

On the Lossiness of the Rabin Trapdoor Function

Yannick Seurin

ANSSI, Paris, France
yannick.seurin@m4x.org

Abstract. Lossy trapdoor functions, introduced by Peikert and Waters (STOC '08), are functions that can be generated in two indistinguishable ways: either the function is injective, and there is a trapdoor to invert it, or the function is lossy, meaning that the size of its range is strictly smaller than the size of its domain. Kakvi and Kiltz (EUROCRYPT 2012) proved that the Full Domain Hash signature scheme based on a lossy trapdoor function has a *tight* security reduction from the lossiness of the trapdoor function. Since Kiltz, O'Neill, and Smith (CRYPTO 2010) showed that the RSA trapdoor function is lossy under the Φ -Hiding assumption of Cachin, Micali, and Stadler (EUROCRYPT '99), this implies that the RSA Full Domain Hash signature scheme has a *tight* security reduction from the Φ -Hiding assumption (for public exponents $e < N^{1/4}$). In this work, we consider the Rabin trapdoor function, *i.e.* modular squaring over \mathbb{Z}_N^* . We show that when adequately restricting its domain (either to the set \mathbb{QR}_N of quadratic residues, or to $(\mathbb{J}_N)^+$, the set of positive integers $1 \leq x \leq (N-1)/2$ with Jacobi symbol $+1$) the Rabin trapdoor function is lossy, the injective mode corresponding to Blum integers $N = pq$ with $p, q \equiv 3 \pmod{4}$, and the lossy mode corresponding to what we call pseudo-Blum integers $N = pq$ with $p, q \equiv 1 \pmod{4}$. This lossiness result holds under a natural extension of the Φ -Hiding assumption to the case $e = 2$ that we call the $2\text{-}\Phi/4$ -Hiding assumption. We then use this result to prove that deterministic variants of Rabin-Williams Full Domain Hash signatures have a tight reduction from the $2\text{-}\Phi/4$ -Hiding assumption. We also show that these schemes are unlikely to have a tight reduction from the factorization problem by extending a previous “meta-reduction” result by Coron (EUROCRYPT 2002), later corrected by Kakvi and Kiltz (EUROCRYPT 2012). These two results therefore answer one of the main questions left open by Bernstein (EUROCRYPT 2008) in his work on Rabin-Williams signatures.

1 Introduction

1.1 Background

Lossy Trapdoor Functions. Lossy Trapdoor Functions (LTF) were introduced by Peikert and Waters [28] and have since then found a wide range of applications in cryptography such as deterministic public-key encryption [8], hedged

public-key encryption [2], and security against selective opening attacks [3, 14] to name a few. Informally, an LTF consists of two families of functions: functions in the first family are injective (and efficiently invertible using some trapdoor), while functions in the second family are non-injective and hence lose information on their input. The key requirement for an LTF is that functions sampled from the first and the second family be computationally indistinguishable. Many constructions of LTF are known from various hardness assumptions such as DDH, LWE, etc. [28]. In particular, Kiltz, O’Neill, and Smith showed [24] that the RSA trapdoor function $f : x \mapsto x^e \bmod N$, where $N = pq$ is an RSA modulus, is lossy under the Φ -Hiding assumption, introduced by Cachin, Micali, and Stadler [9]. When e is coprime with $\phi(N)$ ($\phi(\cdot)$ is Euler’s totient function), f is injective on the domain \mathbb{Z}_N^* , while when e divides $\phi(N)$ (but e^2 does not), f is e -to-1 on \mathbb{Z}_N^* . The Φ -Hiding assumption states that given (N, e) where $e < N^{1/4}$, it is hard to tell whether $\gcd(e, \phi(N)) = 1$ or $e | \phi(N)$, which corresponds to respectively the injective and lossy modes of the RSA function.

Full Domain Hash Signatures. Full Domain Hash (FDH) signatures [4] are a class of signature schemes which can be based on any trapdoor function f : the signature of a message m is computed as $\sigma = f^{-1}(H(m))$, where H is some hash function (the secret signature key is the trapdoor enabling to invert f). For a long time, the only known security result for FDH signatures, due to Coron [11] (improving on a previous result [4]), had been a *non-tight* reduction from the problem of inverting the trapdoor function, losing a factor q_s (the maximal number of signature queries made by the forger). Recently, Kakvi and Kiltz [22] showed that the FDH signature scheme, when based on a trapdoor function which is lossy, has a *tight* reduction from the problem of distinguishing the injective from the lossy mode of the LTF. In particular, this applies to RSA-FDH signatures with public exponents $e < N^{1/4}$, which hence have a tight security reduction from the Φ -Hiding problem.¹ Moreover, in the same paper, Kakvi and Kiltz corrected a previous “meta-reduction” result due to Coron [12] stating that the security reduction of [11] losing a factor q_s is essentially optimal. More precisely, they showed that when the trapdoor function is *certified* (meaning that there is an efficient algorithm distinguishing injective from non-injective members of the function family), any security reduction from inverting the trapdoor function to breaking FDH signatures must lose a factor q_s (unless inverting the trapdoor function is easy). This applies in particular to RSA-FDH signatures with public exponents $e > N^{1/4}$ since RSA is certified for these parameters [23].

1.2 Contributions of this Work

Lossiness of the Rabin Trapdoor Function. We show that the Rabin trapdoor function, *i.e.* modular squaring, is lossy (with exactly one or two bits of lossiness) when adequately restricting its domain. Since any quadratic residue

¹ Tight security reductions are important for adequately setting security parameters, see the discussion of this point in [22].

modulo an RSA modulus $N = pq$ has exactly four square roots, it is not immediately obvious how to render this function injective. It is well known that when N is a so-called Blum integer, *i.e.* $p, q \equiv 3 \pmod{4}$, any quadratic residue has a unique square root which is also a quadratic residue, named its *principal* square root. Hence, in this case, modular squaring defines a permutation over the set of quadratic residues \mathbb{QR}_N . One potential problem with this definition of the injective mode is that the domain of the permutation is (presumably) not efficiently recognizable (this is exactly the Quadratic Residuosity assumption). A different way to restrict the domain of modular squaring is to consider the set $(\mathbb{J}_N)^+$ of integers $1 \leq x \leq (N-1)/2$ with Jacobi symbol $+1$ (which is efficiently recognizable). We show that when restricting its domain to either \mathbb{QR}_N or $(\mathbb{J}_N)^+$ to make it injective, modular squaring becomes an LTF. The lossy mode corresponds to integers $N = pq$ such that $p, q \equiv 1 \pmod{4}$, that we call pseudo-Blum integers. It can be shown that in that case, modular squaring becomes 4-to-1 over \mathbb{QR}_N and 2-to-1 over $(\mathbb{J}_N)^+$. Indistinguishability of the injective and lossy modes is then exactly the problem of distinguishing Blum from pseudo-Blum integers, which is equivalent to tell whether 2 divides $\phi(N)/4$ or not. This can be seen as the extension of the traditional Φ -Hiding assumption to exponent $e = 2$, so that we call this problem the $2\text{-}\Phi/4$ -Hiding problem. Details can be found in Sections 2 and 3.

Application to Rabin-Williams Signatures. We apply our finding to the security of deterministic Rabin-Williams Full Domain Hash signatures. The Rabin signature scheme [29] is one of the oldest provably secure digital signature scheme. Its security relies on the difficulty of computing modular square roots, which is equivalent to factoring integers. Given an RSA modulus $N = pq$, the general principle of Rabin signatures is to first map the message $m \in \{0, 1\}^*$ to a quadratic residue h modulo N using some hash function H , and then return a square root s of h . Since only $1/4$ of integers in \mathbb{Z}_N^* are quadratic residues, directly using $h = H(m) \pmod{N}$ will fail for roughly 3 out of 4 messages. This can be coped with using a randomized padding. The simplest one, Probabilistic Full Domain Hash with ℓ -bit salts (ℓ -PFDH) [12], computes $h = H(r, m)$ for random ℓ -bit salts r , until h is a quadratic residue (r is then included in the signature for verification). A way to avoid this probabilistic method is to use a tweak, as proposed by Williams [31].² For any RSA modulus N , one can find four values $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \mathbb{Z}_N^*$ such that for any $h \in \mathbb{Z}_N^*$, there is a unique $i \in [1; 4]$ such that $\alpha_i^{-1}h \pmod{N}$ is a quadratic residue.³ When $p \equiv 3 \pmod{8}$ and $q \equiv 7 \pmod{8}$, one can use the set of values $\{1, -1, 2, -2\}$. This way, the signature becomes a so-called tweaked square root (α, s) , where s is a square root of

² Williams' paper [31] was primarily concerned with public key encryption. The idea of using a tweak for deterministic signing is implicit in the ISO/IEC 9796 standard published in 1991, and was later made more explicit in a paper by Kurosawa and Ogata [25].

³ The sufficient condition for this is that the pairs of Legendre symbols $((\frac{\alpha_i}{p}), (\frac{\alpha_i}{q}))$ take each of the four values $(1, 1), (-1, 1), (1, -1)$ and $(-1, -1)$ for exactly one α_i .

$\alpha^{-1}H(m) \bmod N$ for the correct value $\alpha \in \{1, -1, 2, -2\}$, and the verification algorithm now checks whether $\alpha s^2 = H(m) \bmod N$. This enables to define FDH Rabin-Williams signatures.

Since any quadratic residue modulo an RSA modulus N has four square roots, one must also specify which (tweaked) square root of the hash to use as the signature. There are basically two ways to proceed. The first one is simply to pick a square root at random. However, when no randomization (or randomization with only a small number of bits) is used in the input to the hash function, one must be careful not to output two non-trivially distinct square roots if the same message is signed twice, since this would reveal the factorization of the modulus N . In consequence, the signature algorithm must either be stateful and store all signatures previously output (which is cumbersome), or generate the bits for deciding which root to use pseudo-randomly.⁴ However, in constrained environments, implementors might be reluctant to pay the additional cost of a pseudorandom function (moreover, how exactly this derandomization is done is not always precisely discussed, and may have security implications as explained in [26]).

The second option is to define some deterministic rule telling which square root to use as the signature. The most popular way to do so is to use for N a Blum integer and to use the principal square root. A variant is to use what we call the absolute principal square root, *i.e.* $|s \bmod N|$, where s is the principal square root represented by an integer in $[-(N-1)/2; (N-1)/2]$. This turns out to also be the unique square root in $(\mathbb{J}_N)^+$. We will call these ways to choose a square root Principal Rabin-Williams (PRW) and Absolute Principal Rabin-Williams (APRW) respectively.⁵ When no randomization in the input to the hash function is used, the signature algorithm then becomes entirely deterministic (without having to appeal to an auxiliary pseudorandom function), which is attractive from an implementation point of view.

Bernstein [7] proposed an extensive study of possible variants of Rabin-Williams signature schemes depending on the length of the salt and the square root selection method. In particular, for FDH signatures, he showed a tight security reduction from the factoring assumption for the probabilistic square root selection method (Fixed Unstructured). On the other hand, for PRW and APRW, only a loose reduction from factoring is known using methods of Coron [11, 7]. Our main result is a tight security reduction from the $2-\Phi/4$ -Hiding problem for the PRW and APRW schemes, building on the results of [22]. Details can be found in Section 4.

Extending the Coron-Kakvi-Kiltz Meta-reduction Result. Recall that Coron’s meta-reduction result [12] as corrected by Kakvi and Kiltz [22] states that when the trapdoor function is *certified*, any security reduction from in-

⁴ This method was called Fixed Unstructured Rabin-Williams in [7], and Probabilistic Rabin-Williams (PRW) in [26].

⁵ PRW was called Fixed Principal in [7] and Deterministic Rabin-Williams (DRW) in [26], while APRW was called Fixed |Principal| in [7].

verting the trapdoor function to breaking FDH signatures must lose a factor q_s . Since this only applies for certified trapdoor functions, this leaves open the question of whether there might exist a tight reduction from inverting the trapdoor function to breaking FDH signatures when the trapdoor function is not certified. In particular, the question whether there exists a tight security reduction from factoring (or equivalently, computing modular square roots) for the PRW and APRW schemes was left as an open problem in [7]. However, we observe that the meta-reduction result still holds (namely, any security reduction from inverting the trapdoor function to breaking FDH signatures must lose a factor q_s) when the underlying trapdoor function is *gap one-way*, meaning that inverting the injective mode of the function is hard even with the help of an oracle distinguishing injective from non-injective modes of the trapdoor function. This implies in particular that if factoring with the help of an oracle solving the $2-\Phi/4$ -Hiding problem is hard, the PRW and APRW signature schemes cannot have a tight security reduction from the factorization problem. This essentially answers the open question of [7] regarding the security reductions for these schemes. Details can be found in Section 5.

1.3 Related and Future Work

Two constructions of lossy trapdoor functions based on modular squaring were previously proposed, however they are slightly more complicated than the basic Rabin trapdoor function. Mol and Yilek [27] gave a construction whose security relies on an assumption close in spirit (though more involved) to the $2-\Phi/4$ -Hiding assumption. Freeman *et al.* [16] gave a construction relying on the Quadratic Residuosity problem.

The cryptographic applications of the set $(\mathbb{J}_N)^+$ when N is a Blum integer were previously considered by Goldwasser *et al.* [19], Fischlin and Schnorr [15], and Hofheinz and Kiltz [21] (in this last paper, it was denoted \mathbb{QR}_N^+ and named *group of signed quadratic residues*). In particular, it was showed in [21] that under the factoring assumption, the Strong Diffie-Hellman problem [1] is hard in this group.

The Coron-Kakvi-Kiltz meta-reduction result [12, 22] was extended by Hofheinz *et al.* [20] to the case where the signature scheme is re-randomizable (rather than with unique signatures).

Kiltz *et al.* [24] showed that lossiness of RSA implies that the RSA-OAEP encryption scheme [5] meets indistinguishability under chosen-plaintext attacks in the standard model (under appropriate assumptions on the hash functions used to instantiate OAEP). An interesting question is whether lossiness of the Rabin trapdoor function can be used to argue about the security of Rabin-OAEP encryption as was done in [24] for RSA. Though from a theoretical point of view the results of [24] apply to OAEP used with any LTF, they provide some meaningful security insurance only when the amount of lossiness is sufficiently high. This requires more careful investigation in the case of Rabin-OAEP. As a first step in this direction, we note that if “multi-primes” pseudo-Blum integers $N = p_1 \cdots p_m$, with $p_1, \dots, p_m \equiv 1 \pmod{4}$ are indistinguishable from 2-primes

pseudo-Blum integers, lossiness of the Rabin trapdoor function with domain $(\mathbb{J}_N)^+$ can be amplified from 1 bit to $m - 1$ bits. Similar arguments were used for RSA in [24].

2 Preliminaries

2.1 General Notation

The set of integers i such that $a \leq i \leq b$ will be denoted $[a; b]$. The security parameter will be denoted k . A function f of the security parameter is said *negligible* if for any $c > 0$, $f(k) \leq 1/k^c$ for sufficiently large k . When S is a non-empty finite set, we write $s \leftarrow_{\$} S$ to mean that a value is sampled uniformly at random from S and assigned to s . By $z \leftarrow \mathcal{A}^{\mathcal{O}_1, \mathcal{O}_2, \dots}(x, y, \dots)$ we denote the operation of running the (possibly probabilistic) algorithm \mathcal{A} on inputs x, y, \dots with access to oracles $\mathcal{O}_1, \mathcal{O}_2, \dots$ (possibly none), and letting z be the output.

2.2 Basic Definitions

Given an (odd for most of what follows) integer N , the multiplicative group of integers modulo N is denoted \mathbb{Z}_N^* . This group has order $\phi(N)$ where $\phi(\cdot)$ is the Euler function. We denote \mathbb{J}_N the subgroup of \mathbb{Z}_N^* of all elements $x \in \mathbb{Z}_N^*$ with Jacobi symbol $(\frac{x}{N}) = 1$. This subgroup has index 2 and order $\phi(N)/2$ in \mathbb{Z}_N^* . Moreover it is efficiently recognizable even without the factorization of N since the Jacobi symbol is efficiently computable given only N . We also denote $\bar{\mathbb{J}}_N$ the coset of elements $x \in \mathbb{Z}_N^*$ such that $(\frac{x}{N}) = -1$. Finally, we denote \mathbb{QR}_N the subgroup of quadratic residues of \mathbb{Z}_N^* . This subgroup is widely believed not to be efficiently recognizable when N is composite and its factorization is unknown: this is the Quadratic Residuosity assumption.

We will represent elements of \mathbb{Z}_N as signed integers in $[-(N-1)/2, (N-1)/2]$. Given an integer x , we denote $|x \bmod N|$ the absolute value of $x \bmod N$. For any subset $S \subset \mathbb{Z}_N$, we denote $S^+ = S \cap [1; (N-1)/2]$ and $S^- = S \cap [-(N-1)/2; -1]$. Note that $(\mathbb{J}_N)^+$, $(\mathbb{J}_N)^-$, $(\bar{\mathbb{J}}_N)^+$ and $(\bar{\mathbb{J}}_N)^-$ form a partition of \mathbb{Z}_N^* .

We call an integer $N = pq$ which is the product of two distinct odd primes a *Blum integer* when $p, q \equiv 3 \pmod{4}$, and a *pseudo-Blum integer* when $p, q \equiv 1 \pmod{4}$, and we denote

$$\begin{aligned} \text{Bl}(k) &= \{(N, p, q) : N = pq, p \neq q \text{ are } [k/2]\text{-bit primes with } p, q \equiv 3 \pmod{4}\} \\ \tilde{\text{Bl}}(k) &= \{(N, p, q) : N = pq, p \neq q \text{ are } [k/2]\text{-bit primes with } p, q \equiv 1 \pmod{4}\} . \end{aligned}$$

We call a Blum integer $N = pq$ such that moreover $p \equiv 3 \pmod{8}$ and $q \equiv 7 \pmod{8}$ a *Williams integer*, and a pseudo-Blum integer such that $p \equiv 5 \pmod{8}$ and $q \equiv 1 \pmod{8}$ a *pseudo-Williams integer*. We denote

$$\begin{aligned} \text{Wi}(k) &= \{(N, p, q) \in \text{Bl}(k) : p \equiv 3 \pmod{8}, q \equiv 7 \pmod{8}\} \\ \tilde{\text{Wi}}(k) &= \{(N, p, q) \in \tilde{\text{Bl}}(k) : p \equiv 5 \pmod{8}, q \equiv 1 \pmod{8}\} . \end{aligned}$$

Note that:

- when N is a Blum integer, $-1 \in \mathbb{J}_N \setminus \mathbb{QR}_N$;
- when N is a pseudo-Blum integer, $-1 \in \mathbb{QR}_N$;
- when N is a Williams or a pseudo-Williams integer, $2 \in \bar{\mathbb{J}}_N$.

A quadratic residue modulo an RSA modulus $N = pq$ has four square roots, two of which are in $(\mathbb{Z}_N^*)^+$ and two of which are in $(\mathbb{Z}_N^*)^-$. The two square roots in $(\mathbb{Z}_N^*)^+$ will be called the *absolute* square roots in what follows. The following lemma will be important when proving lossiness of the Rabin trapdoor function.

Lemma 1. *Let $N = pq$ be a RSA modulus with $N \equiv 1 \pmod{4}$. Let $x \in \mathbb{QR}_N$, and let s_1 and s_2 be the two absolute square roots of x (the two other square roots being $-s_1$ and $-s_2$). Then:*

- if N is a Blum integer, exactly one s_i is in $(\mathbb{J}_N)^+$ and the other is in $(\bar{\mathbb{J}}_N)^+$; moreover if $s_i \in (\mathbb{J}_N)^+$ then either $s_i \in \mathbb{QR}_N$ or $-s_i \in \mathbb{QR}_N$;
- if N is a pseudo-Blum integer, then $s_1, s_2, -s_1, -s_2$ are either all in \mathbb{QR}_N , or all in $\mathbb{J}_N \setminus \mathbb{QR}_N$, or all in $\bar{\mathbb{J}}_N$.

Proof. Consider $x \in \mathbb{QR}_N$. Denote $x_p = x \pmod{p}$ and $x_q = x \pmod{q}$. Let also $\pm r_p$ and $\pm r_q$ denote the two square roots of respectively $x_p \pmod{p}$ and $x_q \pmod{q}$. The four square roots of x modulo N are obtained by combining $\pm r_p$ and $\pm r_q$ by the Chinese Remainder Theorem, *i.e.* there are two integers c_p and c_q such that the four square roots of x are $\pm(pc_p r_q \pm qc_q r_p) \pmod{N}$. Assume that one of the two absolute square roots is $s_1 = (pc_p r_q + qc_q r_p) \pmod{N}$ (the reasoning is similar if it is $-(pc_p r_q + qc_q r_p) \pmod{N}$). Then the other absolute square root satisfies $s_2 = \alpha(pc_p r_q - qc_q r_p) \pmod{N}$, with $\alpha = \pm 1$ so that:

$$\left(\frac{s_2}{p}\right) = \left(\frac{\alpha}{p}\right) \left(\frac{-1}{p}\right) \left(\frac{s_1}{p}\right) \quad \text{and} \quad \left(\frac{s_2}{q}\right) = \left(\frac{\alpha}{q}\right) \left(\frac{s_1}{q}\right) .$$

Consequently:

- when N is a Blum integer, s_1 and s_2 have opposite Jacobi symbols; moreover, assuming $s_1 \in (\mathbb{J}_N)^+$ then since -1 is a non-quadratic residue, either $s_1 \in \mathbb{QR}_N$ or $-s_1 \in \mathbb{QR}_N$;
- when N is a pseudo-Blum integer, we see that

$$\left(\frac{s_1}{p}\right) = \left(\frac{-s_1}{p}\right) = \left(\frac{s_2}{p}\right) = \left(\frac{-s_2}{p}\right)$$

$$\text{and} \quad \left(\frac{s_1}{q}\right) = \left(\frac{-s_1}{q}\right) = \left(\frac{s_2}{q}\right) = \left(\frac{-s_2}{q}\right) ,$$

from which the claim on the localization of the four square roots follows.

This concludes the proof. \square

Hence when N is a Blum integer, the two absolute square roots can easily be distinguished through their Jacobi symbol. In the following, given a Blum integer N and $x \in \mathbb{QR}_N$, we will call the unique square root of x which is in \mathbb{QR}_N the *principal* square root of x , and denote it $\mathbf{psr}(x)$. We will also call the unique square root of x which is in $(\mathbb{J}_N)^+$ the *absolute principal* square root of x , and will denote it $|\mathbf{psr}|(x)$. The notation is chosen so that $|\mathbf{psr}|(x) = \mathbf{psr}(x) \pmod{N}$.

Tweaked Square Roots. Let N be a Williams integer. Then for any $x \in \mathbb{Z}_N^*$ there is a unique $\alpha \in \{1, -1, 2, -2\}$ such that $\alpha^{-1}x \bmod N$ is a quadratic residue.⁶ The four pairs $(\alpha, s_i)_{i=1, \dots, 4}$ where $(s_i)_{i=1, \dots, 4}$ are the four square roots of $\alpha^{-1}x \bmod N$ are named the *tweaked square roots* of x , and α is named the *tweak*. Hence, (α, s) with $\alpha \in \{1, -1, 2, -2\}$ is a tweaked square root of $x \in \mathbb{Z}_N^*$ iff $\alpha s^2 = x \bmod N$. By extension, the principal tweaked square root of x is the unique tweaked square root (α, s) such that $s \in \mathbb{Q}\mathbb{R}_N$, and the absolute principal tweaked square root is the unique tweaked square root (α, s) such that $s \in (\mathbb{J}_N)^+$. Overloading the notation, they will be denoted respectively $\text{psr}(x)$ and $|\text{psr}|(x)$.

2.3 Trapdoor Functions

We recall some formal definitions associated with trapdoor functions (we follow closely the ones of [22]). We also introduce the concept of *gap one-way trapdoor function*, which is informally a trapdoor function which is hard to invert even when given access to an oracle which tells whether a member of the family is injective or lossy.

Definition 1 (Trapdoor Function). A trapdoor function (TDF) is a tuple of polynomial-time algorithms $\text{TDF} = (\text{InjGen}, \text{Eval}, \text{Invert})$ with the following properties:

- $\text{InjGen}(1^k)$: a probabilistic algorithm which on input the security parameter 1^k , outputs a public description pub (with implicitly understood domain \mathcal{D}_{pub}) and a trapdoor td ;
- $\text{Eval}(\text{pub}, x)$: a deterministic algorithm which on input pub and a point $x \in \mathcal{D}_{\text{pub}}$, outputs a point $y \in \{0, 1\}^*$; we denote $f_{\text{pub}} : x \mapsto \text{Eval}(\text{pub}, x)$;
- $\text{Invert}(\text{td}, y)$: a deterministic algorithm which on input td and a point $y \in \{0, 1\}^*$, outputs a point $x \in \mathcal{D}_{\text{pub}}$ when $y \in f_{\text{pub}}(\mathcal{D}_{\text{pub}})$ (and \perp otherwise).

We require that for any k and any (pub, td) possibly output by $\text{InjGen}(1^k)$, the function $f_{\text{pub}} : x \mapsto \text{Eval}(\text{pub}, x)$ be injective, and $y \mapsto \text{Invert}(\text{td}, y)$ be its inverse f_{pub}^{-1} . We also require that \mathcal{D}_{pub} and $f_{\text{pub}}(\mathcal{D}_{\text{pub}})$ be efficiently samplable.

Definition 2 (One-Way TDF). A trapdoor function $\text{TDF} = (\text{InjGen}, \text{Eval}, \text{Invert})$ is said to be (t, ε) -one-way if for any adversary \mathcal{A} running in time at most t , one has:

$$\Pr [\text{pub} \leftarrow \text{InjGen}(1^k), x \leftarrow_{\mathcal{S}} \mathcal{D}_{\text{pub}}, x' \leftarrow \mathcal{A}(\text{pub}, \text{Eval}(\text{pub}, x)) : x' = x] \leq \varepsilon .$$

Definition 3 (Certified TDF). A trapdoor function $\text{TDF} = (\text{InjGen}, \text{Eval}, \text{Invert})$ is said to be certified if there exists a deterministic polynomial-time algorithm Certify which, on input an arbitrary string pub (not necessarily generated by InjGen) returns 1 iff the function $x \mapsto \text{Eval}(\text{pub}, x)$ is injective over \mathcal{D}_{pub} .

⁶ This follows easily from the fact that the pairs of Legendre symbols $((\frac{\alpha}{p}), (\frac{\alpha}{q}))$ for $\alpha = 1, -1, 2$, and -2 are respectively $(1, 1)$, $(-1, -1)$, $(-1, 1)$ and $(1, -1)$.

Definition 4 (Lossy TDF). A lossy trapdoor function (LTF) with absolute lossiness ℓ is a tuple of algorithms $\text{LTF} = (\text{InjGen}, \text{LossyGen}, \text{Eval}, \text{Invert})$ such that $(\text{InjGen}, \text{Eval}, \text{Invert})$ is a TDF as per Definition 1, and moreover LossyGen is a probabilistic algorithm which on input 1^k , outputs a public description pub' such that the range of the function $f_{\text{pub}'} : x \mapsto \text{Eval}(\text{pub}', x)$ over $\mathcal{D}_{\text{pub}'}$ satisfies:

$$\frac{|\mathcal{D}_{\text{pub}'}|}{|f_{\text{pub}'}(\mathcal{D}_{\text{pub}'})|} \geq \ell .$$

We say that LTF is (t, ε) -secure if for any adversary \mathcal{A} running in time at most t , the following advantage is less than ε :

$$\left| \Pr[(\text{pub}, \text{td}) \leftarrow \text{InjGen}(1^k) : 1 \leftarrow \mathcal{A}(\text{pub})] - \Pr[\text{pub}' \leftarrow \text{LossyGen}(1^k) : 1 \leftarrow \mathcal{A}(\text{pub}')] \right| .$$

We say that LTF is a regular (ℓ, t, ε) -lossy trapdoor function if LTF is (t, ε) -secure and all functions generated by LossyGen are ℓ -to-1 on $\mathcal{D}_{\text{pub}'}$.

Remark 1. One can easily show that if TDF is a regular (ℓ, t, ε) -lossy TDF, then it is (t', ε') -one way with $t' \simeq t$ and $\varepsilon' \leq \varepsilon + 1/\ell$. Note in particular that asymptotically, if $\ell = \mathcal{O}(1)$ is constant (as is the case for the trapdoor functions considered in this paper), this only implies that TDF is weakly one-way [18].

Definition 5 (Gap One-Way TDF). A trapdoor function $\text{TDF} = (\text{InjGen}, \text{Eval}, \text{Invert})$ is said (t, ε, n) -gap one-way if for any adversary \mathcal{A} running in time at most t and making at most n queries to a $\text{Certify}(\cdot)$ oracle which on input a string pub , returns 1 iff the function $x \mapsto \text{Eval}(\text{pub}, x)$ is injective over \mathcal{D}_{pub} , one has:

$$\Pr \left[\text{pub} \leftarrow \text{InjGen}(1^k), x \leftarrow_{\S} \mathcal{D}_{\text{pub}}, x' \leftarrow \mathcal{A}^{\text{Certify}(\cdot)}(\text{pub}, \text{Eval}(\text{pub}, x)) : x' = x \right] \leq \varepsilon .$$

Informally, for a lossy TDF, being gap one-way means that inverting the injective mode of the function cannot be black-box reduced to the lossiness of the TDF. Note that for a certified TDF, being gap one-way is equivalent to being one-way since the Certify oracle can be efficiently implemented.

2.4 Signature Schemes

A signature scheme Σ is a tuple of algorithms $(\Sigma.\text{KeyGen}, \Sigma.\text{Sig}, \Sigma.\text{Ver})$ where $\Sigma.\text{KeyGen}(1^k)$ outputs a pair of public/secret key (pk, sk) , $\Sigma.\text{Sig}(\text{sk}, m)$, on input a secret key sk and a message $m \in \{0, 1\}^*$, outputs a signature σ , and $\Sigma.\text{Ver}(\text{pk}, m, \sigma)$, on input a public key pk , a message m , and a purported signature σ , either outputs 1 (accepts) or 0 (rejects). A signature scheme is said to have *unique signatures* if for all k , for any public key pk possibly output by $\text{KeyGen}(1^k)$, and any message $m \in \{0, 1\}^*$, there is exactly one string σ such that $\text{Ver}(\text{pk}, m, \sigma)$ accepts. The usual security definition for a signature scheme is existential unforgeability under chosen-message attacks (EUF-CMA security). We recall this definition in the full version of the paper [30].

FDH Signatures Based on an Arbitrary TDF. Let $\text{TDF} = (\text{InjGen}, \text{Eval}, \text{Invert})$ be a trapdoor function. The Full Domain Hash signature scheme TDF-FDH is defined as follows: the key generation algorithm $\text{KeyGen}(1^k)$ runs $\text{InjGen}(1^k)$ to obtain (pub, td) , selects a random hash function $\mathbf{H} : \{0, 1\}^* \rightarrow f_{\text{pub}}(\mathcal{D}_{\text{pub}})$, and sets $\text{pk} = (\text{pub}, \mathbf{H})$ and $\text{sk} = \text{td}$. The signature algorithm, on input td and m , computes $h = \mathbf{H}(m)$ and returns $\sigma = \text{Invert}(\text{td}, h)$. The verification algorithm, on input pub , m and σ , checks that $\text{Eval}(\text{pub}, \sigma) = \mathbf{H}(m)$. This scheme can be shown EUF-CMA secure in the Random Oracle Model under the assumption that TDF is (strongly) one-way [4, 11], but the reduction loses a factor q_s , where q_s is the maximal number of signature queries of the adversary, and this loss cannot be avoided assuming that TDF is certified [12, 22].

3 The $2\text{-}\Phi/4$ -Hiding Assumption and Lossiness of the Rabin Trapdoor Function

3.1 Definition

We introduce the $2\text{-}\Phi/4$ -Hiding assumption, an extension of the traditional Φ -Hiding assumption to the case $e = 2$. The Φ -Hiding assumption, introduced by Cachin *et al.* in [9], roughly states that given an RSA modulus $N = pq$ and a random prime $3 \leq e < N^{1/4}$, it is hard to distinguish whether e divides $\phi(N)$ or not (when $e \geq N^{1/4}$ and $e|\phi(N)$, N can be factored using Coppersmith’s method for finding small roots of univariate modular equations [10, 9]). Kiltz *et al.* [24] were the first to observe that the Φ -Hiding assumption can be interpreted in terms of lossiness of the RSA trapdoor permutation.

The original definition of the Φ -Hiding assumption was formulated for primes e randomly drawn in $[3; N^{1/4}]$. Since in practice RSA is often used with a fixed, small prime e (*e.g.* $e = 3$ or $e = 2^{16} + 1$), Kakvi and Kiltz [22] introduced the Fixed-Prime Φ -Hiding assumption, which states, for a fixed prime e , that it is hard, given an RSA modulus $N = pq$, to distinguish whether e divides $\phi(N)$ or not (the exact statement of the assumption is slightly different for $e = 3$ and $e > 3$ in order to avoid trivial distinguishers). The $2\text{-}\Phi/4$ -Hiding assumption is the extension of the Fixed-Prime Φ -Hiding assumption to the case $e = 2$. Since for an RSA modulus N (more generally for any number which has at least two distinct prime factors) one always has that 4 divides $\phi(N)$, the problem will be to distinguish whether 2 divides $\phi(N)/4$ or not. Moreover, when $N \equiv 3 \pmod{4}$, one can check that 2 always divides $\phi(N)/4$, so that the instances will be restricted to RSA moduli such that $N \equiv 1 \pmod{4}$. As a matter of fact, distinguishing whether 2 divides $\phi(N)/4$ or not when $N \equiv 1 \pmod{4}$ turns out to be equivalent to distinguishing Blum integers from pseudo-Blum integers. Indeed, if N is a Blum integer, then $p = 4p' + 3$ and $q = 4q' + 3$, so that $\phi(N) = 4(2p' + 1)(2q' + 1)$ and $2 \nmid (\phi(N)/4)$. On the other hand, if N is a pseudo-Blum integer, then $p = 4p' + 1$ and $q = 4q' + 1$, so that $\phi(N) = 16p'q'$ and $2 | (\phi(N)/4)$. We now precisely formalize the assumption.

Definition 6 (2- $\Phi/4$ -Hiding Assumption). We say that the 2- $\Phi/4$ -Hiding problem is (t, ε) -hard if for any algorithm \mathcal{A} running in time at most t , the following advantage is less than ε :

$$\text{Adv}^{2-\Phi/4}(\mathcal{A}) \stackrel{\text{def}}{=} \left| \Pr[(N, p, q) \leftarrow_{\S} \text{Bl}(k) : 1 \leftarrow \mathcal{A}(N)] - \Pr[(N, p, q) \leftarrow \widetilde{\text{Bl}}(k) : 1 \leftarrow \mathcal{A}(N)] \right| .$$

A variant of this problem is obtained by switching from Blum integers to Williams integers, *i.e.* replacing $\text{Bl}(k)$ and $\widetilde{\text{Bl}}(k)$ in the above definition by respectively $\text{Wi}(k)$ and $\widetilde{\text{Wi}}(k)$. Clearly, the hardness of this variant is polynomially related to the hardness of the original problem, under the plausible assumption that roughly half of Blum, resp. pseudo-Blum integers are Williams, resp. pseudo-Williams integers.

3.2 Lossiness of the Rabin and Rabin-Williams Trapdoor Functions

We now show that the 2- $\Phi/4$ -Hiding assumption implies that squaring is a lossy trapdoor function over the domains \mathbb{QR}_N or $(\mathbb{J}_N)^+$, for $N \equiv 1 \pmod{4}$, with respectively two bits or one bit of lossiness. The injective mode corresponds to N being a Blum integer, and the lossy mode corresponds to N being a pseudo-Blum integer.

The Rabin LTFs. We first define two related LTFs, that we name respectively the Principal Rabin LTF PR-LTF and the Absolute Principal Rabin LTF APR-LTF as follows:

- on input 1^k , PR-LTF.InjGen and APR-LTF.InjGen both draw $(N, p, q) \leftarrow_{\S} \text{Bl}(k)$, and output $\text{pub} = N$ and $\text{td} = (p, q)$;
- on input 1^k , PR-LTF.LossyGen and APR-LTF.LossyGen both draw $(N, p, q) \leftarrow_{\S} \widetilde{\text{Bl}}(k)$, and output $\text{pub}' = N$;
- the domain is $\mathcal{D}_N = \mathbb{QR}_N$ for PR-LTF, and $\mathcal{D}_N = (\mathbb{J}_N)^+$ for APR-LTF; the evaluation algorithms PR-LTF.Eval(N, x) and APR-LTF.Eval(N, x) both output $f_N(x) = x^2 \pmod{N}$; in both cases $f_N(\mathcal{D}_N) = \mathbb{QR}_N$ in injective mode;
- the inversion algorithm PR-LTF.Invert($(p, q), y$) outputs the principal square root $\text{psr}(y)$, while APR-LTF.Invert($(p, q), y$) outputs the absolute principal square root $|\text{psr}|(y)$ (for N a Blum integer and $y \in \mathbb{QR}_N$).

Theorem 1. Assuming the 2- $\Phi/4$ -Hiding problem is (t, ε) -hard, the Principal Rabin trapdoor function PR-LTF is a regular $(4, t, \varepsilon)$ -LTF, while the Absolute Principal Rabin trapdoor function APR-LTF is a regular $(2, t, \varepsilon)$ -LTF.

Proof. Indistinguishability of the injective and lossy modes is exactly the 2- $\Phi/4$ -Hiding problem. It follows from Lemma 1 that when N is a Blum integer, any $y \in \mathbb{QR}_N$ has exactly one pre-image in \mathbb{QR}_N or $(\mathbb{J}_N)^+$, while when N is pseudo-Blum integer, any y in the range $f_N(\mathbb{QR}_N)$ has exactly 4 pre-images in \mathbb{QR}_N , and any y in the range $f_N((\mathbb{J}_N)^+)$ has exactly 2 pre-images in $(\mathbb{J}_N)^+$. \square

The Rabin-Williams LTFs. The PR-LTF and APR-LTF LTFs can be straightforwardly extended to what we call the Principal Rabin-Williams LTF PRW-LTF and Absolute Principal Rabin-Williams LTF APRW-LTF as follows:

- on input 1^k , PRW-LTF.InjGen and APRW-LTF.InjGen both draw a random Williams integer $(N, p, q) \leftarrow_{\S} \text{Wi}(k)$, and output $\text{pub} = N$ and $\text{td} = (p, q)$;
- on input 1^k , PRW-LTF.LossyGen and APRW-LTF.LossyGen both draw a random pseudo-Williams integer $(N, p, q) \leftarrow_{\S} \widetilde{\text{Wi}}(k)$ and output $\text{pub}' = N$;
- the domain of PRW-LTF is $\mathcal{D}_N = \{1, -1, 2, -2\} \times \mathbb{QR}_N$, while the domain of APRW-LTF is $\mathcal{D}_N = \{1, -1, 2, -2\} \times (\mathbb{J}_N)^+$; the evaluation algorithms PRW-LTF.Eval($N, (\alpha, x)$) and APRW-LTF.Eval($N, (\alpha, x)$) compute the function $f_N(\alpha, x) = \alpha x^2 \bmod N$; in both cases $f_N(\mathcal{D}_N) = \mathbb{Z}_N^*$ in injective mode;
- algorithm PRW-LTF.Invert($(p, q), y$) computes the principal tweaked square root $\text{psr}(y)$, while APRW-LTF.Invert($(p, q), y$) computes the absolute principal tweaked square root $|\text{psr}|(y)$ (for N a Williams integer and $y \in \mathbb{Z}_N^*$).

Theorem 2. *Under the assumption that Williams and pseudo-Williams integers are (t, ε) -indistinguishable, the Principal Rabin-Williams trapdoor function PRW-LTF is a regular $(4, t, \varepsilon)$ -LTF, while the Absolute Principal Rabin-Williams trapdoor function APRW-LTF is a regular $(2, t, \varepsilon)$ -LTF.*

Proof. Indistinguishability of the injective and lossy modes is exactly indistinguishability of Williams and pseudo-Williams integers, which follows from the $2\text{-}\Phi/4$ -Hiding assumption and the additional (reasonable) assumption that roughly half of Blum, resp. pseudo-Blum integers, are Williams, resp. pseudo-Williams integers. Injectivity of f_N for both PRW-LTF and APRW-LTF follows directly from Lemma 1 and the discussion about tweaked square roots in Section 2. Assume now that N is a pseudo-Williams integer, and let $y \in f_N(\mathcal{D}_N)$ with $\mathcal{D}_N = \{1, -1, 2, -2\} \times \mathbb{QR}_N$. We show that y has exactly 4 pre-images in \mathcal{D}_N , which will establish that PRW-LTF is 4-to-1 on \mathcal{D}_N . Let $(\alpha, x) \in \mathcal{D}_N$ be such that $\alpha x^2 = y \bmod N$. Then by Lemma 1, y has at least 4 pre-images in \mathcal{D}_N , all with the same tweak α . Assume that y has an extra pre-image $(\alpha', x') \in \mathcal{D}_N$ with $\alpha' \neq \alpha$. Note that when $N = pq$ is a pseudo-Williams integer (*i.e.* $p \equiv 5 \pmod{8}$ and $q \equiv 1 \pmod{8}$), the pairs of Legendre symbols $\left(\frac{\alpha}{p}\right), \left(\frac{\alpha}{q}\right)$ for $\alpha = 1, -1, 2,$ and -2 are respectively $(1, 1), (1, 1), (-1, 1)$ and $(-1, 1)$. Hence it must be that $\alpha' = -\alpha$, so that $x^2 = -(x')^2 \bmod N$. Let a be any square root of -1 modulo N . Since $a^2 = -1 \bmod N$, we observe (denoting $p = 8p' + 5$ and $q = 8q' + 1$) that:

$$\begin{aligned} \left(\frac{a}{p}\right) &\equiv a^{\frac{p-1}{2}} \equiv a^{\frac{8p'+4}{2}} \equiv (-1)^{2p'+1} \equiv -1 \pmod{p} \\ \left(\frac{a}{q}\right) &\equiv a^{\frac{q-1}{2}} \equiv a^{\frac{8q'}{2}} \equiv (-1)^{2q'} \equiv 1 \pmod{q} \end{aligned}$$

so that $a \in \mathbb{J}_N$. Hence, we have that $x^2 = (ax')^2 \bmod N$, with $x, x' \in \mathbb{QR}_N$. Yet by Lemma 1, one should have $ax' \in \mathbb{QR}_N$ as well, which is impossible since $a \in \mathbb{J}_N$. Hence y has exactly 4 pre-images in \mathcal{D}_N .

The proof that APRW-LTF is 2-to-1 on $\mathcal{D}_N = \{1, -1, 2, -2\} \times (\mathbb{J}_N)^+$ is very similar. See the full version of the paper [30]. \square

4 Application to Rabin-Williams Signatures

There are two very close ways to define deterministic Rabin-Williams FDH signatures, called *principal* and *|principal|* in the terminology of Bernstein [7]. We will use the name *Absolute Principal* Rabin-Williams signatures for the latter in this paper. Before defining precisely these schemes, we stress that the exact definition of the verification algorithm is important, especially with respect to how a forgery is defined (since a forgery is exactly a string which is accepted by the verification algorithm). Hence, to be more precise, we will define in total four “real” signature schemes: Principal Rabin-Williams (PRW), Absolute Principal Rabin-Williams (APRW), as well as two slightly different variants that we call PRW* and APRW*, which differ from respectively PRW and APRW only in their verification algorithm. We will also define a “theoretical” scheme PRW** where the verification algorithm is inefficient (this will be necessary to establish a clean security reduction). For the five schemes, the signing algorithm first hashes the message $h = \mathbf{H}(m)$; then, for the PRW, PRW*, and PRW** schemes, the signing algorithm returns the principal tweaked square root of h , whereas for the APRW and APRW* schemes, the signing algorithm returns the absolute principal tweaked square root of h . In all the following, we assume that if h is not coprime with N , the signing algorithm outputs some fixed signature, *e.g.* $(1, 1)$. Since this happens only with negligible probability when \mathbf{H} is modeled as a random oracle, this does not affect the security analysis.

We now proceed to the formal definition. First, all the schemes share exactly the same key generation algorithm:

- (A)PRW(*, **).KeyGen(1^k): on input the security parameter 1^k , draw uniformly at random $(N, p, q) \leftarrow_s \text{Wi}(k)$. Select a hash function $\mathbf{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_N$. The public key is $\text{pk} = (N, \mathbf{H})$ and the secret key is $\text{sk} = (p, q)$.

Note that the hash function will usually be selected once for each security parameter k and common to all public keys, but this affects the security proof only up to negligible terms, see Bernstein [7].

The signing algorithm for PRW, PRW*, and PRW** on one hand, and for APRW and APRW* on the other hand, are the same, and are defined as follows:

- PRW(*, **).Sig(sk, m): To sign a message m , compute $h = \mathbf{H}(m)$, and output the principal tweaked square root $\sigma = (\alpha, s) = \text{psr}(h)$.
- APRW(*).Sig(sk, m): To sign a message m , compute $h = \mathbf{H}(m)$, and output the absolute principal tweaked square root $\sigma = (\alpha, s) = |\text{psr}|(h)$.

The verification algorithms for the five schemes are very close, and differ only with respect to an additional check on the Jacobi symbol of the signature made for PRW* and APRW*, and on the quadratic residuosity of the signature for PRW**. They are defined as follows:

- (A)PRW^(*,**).Ver(pk, m, σ): To check a purported signature σ = (α, s) on message m, first ensure that s ∈ S, and then check that αs² = H(m) mod N. Accept if this holds, and reject otherwise;

where the set S is defined as:

- S = Z_N^{*} for PRW, S = J_N for PRW*, and S = QR_N for PRW**;
- S = (Z_N^{*})⁺ for APRW and S = (J_N)⁺ for APRW*.

Note that the verification algorithm is (presumably) inefficient for PRW** since it needs to decide whether the signature is indeed the principal square root, *i.e.* a quadratic residue.

The following claims are straightforward:

- in PRW, each message has exactly four valid signatures: (α, s₁) = |psr|(H(m)), (α, −s₁), and (α, s₂), (α, −s₂) with s₂ ∈ (J_N)⁺;
- in PRW*, each message has exactly two valid signatures: (α, s) = |psr|(H(m)) and (α, −s);
- in PRW**, each message has a unique valid signature: (α, s) = psr(H(m));
- in APRW, each message has exactly two valid signatures: |psr|(H(m)) and (α, s₂) with s₂ ∈ (J_N)⁺;
- in APRW*, each message has a unique valid signature: |psr|(H(m)).

We now relate the security of PRW, PRW*, and PRW** on one hand, and APRW and APRW* on the other hand.

Lemma 2. *The security of PRW, PRW* and PRW** on one hand, and APRW and APRW* on the other hand, is related as depicted in Figure 1, where an arrow labeled (t, f(ε)) from scheme A to scheme B means that if scheme A is (t, ε, q_h, q_s)-EUF-CMA secure in the ROM, then scheme B is (t', f(ε), q_h, q_s)-EUF-CMA secure for t' ≈ t.*

Proof. See the full version of the paper [30]. □

Hence, one can see that PRW and PRW* on one hand, and APRW and APRW* on the other hand, have the same security up to a factor 2. In other words, omitting the additional check on the Jacobi symbol has negligible impact on security. Since computing a Jacobi symbol might be costly (in particular, it is more expensive than modular squaring), we see that PRW and APRW are superior in terms of security/efficiency trade-off.

In the following, we give a tight reduction for PRW** and APRW* from the 2-Φ/4-Hiding assumption, which extends to PRW and APRW by Lemma 2. It is easy to see that the PRW**, resp. APRW* signature scheme is exactly the instantiation of the generic TDF-FDH scheme recalled in Section 2.4 with PRW-LTF, resp. APRW-LTF. In order to conclude about the security of these schemes, we appeal to the main result of [22]. This theorem was originally stated for trapdoor *permutations*, but it can be straightforwardly extended to trapdoor functions such that D_{pub} and f_{pub}(D_{pub}) are efficiently samplable.




Fig. 1. Set of reductions proved in Lemma 2. An arrow labeled $(t, f(\varepsilon))$ from scheme A to scheme B means that if scheme A is $(t, \varepsilon, q_h, q_s)$ -EUF-CMA secure in the ROM, then scheme B is $(t', f(\varepsilon), q_h, q_s)$ -EUF-CMA secure for $t' \simeq t$. The reduction from 2- $\Phi/4$ -Hiding to breaking PRW** and APRW* is Theorem 4.

Theorem 3 ([22]). *Assume LTF is a regular (ℓ, t', ε') -LTF for $\ell \geq 2$. Then for any (q_h, q_s) , the TDF-FDH signature scheme instantiated with LTF is $(t, \varepsilon, q_h, q_s)$ -EUF-CMA secure in the ROM, where*

$$\varepsilon = \left(\frac{2\ell - 1}{\ell - 1} \right) \varepsilon' \quad \text{and} \quad t = t' - q_h T_{\text{Eval}} ,$$

where T_{Eval} is the time to run algorithm Eval of LTF.

Theorem 4. *Assuming the 2- $\Phi/4$ -Hiding problem is (t', ε') -hard, then for any (q_h, q_s) , the PRW** signature scheme is $(t, \varepsilon, q_h, q_s)$ -EUF-CMA secure, where $\varepsilon = 7\varepsilon'/3$ and $t = t' - \mathcal{O}(q_h k^3)$, and the APRW* signature scheme is $(t, \varepsilon, q_h, q_s)$ -EUF-CMA secure, where $\varepsilon = 3\varepsilon'$ and $t = t' - \mathcal{O}(q_h k^3)$.*

Proof. This follows directly from Theorems 2 and 3 (noting that \mathbb{QR}_N and $(\mathbb{J}_N)^+$ are efficiently samplable). Combined with Lemma 2, this yields tight security reductions for PRW and APRW (see Figure 1 for a clear picture). \square

Remark 2. The global security reduction from the 2- $\Phi/4$ -Hiding assumption to breaking the signature scheme is slightly looser for PRW (factor 28/3) than for APRW (factor 6 = 18/3). We also remark that a PRW signature oracle is (potentially) slightly more powerful than an APRW signature oracle because it reveals some non-trivial information regarding the quadratic residuosity of the square roots of the hash of the message (whereas this information, which is unnecessary for verifying signatures, is “canceled” in an APRW signature oracle). Since APRW signatures are not more costly than PRW signatures (and even slightly more communication efficient), these two observations make a case in favor of APRW signatures.

As explained in [22], these results can be extended to PSS-R [6], allowing a smaller overhead of the randomized signature under the 2- $\Phi/4$ -Hiding assumption. It seems also likely (though we have not checked the details) that the same techniques can be used to prove a tight security reduction from the 2- $\Phi/4$ -Hiding assumption for Rabin-Williams Partial Domain Hash signatures [13, 17].

5 Extending the Coron-Kakvi-Kiltz Meta-reduction Result

In this section, we complete the picture of the security of FDH signatures by extending Coron’s meta-reduction result [11] as corrected by Kakvi and Kiltz [22]. In a nutshell, this result says that if a trapdoor function TDF is certified, then any reduction from inverting the trapdoor function to breaking the EUF-CMA security of the TDF-FDH signature scheme must lose a factor q_s (the maximal number of signature queries made by the forger) in its time-to-success ratio. The theorem below extends this to trapdoor functions which are not necessarily certified, assuming that TDF is gap one-way. The proof is straightforwardly adapted from the one of [11, 22]: when simulating the forger, the meta-reduction checks as a preliminary step that the public key received from the reduction contains a parameter `pub` which defines an injective function. When TDF is certified, this can be done efficiently by the meta-reduction itself. In the variant below, the meta-reduction uses a `Certify` oracle for this step, hence breaking the gap one-wayness (rather than classical one-wayness) of the trapdoor function.

Theorem 5. *Let TDF be a trapdoor function. Let $t_R, \varepsilon_R, n, \varepsilon_F, q_h, q_s$ be functions of the security parameter with $q_h > q_s$. Assume there exists a reduction \mathcal{R} which $(t_R, \varepsilon_R, n, \varepsilon_F, q_h, q_s)$ -reduces breaking the one-wayness of TDF to breaking EUF-CMA security of the TDF-FDH signature scheme. Then there exists a meta-reduction \mathcal{M} which (t_M, ε_M, n) -breaks the gap one-wayness of TDF, where:*

$$t_M \leq (n + 1)t_R$$

$$\varepsilon_M \geq \varepsilon_R - \varepsilon_F \cdot \frac{n \cdot \exp(-1)}{q_s} \left(1 - \frac{q_s}{q_h}\right)^{-1}.$$

Proof. A precise definition of a (black-box) reduction and a sketch of the proof are provided in the full version of the paper [30]. \square

Remark 3. Theorem 5 above can be straightforwardly extended to any non-interactive computational problem which is hard relative to a `Certify` oracle (instead of the one-wayness of the underlying trapdoor function).

Consequences for RSA and Rabin-Williams FDH Signatures. We know by Theorem 8 of [22] that RSA-FDH with public exponents $e < N^{1/4}$ has a tight security reduction from the Φ -Hiding assumption. By Theorem 7 of [22] we also know that RSA-FDH with public exponents $e > N^{1/4}$ cannot have a tight security reduction from the problem of inverting RSA —nor any non-interactive hard problem— since RSA is certified for this class of exponents [23]. Theorem 5 above implies that it is unlikely as well that RSA-FDH with $e < N^{1/4}$ can have a tight security reduction from inverting RSA: unless inverting RSA with the help of an oracle solving the Φ -Hiding problem is easy, any reduction from inverting RSA to breaking the EUF-CMA security of RSA-FDH with $e < N^{1/4}$ must lose a factor q_s .

This extends to Rabin-Williams FDH signatures as follows: unless computing modular square roots (or equivalently factoring) with the help of an oracle solving the $2\text{-}\Phi/4\text{-Hiding}$ problem is easy, any reduction from factoring to breaking the EUF-CMA security of the PRW and APRW schemes must lose a factor q_s .

References

- [1] M. Abdalla, M. Bellare, and P. Rogaway. The Oracle Diffie-Hellman Assumptions and an Analysis of DHIES. In D. Naccache, editor, *Topics in Cryptology - CT-RSA 2001*, volume 2020 of *LNCS*, pages 143–158. Springer, 2001.
- [2] M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged Public-Key Encryption: How to Protect against Bad Randomness. In M. Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 232–249. Springer, 2009.
- [3] M. Bellare, D. Hofheinz, and S. Yilek. Possibility and Impossibility Results for Encryption and Commitment Secure under Selective Opening. In A. Joux, editor, *Advances in Cryptology - EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 1–35. Springer, 2009.
- [4] M. Bellare and P. Rogaway. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
- [5] M. Bellare and P. Rogaway. Optimal Asymmetric Encryption. In A. D. Santis, editor, *Advances in Cryptology - EUROCRYPT '94*, volume 950 of *LNCS*, pages 92–111. Springer, 1994.
- [6] M. Bellare and P. Rogaway. The Exact Security of Digital Signatures - How to Sign with RSA and Rabin. In U. M. Maurer, editor, *Advances in Cryptology - EUROCRYPT '96*, volume 1070 of *LNCS*, pages 399–416. Springer, 1996.
- [7] D. J. Bernstein. Proving Tight Security for Rabin-Williams Signatures. In N. P. Smart, editor, *Advances in Cryptology - EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 70–87. Springer, 2008.
- [8] A. Boldyreva, S. Fehr, and A. O'Neill. On Notions of Security for Deterministic Encryption, and Efficient Constructions without Random Oracles. In D. Wagner, editor, *Advances in Cryptology - CRYPTO 2008*, volume 5157 of *LNCS*, pages 335–359. Springer, 2008.
- [9] C. Cachin, S. Micali, and M. Stadler. Computationally Private Information Retrieval with Polylogarithmic Communication. In J. Stern, editor, *Advances in Cryptology - EUROCRYPT '99*, volume 1592 of *LNCS*, pages 402–414. Springer, 1999.
- [10] D. Coppersmith. Finding a Small Root of a Univariate Modular Equation. In U. M. Maurer, editor, *Advances in Cryptology - EUROCRYPT '96*, volume 1070 of *LNCS*, pages 155–165. Springer, 1996.
- [11] J.-S. Coron. On the Exact Security of Full Domain Hash. In M. Bellare, editor, *Advances in Cryptology - CRYPTO 2000*, volume 1880 of *LNCS*, pages 229–235. Springer, 2000.
- [12] J.-S. Coron. Optimal Security Proofs for PSS and Other Signature Schemes. In L. R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 272–287. Springer, 2002.
- [13] J.-S. Coron. Security Proof for Partial-Domain Hash Signature Schemes. In M. Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *LNCS*, pages 613–626. Springer, 2002.

- [14] S. Fehr, D. Hofheinz, E. Kiltz, and H. Wee. Encryption Schemes Secure against Chosen-Ciphertext Selective Opening Attacks. In H. Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 381–402. Springer, 2010.
- [15] R. Fischlin and C.-P. Schnorr. Stronger Security Proofs for RSA and Rabin Bits. *Journal of Cryptology*, 13(2):221–244, 2000.
- [16] D. M. Freeman, O. Goldreich, E. Kiltz, A. Rosen, and G. Segev. More Constructions of Lossy and Correlation-Secure Trapdoor Functions. In P. Q. Nguyen and D. Pointcheval, editors, *Public Key Cryptography - PKC 2010*, volume 6056 of *LNCS*, pages 279–295. Springer, 2010.
- [17] C. Gentry. How to Compress Rabin Ciphertexts and Signatures (and More). In M. K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004*, volume 3152 of *LNCS*, pages 179–200. Springer, 2004.
- [18] O. Goldreich. *Foundations of Cryptography - Volume 1, Basic Tools*. Cambridge University Press, 2001.
- [19] S. Goldwasser, S. Micali, and R. L. Rivest. A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
- [20] D. Hofheinz, T. Jager, and E. Knapp. Waters Signatures with Optimal Security Reduction. In M. Fischlin, J. Buchmann, and M. Manulis, editors, *Public Key Cryptography - PKC 2012*, volume 7293 of *LNCS*, pages 66–83. Springer, 2012.
- [21] D. Hofheinz and E. Kiltz. The Group of Signed Quadratic Residues and Applications. In S. Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *LNCS*, pages 637–653. Springer, 2009.
- [22] S. A. Kakvi and E. Kiltz. Optimal Security Proofs for Full Domain Hash, Revisited. In D. Pointcheval and T. Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 537–553. Springer, 2012.
- [23] S. A. Kakvi, E. Kiltz, and A. May. Certifying RSA. In X. Wang and K. Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 404–414. Springer, 2012.
- [24] E. Kiltz, A. O’Neill, and A. Smith. Instantiability of RSA-OAEP under Chosen-Plaintext Attack. In T. Rabin, editor, *Advances in Cryptology - CRYPTO 2010*, volume 6223 of *LNCS*, pages 295–313. Springer, 2010.
- [25] K. Kurosawa and W. Ogata. Efficient Rabin-type Digital Signature Scheme. *Des. Codes Cryptography*, 16(1):53–64, 1999.
- [26] G. Leurent and P. Q. Nguyen. How Risky Is the Random-Oracle Model? In S. Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *LNCS*, pages 445–464. Springer, 2009.
- [27] P. Mol and S. Yilek. Chosen-Ciphertext Security from Slightly Lossy Trapdoor Functions. In P. Q. Nguyen and D. Pointcheval, editors, *Public Key Cryptography - PKC 2010*, volume 6056 of *LNCS*, pages 296–311. Springer, 2010.
- [28] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In C. Dwork, editor, *Symposium on Theory of Computing - STOC 2008*, pages 187–196. ACM, 2008.
- [29] M. O. Rabin. Digitalized signatures and public-key functions as intractable as factorization. Technical Report 212, MIT Laboratory for Computer Science, 1979.
- [30] Y. Seurin. On the Lossiness of the Rabin Trapdoor Function. Full version of this paper. Available at <http://eprint.iacr.org/2013/256>.
- [31] H. C. Williams. A modification of the RSA public-key encryption procedure. *IEEE Transactions on Information Theory*, 26(6):726–729, 1980.