A} :

$$\begin{aligned}
\varepsilon_{\mathcal{A}} &\geq \Pr[\mathcal{R}^{\mathcal{F}_{sim}}(I) \rightarrow S : \text{Verify}(I, S, w) = 1] \\
&\geq \Pr[\mathcal{R}^{\mathcal{F}_{real}}(I) \rightarrow S : \text{Verify}(I, S, w) = 1] - \varepsilon_{\mathcal{F}} \cdot \frac{r \cdot \exp(-1)}{q_s} \left(1 - \frac{q_s}{2^\ell}\right)^{-1} \\
&= \varepsilon_{\mathcal{R}} - \varepsilon_{\mathcal{F}} \cdot \frac{r \cdot \exp(-1)}{q_s} \left(1 - \frac{q_s}{2^\ell}\right)^{-1}.
\end{aligned}$$

Now we consider the case where the reduction can rewind the forger to some state St . This is equivalent to restarting the forger with the same random coins and giving it the same input as before until it reaches the required state St . If the reduction rewinds the forger in the $(n+1)^{th}$ round and sends the same public key to the forger, the forger will make the same signature queries and submit the same forgery as in the previous round. This can be simulated by \mathcal{A} by setting $\mathcal{U}_{n+1} = \mathcal{U}_n, \beta_{n+1} = \beta_n, \alpha_{n+1}^{rw} = \alpha_n^{rw}, \alpha_{n+1}^{first} = \alpha_n^{first}$ and proceeding as before. However, in the case where the reduction sends a new public key, then the forger will make different queries and submit a different forgery. Adversary \mathcal{A} can simulate this by picking new messages and new values for $\beta_{n+1}, \alpha_{n+1}^{rw}, \alpha_{n+1}^{first}$, as it normally does, and proceeding as before. We see that in both cases the transcript the reduction sees in the $(n+1)^{th}$ round when interacting with \mathcal{A} is the same as when interacting with a real forger \mathcal{F}_{real} , except when $\alpha_{n+1}^{rw} \in \mathcal{Q}_{n+1} \wedge \alpha_{n+1}^{first} \notin \mathcal{Q}_{n+1}$. By a similar argument as above, we get the same bound.

We see that the adversary runs each round of the reduction twice, thus giving us our time bound. \blacksquare

Remark 5.5 We need that the forgers are run sequentially, otherwise the proof might not go through. When the reduction is allowed to run the forgers in parallel, it may interleave the executions of the forgers and make the replies to one dependant on the queries of the others. If this is the case, then real forgers would not have a problem producing a forgery. However, the simulation of the forger as described above may not be able to do so. Due to the fact that the reduction may change its responses based on the queries, the simulated forgers may never be able to get a forgery and then submit it.

As an example, consider the following execution of the adversary described above, with only 2 parallel forgers. The adversary sends the instance I to the reduction. The reduction in reply sends a public key pk_1 for the first forger \mathcal{F}_1 . The first forger defines the universe \mathcal{U}_1 and sequences $\alpha_1^{rw}, \alpha_1^{first}$ as defined above. It then proceeds to make signing queries on messages indexed α_1^{first} . Based on these queries, \mathcal{R} computes a public key pk_2 for a second forger \mathcal{F}_2 . The second forger begins by defining its universe \mathcal{U}_2 and sequences $\alpha_2^{rw}, \alpha_2^{first}$. Forger \mathcal{F}_2 then proceeds to make signing queries on the messages indexed by α_2^{first} . Based on this the reduction \mathcal{R} computes the responses S_1^{first} and sends it to forger \mathcal{F}_1 . Having received this, adversary \mathcal{A} rewinds \mathcal{R} to before it answered the signature queries. At this point forger \mathcal{F}_1 will make signature queries to the messages indexed by α_1^{rw} . Using these new queries, the reduction \mathcal{R} will compute a new public key \widehat{pk}_2 and use it to initialize forger \mathcal{F}_2 . Due to the fact that $\widehat{pk}_2 \neq pk_2$, forger \mathcal{F}_2 will have to define a new universe $\widehat{\mathcal{U}}_2$ and new sequences $\widehat{\alpha}_2^{first}, \widehat{\alpha}_2^{rw}$. At this point forger \mathcal{F}_2 will make signature queries on the messages indexed by $\widehat{\alpha}_2^{first}$. Based on these queries, the reduction \mathcal{R} will compute the responses to forger \mathcal{F}_1 . At this point it is not clear how to proceed, as the internal state of the reduction has changed. In particular the set \mathcal{Q}'_1 may now exclude α_1^{rw} , which would mean that we can no longer apply Lemma 5.1. Extending Theorem 5.4 to parallel forgers remains an open problem.

6 The Probabilistic Signature Scheme

6.1 The Scheme

Let $\text{TDP} = (\text{Gen}, \text{Eval}, \text{Invert})$ be a trapdoor permutation. For simplicity we assume that $\text{Dom}_{pub} = \{0, 1\}^k$, for all pub output by Gen . (If that is not the case and we have, e.g., $\text{Dom}_{pub} = \mathbb{Z}_N^*$, the scheme can be adapted accordingly [BR96].) We now recall the Probabilistic Signature Scheme [BR96] (PSS) PSS is parametrized not only by the security parameter, but also by two additional integers k_0 and k_1 . The first parameter k_0 defines the size of randomness in bits and the second defines the domains of the

two hash functions. PSS uses two hash functions $H : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and $G : \{0, 1\}^{k_1} \rightarrow \{0, 1\}^{k-k_1}$. We also require two additional functions G_1 and G_2 . $G_1(\omega)$ returns the first k_0 bits of the output of $G(\omega)$ and $G_2(\omega)$ returns the remaining $(k - k_1) - k_0$ bits of $G(\omega)$. The scheme TDP-PSS $[k_0, k_1]$ is defined in Figure 4.

<pre> procedure KeyGen $(pub, td) \leftarrow_s \text{Gen}(1^k)$ Pick hash functions $H : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$, $G : \{0, 1\}^{k_1} \rightarrow \{0, 1\}^{k-k_1}$ return $(pk = (H, G, pub), sk = td)$ procedure Sign(sk, m) $r \in_R \{0, 1\}^{k_0}$ $\omega \leftarrow H(m r)$ $r^* \leftarrow G_1(\omega) \oplus r$ $y = \omega r^* G_2(\omega)$ return $\sigma = \text{Invert}(td, y)$ procedure Verify(pk, m, σ) $y = \text{Eval}(pub, \sigma)$ parse y as $\omega r^* \gamma$ $r = r^* \oplus G_1(\omega)$ if $H(m r) = \omega \wedge G_2(\omega) = \gamma$ return 1 else return 0 </pre>	TDP-PSS
--	---------

Figure 4: The Trapdoor Permutation Probabilistic Signature Scheme TDP-PSS.

6.2 Classical Security Results of TDP-PSS

The original reduction by Bellare and Rogaway from one-wayness of TDP with a security loss of q_h [BR96] was later improved by Coron to a factor of q_s [Cor02] for the case of the RSA trapdoor permutation.

Theorem 6.1 (Coron [Cor02]) *Assume the trapdoor permutation RSA is (t', ε') -hard to invert. Then for any (q_h, q_s) , RSA-PSS $[k_0, k_1]$ is $(q_h, q_s, t, \varepsilon)$ -UF-CMA secure in the Random Oracle Model, where*

$$\begin{aligned} \varepsilon &= \varepsilon' (1 + 6 \cdot q_s \cdot 2^{-k_0}) + 2 \cdot (q_h + q_s)^2 \cdot 2^{-k_1} \\ t &= t' - (q_h + q_s + 1) \cdot k_1 \cdot \mathcal{O}(k^3). \end{aligned}$$

6.3 A Tight Security Proof for TDP-PSS

We now present a tight reduction which reduces breaking TDP-PSS to the lossiness of the underlying trapdoor permutation TDP.

Theorem 6.2 *Assume $\text{TDP} = (\text{Gen}, \text{Eval}, \text{Invert})$ is a regular (l, t', ε') -lossy trapdoor permutation for $l \geq 2$. Then, for any (q_h, q_s, k_0, k_1) , TDP-PSS $[k_0, k_1]$ is $(t, \varepsilon, q_h, q_s)$ -UF-CMA secure in the Random Oracle Model, where*

$$\begin{aligned} \varepsilon &= \left(\frac{2l-1}{l-1} \right) \cdot \varepsilon' + \frac{(q_h + q_s)^2}{2^{k_1}} \\ t &= t' - q_h \cdot T_{\text{TDP}}, \end{aligned}$$

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

Remark 6.3 We stress that the security reduction of Theorem 6.2 is independent of the size of the randomness k_0 . In particular, this allows to set $k_0 = 0$.

Proof: Let \mathcal{A} be an adversary that runs in time t against TDP-PSS executed in the UF-CMA experiment described in G_0 in Figure 5 with $\varepsilon = \Pr[G_0^{\mathcal{A}} \Rightarrow 1] = \Pr[\text{UF-CMA}^{\mathcal{A}} \Rightarrow 1]$. Here we assume wlog that \mathcal{A} always makes a query to $H(m, \cdot)$ before calling $\text{Sign}(m)$ or $\text{Finalize}(m, \cdot)$.

procedure Initialize $(pub, td) \leftarrow_s \text{Gen}(1^k)$ $(pub, \perp) \leftarrow_s \text{LossyGen}(1^k)$ Return $pk = pub$	Games $G_0 - G_4$ $//G_0, G_1, G_4$ $//G_2, G_3$	procedure $G(\omega)$ if $\mathcal{W}[\omega]$ is defined then fetch $\alpha = \mathcal{W}[\omega]$ else $\alpha \in_R \{0, 1\}^{k-k_1}$ $\mathcal{W}[\omega] := \alpha$ end if return α	Game $G_0 - G_4$
procedure $H(m, r)$ if $\mathcal{H}[m, r]$ then then fetch $\omega_{m,r} = \mathcal{H}[m, r]$ else $\omega_{(m,r)} \in_R \{0, 1\}^{k_1}$ $\mathcal{H}[m, r] := \omega_{(m,r)}$ end if return $\omega_{(m,r)}$	Game G_0	procedure $H(m, r)$ if $\mathcal{H}[m, r]$ is defined then fetch $(\omega_{(m,r)}, \sigma_{(m,r)}) = \mathcal{H}[m, r]$ else $\sigma_{(m,r)} \in_R \text{Dom}_{pub}$ $y_{(m,r)} = \text{Eval}(pub, \sigma_{(m,r)})$ Parse $y_{(m,r)} = \omega_{(m,r)} r_{(m,r)}^* \gamma_{(m,r)}$ if $\mathcal{W}[\omega_{(m,r)}]$ is defined then return \perp $\alpha_{m,r} = (r_{(m,r)}^* \oplus r) \gamma_{m,r}$ $\mathcal{W}[\omega_{(m,r)}] := \alpha_{m,r}$ $\mathcal{H}[m, r] := (\omega_{(m,r)}, \sigma_{(m,r)})$ end if return $\omega_{(m,r)}$	Games $G_1 - G_4$
procedure $\text{Sign}(m)$ $\mathcal{M} \leftarrow \mathcal{M} \cup (m)$ $r \in_R \{0, 1\}^{k_0}$ $\omega = H(m, r)$ $r^* = G_1(\omega) \oplus r$ $y = \omega r^* G_2(\omega)$ return $\sigma_m = \text{Invert}(td, y)$	Game G_0	procedure $\text{Sign}(m)$ $\mathcal{M} \leftarrow \mathcal{M} \cup (m)$ $r \in_R \{0, 1\}^{k_0}$ Call $H(m, r)$ Fetch $(\omega_{m,r}, \sigma_{m,r}) = \mathcal{H}[m, r]$, return $\sigma_{m,r}$	Games $G_1 - G_4$
Procedure $\text{Finalize}(m^*, \sigma^*)$ $y = \text{Eval}(pub, \sigma^*); y = \omega r^* \gamma$ $r = r^* \oplus G_1(\omega); \text{Call } H(m, r)$ Fetch $\mathcal{H}[(m^*, r)]$ if $\sigma_{(m^*, r)} = \sigma^*$ then $\text{BAD} = \text{true}$ return 0 if $(H(m^*, r) = \omega) \wedge (G_2(\omega) = \gamma) \wedge (m^* \notin \mathcal{M})$ return 1 else return 0	Game $G_0 - G_4$ $//G_3, G_4$ $//G_3, G_4$		

Table 5: Games for the proof of Theorem 6.2.

Now consider Game G_1 from Figure 5.

Lemma 6.4 $\Pr[G_1^{\mathcal{A}} \Rightarrow 1] \geq \Pr[G_0^{\mathcal{A}} \Rightarrow 1] - \frac{(q_h + q_s)^2}{2^{k_1}}$.

Proof: In G_0 , we modelled the hash functions as a random oracle. In G_1 we modify the H oracle and the signing oracle. On any (m, r) the H oracle now works by evaluating the permutation on a random element $\sigma_{(m,r)} \in \text{Dom}_{pub}$. The result is then parsed as $\omega_{(m,r)} || r_{(m,r)}^* || \gamma_{(m,r)}$ and we check if $\omega_{(m,r)}$ has been previously queried to the G oracle. If it has, then we abort, otherwise we store the randomness and the values $\omega_{(m,r)}$ and $\sigma_{(m,r)}$. Then we compute the output of $G(\omega_{(m,r)})$ and store it. The signing oracle is then modified to pick a random value r , for which we have already have computed $H(m || r)$, and return the element $\sigma_{(m,r)}$. Note that signing no longer requires the trapdoor td . It can be seen that all our signatures will verify due to the manner in which we compute the responses to queries to H and G . Thus our simulation of the signatures is correct. Since TDP is a permutation, the distribution of our H -Oracle queries in G_1 is the same as in G_0 , except when the reduction aborts, due to a collision in the G -Oracle. The probability of a collision is at most $\frac{(q_s + q_h)^2}{2^{k_1}}$, since the adversary makes at most $q_h + q_s$ queries, implicit or explicit, to the G -oracle, giving at most $(q_h + q_s)^2$ possible collisions from a total of 2^{k_1} possible choices. Thus we have $\Pr[G_1^{\mathcal{A}} \Rightarrow 1] \geq \Pr[G_0^{\mathcal{A}} \Rightarrow 1] - \frac{(q_h + q_s)^2}{2^{k_1}}$. ■

Lemma 6.5 *There exists a distinguisher \mathcal{D}_1 against the lossiness of TDP, which runs in time $t = t_{\mathcal{A}} + (q_h + q_s) \cdot T_{\text{TDP}}$ and that $\Pr[\mathbf{G}_1^{\mathcal{A}} \Rightarrow 1] - \Pr[\mathbf{G}_2^{\mathcal{A}} \Rightarrow 1] = \mathbf{Adv}_{\text{TDP}}^l(\mathcal{D}_1)$.*

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 $\text{Sign}(\cdot), \text{H}(\cdot, \cdot), \text{G}(\cdot)$ as described in games \mathbf{G}_1 & \mathbf{G}_2 , for which it requires time $(q_h + q_s) \cdot T_{\text{TDP}}$. 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 $\Pr[\mathbf{L}_0^{\mathcal{D}_1} \Rightarrow 1] = 1 - \Pr[\mathbf{G}_1^{\mathcal{A}} \Rightarrow 1]$. Similarly, we have $\Pr[\mathbf{L}_1^{\mathcal{D}_1} \Rightarrow 1] = 1 - \Pr[\mathbf{G}_2^{\mathcal{A}} \Rightarrow 1]$. Hence we have $\Pr[\mathbf{G}_1^{\mathcal{A}} \Rightarrow 1] - \Pr[\mathbf{G}_2^{\mathcal{A}} \Rightarrow 1] = (1 - \Pr[\mathbf{L}_0^{\mathcal{D}_1} \Rightarrow 1]) - (1 - \Pr[\mathbf{L}_1^{\mathcal{D}_1} \Rightarrow 1]) = \Pr[\mathbf{L}_1^{\mathcal{D}_1} \Rightarrow 1] - \Pr[\mathbf{L}_0^{\mathcal{D}_1} \Rightarrow 1] = \mathbf{Adv}_{\text{TDP}}^l(\mathcal{D}_1)$. ■

Lemma 6.6 $\Pr[\mathbf{G}_3^{\mathcal{A}} \Rightarrow 1] = \left(\frac{l-1}{l}\right) \Pr[\mathbf{G}_2^{\mathcal{A}} \Rightarrow 1]$.

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

Lemma 6.7 *There exists a distinguisher \mathcal{D}_2 against the lossiness of TDP, which runs in time $t = t_{\mathcal{A}} + (q_h + q_s) \cdot T_{\text{TDP}}$ and that $\Pr[\mathbf{G}_3^{\mathcal{A}} \Rightarrow 1] - \Pr[\mathbf{G}_4^{\mathcal{A}} \Rightarrow 1] = \mathbf{Adv}_{\text{TDP}}^l(\mathcal{D}_2)$.*

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 $(q_h + q_s) \cdot T_{\text{TDP}}$. After \mathcal{A} calls `Finalize`, \mathcal{D}_2 returns the output of `Finalize`. We can see that $\Pr[\mathbf{G}_4^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathbf{L}_0^{\mathcal{D}_2} \Rightarrow 1]$. Similarly, we have $\Pr[\mathbf{G}_3^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathbf{L}_1^{\mathcal{D}_2} \Rightarrow 1]$. Hence we have $\Pr[\mathbf{G}_3^{\mathcal{A}} \Rightarrow 1] - \Pr[\mathbf{G}_4^{\mathcal{A}} \Rightarrow 1] = \Pr[\mathbf{L}_1^{\mathcal{D}_2} \Rightarrow 1] - \Pr[\mathbf{L}_0^{\mathcal{D}_2} \Rightarrow 1] = \mathbf{Adv}_{\text{TDP}}^l(\mathcal{D}_2)$. ■

Lemma 6.8 $\Pr[\mathbf{G}_4^{\mathcal{A}} \Rightarrow 1] = 0$.

Proof: In \mathbf{G}_4 we again use the original `KeyGen` such that `Eval(pub, ·)` 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 $\Pr[\mathbf{G}_4^{\mathcal{A}} \Rightarrow 1] = 0$. ■

We combine Lemmas 6.4 to 6.8 to get:

$$\Pr[\mathbf{G}_0^{\mathcal{A}} \Rightarrow 1] \leq \mathbf{Adv}_{\text{TDP}}^l(\mathcal{D}_1) + \left(\frac{l}{l-1}\right) \mathbf{Adv}_{\text{TDP}}^l(\mathcal{D}_2) + \frac{(q_s + q_h)^2}{2^{k_1}}.$$

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:

$$\varepsilon \leq \frac{2l-1}{l-1} \cdot \varepsilon' + \frac{(q_h + q_s)^2}{2^{k_1}}.$$

This completes the proof. ■

6.4 PSS with message recovery

We now recall the PSS with message recovery (PSS-R) scheme, which is a digital signature scheme with message recovery based on PSS.

Definition 6.9 A digital signature scheme with message recovery is a triple of probabilistic algorithms $\text{SIG-R} = (\text{KeyGen}, \text{Sign}, \text{Recover})$, where:

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 an “enhanced signature” σ .
3. **Recover** is a deterministic algorithm, which on input of a public key and an “enhanced signature” σ outputs either the corresponding message m (accept) or the special symbol \perp (reject).

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

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

We describe TDP-PSS-R in Figure 5 using the same notation as in Section 6. The overhead of signature scheme with recover is defined as the difference in size between the message and the signature. With PSS-R, we can sign and recover messages of size at most $n = k - k_0 - k_1$ bits long with a k bit signature. Thus we have a total overhead of $k_1 + k_0$ bits.

<pre> procedure KeyGen $(pub, td) \leftarrow_s \text{Gen}(1^k)$ Pick hash functions $H : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$, $G : \{0, 1\}^{k_1} \rightarrow \{0, 1\}^{k-k_1}$ return $(pk = (H, G, pub), sk = td)$ procedure Sign(sk, m) $r \in_R \{0, 1\}^{k_0}$ $\omega \leftarrow H(m r)$ $r^* \leftarrow G_1(\omega) \oplus r$ $m^* \leftarrow G_2(\omega) \oplus m$ $y = \omega r^* m^*$ return $\sigma = \text{Invert}(td, y)$ procedure Recover(pk, σ) $y = \text{Eval}(pub, \sigma)$ parse y as $\omega r^* m^*$ $m = m^* \oplus G_2(\omega)$ $r = r^* \oplus G_1(\omega)$ if $H(m r) = \omega$ return m else return \perp </pre>	<p>TDP-PSS-R</p> <p>// $\sigma = f_{pub}^{-1}(y)$</p>
--	--

Figure 5: The Trapdoor Probabilistic Signature Scheme with Message Recovery TDP-PSS-R.

Theorem 6.10 Assume $\text{TDP} = (\text{Gen}, \text{Eval}, \text{Invert})$ is a regular (l, t', ε') -lossy trapdoor permutation for $l \geq 2$. Then, for any (q_h, q_s, k_0, k_1) , $\text{TDP-PSS-R}[k_0, k_1]$ is $(t, \varepsilon, q_h, q_s)$ -UF-CMA secure in the Random Oracle Model, where

$$\varepsilon = \left(\frac{2l-1}{l-1} \right) \cdot \varepsilon' + \frac{(q_h + q_s)^2}{2^{k_1}}$$

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

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

The proof of security for PSS-R is similar to that of PSS and hence we do not present it.

We now discuss the choices of parameters k_0, k_1 . First of we notice that the security of TDP-PSS-R does not depend at all on the size of the randomness k_0 . Hence, we can simply pick $k_0 = 0$. We note that this does not contradict the result of Coron [Cor02] that states if we use $k_0 = 0$, then our reduction must lose a factor of q_s . This is due to fact that this only holds if TDP-PSS is instantiated with a certified trapdoor permutation.

For the choice of k_1 , we need the notion of “work factor” due to Bellare and Ristenpart [BR09], whose definition we now recall.

Definition 6.11 The work factor WF of an adversary \mathcal{A} with a success probability of $\varepsilon_{\mathcal{A}}$ and running time $t_{\mathcal{A}}$ is given by:

$$\text{WF}(\mathcal{A}) = \frac{t_{\mathcal{A}}}{\varepsilon_{\mathcal{A}}}.$$

We say that a scheme achieves κ bit security if for all adversaries \mathcal{A} , we have $\text{WF}(\mathcal{A}) \geq 2^{\kappa+1}$. This condition determines our choice of parameters. Using this and the formulae from Theorem 6.10, we can compute the value for k_1 such that TDP-PSS has κ -bit security. We see that we need:

$$\frac{t'}{\varepsilon'} \frac{l-1}{2l-1} \geq 2^{\kappa} \quad \text{and} \quad \frac{t \cdot 2^{k_1}}{(q_s + q_h)^2} \geq 2^{\kappa}.$$

The first requirement is satisfied by our assumption that we have a regular (l, t', ε') -lossy trapdoor permutation. If we look at the second inequality, we see that we require $k_1 \geq \kappa + 2 \log(q_h + q_s) - \log t$. However, we also have that $t \geq q_s + q_h$, hence giving us $k_1 \geq \kappa + \log(q_h + q_s)$.

We can perform similar calculations for the proofs of Bellare-Rogaway and Coron, but we omit them here. We compute some concrete figures for the overhead imposed by each security proof and present them in Table 6. We take $\kappa = 80, q_h = 2^{80}, q_s = 2^{30}$.

Security Proof	k_0	k_1	Total overhead
Bellare-Rogaway [BR96]	160	160	320
Coron [Cor02]	30	160	190
This work	0	160	160

Table 6: Total overhead using RSA-PSS-R for 80 bit security.

Acknowledgements

We thank Mihir Bellare, Dennis Hofheinz, and Bertram Poettering for valuable comments on an earlier draft. The authors are supported by a Sofja Kovalevskaja Award of the Alexander von Humboldt Foundation, funded by the German Federal Ministry for Education and Research.

References

- [AGO11] Masayuki Abe, Jens Groth, and Miyako Ohkubo. Separating short structure-preserving signatures from non-interactive assumptions. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 628–646. Springer, 2011. (Cited on page 14.)
- [Ber08] Daniel J. Bernstein. Proving tight security for Rabin-Williams signatures. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 70–87. Springer, April 2008. (Cited on page 3, 4.)

- [BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 514–532. Springer, December 2001. (Cited on page 4.)
- [BNPS03] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael 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. (Cited on page 4.)
- [BR93] Mihir Bellare and Phillip 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, November 1993. (Cited on page 1.)
- [BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *Advances in Cryptology – EUROCRYPT’96*, volume 1070 of *Lecture Notes in Computer Science*, pages 399–416. Springer, May 1996. (Cited on page 1, 2, 3, 7, 16, 17, 21.)
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, May / June 2006. (Cited on page 5.)
- [BR09] Mihir Bellare and Thomas Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for Waters’ IBE scheme. In Antoine Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 407–424. Springer, April 2009. (Cited on page 21.)
- [BY96] Mihir Bellare and Moti Yung. Certifying permutations: Noninteractive zero-knowledge based on any trapdoor permutation. *Journal of Cryptology*, 9(3):149–166, 1996. (Cited on page 2, 6.)
- [Cac99] Christian 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, November 1999. (Cited on page 2.)
- [CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In Jacques Stern, editor, *Advances in Cryptology – EUROCRYPT’99*, volume 1592 of *Lecture Notes in Computer Science*, pages 402–414. Springer, May 1999. (Cited on page 2, 3, 10, 11.)
- [CMS11] Sanjit Chatterjee, Alfred Menezes, and Palash Sarkar. Another look at tightness. *Cryptology ePrint Archive*, Report 2011/442, 2011. <http://eprint.iacr.org/>. (Cited on page 1.)
- [Cop96] Don Coppersmith. Finding a small root of a univariate modular equation. In Ueli M. Maurer, editor, *Advances in Cryptology – EUROCRYPT’96*, volume 1070 of *Lecture Notes in Computer Science*, pages 155–165. Springer, May 1996. (Cited on page 10.)
- [Cor00] Jean-Sébastien Coron. On the exact security of full domain hash. In Mihir Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235. Springer, August 2000. (Cited on page 1, 7.)
- [Cor01] Jean-Sébastien Coron. Optimal security proofs for pss and other signature schemes. *Cryptology ePrint Archive*, Report 2001/062, 2001. <http://eprint.iacr.org/>. (Cited on page 1, 3, 4, 7, 13.)
- [Cor02] Jean-Sébastien Coron. Optimal security proofs for PSS and other signature schemes. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 272–287. Springer, April / May 2002. (Cited on page 1, 2, 4, 7, 17, 21.)

- [GJKW07] Eu-Jin Goh, Stanislaw Jarecki, Jonathan Katz, and Nan Wang. Efficient signature schemes with tight reductions to the Diffie-Hellman problems. *Journal of Cryptology*, 20(4):493–514, October 2007. (Cited on page 3, 4.)
- [GMR05] Craig Gentry, Philip D. Mackenzie, and Zulfikar Ramzan. Password authenticated key exchange using hidden smooth subgroups. In Vijayalakshmi Atluri, Catherine Meadows, and Ari Juels, editors, *ACM CCS 05: 12th Conference on Computer and Communications Security*, pages 299–309. ACM Press, November 2005. (Cited on page 2.)
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206. ACM Press, May 2008. (Cited on page 3, 4.)
- [GR05] Craig Gentry and Zulfikar Ramzan. Single-database private information retrieval with constant communication rate. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti 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. (Cited on page 2.)
- [HO08] Brett Hemenway and Rafail Ostrovsky. Public-key locally-decodable codes. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 126–143. Springer, August 2008. (Cited on page 2.)
- [IEE01] 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. (Cited on page 3.)
- [KM07] Neal Koblitz and Alfred J. Menezes. Another look at “provable security”. *Journal of Cryptology*, 20(1):3–37, January 2007. (Cited on page 4.)
- [KOS10] Eike Kiltz, Adam O’Neill, and Adam Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 295–313. Springer, August 2010. (Cited on page 2, 4, 10, 11.)
- [LMRS04] Anna Lysyanskaya, Silvio Micali, Leonid Reyzin, and Hovav Shacham. Sequential aggregate signatures from trapdoor permutations. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 74–90. Springer, May 2004. (Cited on page 2.)
- [Mic00] Silvio Micali. Computationally sound proofs. *SIAM J. Comput.*, 30:1253–1298, October 2000. (Cited on page 2.)
- [PKC98] PKCS #1: RSA cryptography standard. RSA Data Security, Inc., September 1998. Version 2.0. (Cited on page 1, 3.)
- [PV06] Pascal Paillier and Jorge L. Villar. Trading one-wayness against chosen-ciphertext security in factoring-based encryption. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology – ASIACRYPT 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 252–266. Springer, December 2006. (Cited on page 3.)
- [PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 187–196. ACM Press, May 2008. (Cited on page 2, 6.)
- [SF08] Christian Schridde and Bernd Freisleben. On the validity of the phi-hiding assumption in cryptographic protocols. In Josef Pieprzyk, editor, *Advances in Cryptology – ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 344–354. Springer, December 2008. (Cited on page 2.)

[Sma10] Nigel Smart. Ecrypt ii yearly report on algorithms and key sizes (2009-2010). *Framework*, page 116, March 2010. (Cited on page 1.)