

How Risky is the Random-Oracle Model?

Gaëtan Leurent* and Phong Q. Nguyen†

July 23, 2009

Abstract

RSA-FDH and many other schemes secure in the Random-Oracle Model (ROM) require a hash function with output size larger than standard sizes. We show that the random-oracle instantiations proposed in the literature for such cases are weaker than a random oracle, including the proposals by Bellare and Rogaway from 1993 and 1996, and the ones implicit in IEEE P1363 and PKCS standards: for instance, there is a practical 2^{30} preimage attack on BR93 for 1024-bit digests. Next, we study the security impact of hash function defects for ROM signatures. As an extreme case, we note that any hash collision would suffice to disclose the master key in the ID-based cryptosystem by Boneh *et al.* from FOCS '07, and the secret key in the Rabin-Williams signature for which Bernstein proved tight security at EUROCRYPT '08. Interestingly, for both of these schemes, a slight modification can prevent these attacks, while preserving the ROM security result. We give evidence that in the case of RSA and Rabin/Rabin-Williams, an appropriate PSS padding is more robust than all other paddings known.

1 Introduction

The Random-Oracle Model (ROM), in which hash functions are viewed as random oracles, goes back to at least Fiat and Shamir [21]. Popularized by Bellare and Rogaway [7], it has been the subject of much debate. On the one hand, it is widespread in research papers and standards, because ROM schemes are usually more efficient, their security proofs can be based on well-studied assumptions, and can sometimes be tight. In fact, many standardized public-key schemes are, at best, only proven secure in the ROM, *e.g.* RSA-OAEP [8] and RSA-PSS [9].

On the other hand, the ROM is not an assumption (on the hash function): there is no random-oracle definition which a public function can hope to satisfy. A security proof in the standard model (SM) precisely gives assumptions which are sufficient to ensure security properties. By contrast, a ROM security proof only shows that an attack which does not break the proof assumptions must exploit a property not satisfied by the random-oracle simulation of the security proof. This allowed Canetti *et al.* [14] to build signature and encryption schemes which are secure in the ROM, but insecure for any (efficient) implementation of the random oracle, because such implementations can be simulated by a Universal Turing machine, and therefore be distinguished [41] from a random oracle. However, the constructions [14, 41, 25, 4] showing the limitations of the ROM are arguably “unnatural” and significantly differ from real-world constructions. Still, one should be careful with idealized security models: like all Merkle-Damgård/Davies-Meyer hash functions, MD5 was

*DGA and ENS, France. <http://www.eleves.ens.fr/leurent/>

†INRIA and ENS, France. <http://www.di.ens.fr/~pnguyen/>

provably collision-resistant [61, 11] (up to the birthday bound) in the ideal cipher model (with respect to the block cipher underlying the compression function); yet, computing MD5 collisions only costs a few seconds now [54].

This stresses the importance of studying the actual security of schemes proven in the ROM. Unfortunately, very few ROM schemes have also been proven secure in the SM [53, 12]; and for several cases, there is evidence that a security proof in the SM is unlikely [44, 49, 19, 33]. Recent breakthroughs in the cryptanalysis of hash functions [59, 58, 56] have shown that standard hash functions like MD5 or SHA-1 are far from behaving like random oracles. However, the impact on the public-key world has been limited so far, with the exception of [56], which constructs two colliding X.509 certificates for different identities and public keys, and has recently been extended in [54] to construct a rogue CA certificate.

But to study the actual security, one needs to know how the random oracle will be instantiated in practice, should the scheme ever be used. Often, the random-oracle output size matches that of standard hash functions (like 160 bits for SHA-1) or the upcoming SHA-3. In this case, standard hash functions are most likely to be used, despite well-known properties of MD-iterated hash functions (such as the derivation of $h(m_1||m_2)$ from $h(m_1)$ and m_2) which make them easily differentiable from random oracles. But RSA-FDH [9] and many other ROM schemes (such as [29, 16, 17, 32, 10, 23, 13]) actually require a “non-standard” hash function. First, the output may not be a uniformly distributed bitstring: it could be residue classes, or elliptic curve points, *etc.*, fortunately it is known how to deal with such situations given an instantiation with arbitrary output $\{0, 1\}^n$. But if the output bit-length is larger than standard sizes (*e.g.* RSA-FDH which needs at least 1024 bits), it is unclear how the oracle will be instantiated. To the best of our knowledge, the only proposals of random-oracle instantiations supporting arbitrary output bit-length are the following: two historical instantiations proposed by Bellare and Rogaway in their seminal papers [7] (on the ROM) and [9] (on FDH and PSS), recent constructions by Coron *et al.* in the full version of [18], and the instantiations implicit in PKCS#1 v2.1 [51] and IEEE P1363 [28] standards. It seems that none of these instantiations have been analyzed in the literature, except [18] in the indistinguishability framework of Maurer *et al.* [41].

This also raises the question of the impact of potential defects in random-oracle instantiations. When an article provides a ROM security proof, it usually does not say how to instantiate the random oracle, neither what might happen if the hash function does not behave like a random oracle. Assume that Alice implements a scheme ROM-secure under a well-known computational assumption. Several years later, the assumption still stands, but Alice learns that her random-oracle implementation is not as perfect as she thought: Should Alice worry? Are the risks confined to chosen-message existential forgery and ciphertext distinguishability? If Alice had the choice between two (equally efficient) schemes secure in the ROM under the same assumption, maybe Alice would rather choose the least risky one, in terms of robustness to hash function defects: the scheme with less stringent security requirements on the hash function, and the least affected if ever the requirements are not satisfied.

OUR RESULTS. We analyze the main proposals [7, 9, 51, 28] of random-oracle instantiations supporting arbitrary output bit-length. While none of these proposals made it clear what was exactly the expected security guarantee, one might argue that an instantiation with output bit-length n should offer $2^{n/2}$ resistance to collisions, and 2^n resistance to preimages, as required for SHA-3. We show that the proposals fall short of those bounds: for instance, for 1024-bit digests, we give a 2^{30} preimage attack on BR93 [7] and a 2^{106} collision attack on BR96 [9]. We note

that the instantiations implicit in PKCS [51] and IEEE [28] standards are not collision-resistant: independently of the output size, collisions follow directly from SHA-1 collisions, which cost only 2^{61} according to [31, 58]. And we show that when applied to the compression functions of MD5 or SHA-1, the theoretical constructions of Coron *et al.* [18] are no more collision-resistant than MD5 or SHA-1 themselves. This highlights the difficulty of instantiating/simulating a random oracle, and motivates the study of the impact of hash defects on schemes secure in the ROM.

As a second contribution, we show that, while the ROM is useful to detect structural flaws, it can hide very different security requirements on the hash function, and very different damages in case of hash defects, independently of the computational assumption and the tightness of the security reduction. We illustrate this phenomenon with a well-known class of ROM schemes: padding-based signatures from trapdoor one-way functions, such as RSA, Rabin, Rabin-Williams and ESIGN. While it is often believed that a hash collision may at worst give rise to an existential forgery, we show that for several secure signatures proposed in the literature [9, 29, 13, 10, 23], collisions or slight hash function defects can have much more dramatic consequences, namely key-recovery attacks. However, this does not contradict the ROM security proofs, and does not mean that such signatures cannot be secure with a proper hash function. Our most interesting examples are related to Rabin and Rabin-Williams signatures, but the issues are not restricted to factoring-based schemes: in Appendix C, we show similar problems for a recent lattice-based signature scheme [23]. For instance, we remark that any hash collision discloses the master key in the ID-based cryptosystem of Boneh *et al.* [13], and the secret key in the Rabin-Williams signature scheme for which Bernstein [10] recently proved tight security, which was not mentioned in either [13, 10]. Interestingly, we show that a slight modification of the signing process can prevent our collision-based key-recovery attacks, while preserving the ROM security result. We give evidence that in the case of RSA and Rabin/Rabin-Williams, an appropriate PSS padding (with large salt) is more robust than all other paddings known, especially deterministic ones and randomized paddings with small randomness, including Katz-Wang [32] which achieve tightness.

To put things into perspective, consider the Rabin-Williams signature included in the IEEE P1363 standard [28]: it can be optionally deterministic or randomized, but the deterministic version has no tight reduction. Since tightness was a key factor in standardizing RSA-PSS over RSA-FDH, one might be tempted to replace this deterministic Rabin-Williams by its “tight” variant analyzed by Bernstein [10], but doing so would have led to a chosen-message key-recovery attack from any SHA-1 collision because such collisions provide collisions on the RO-instantiation of the P1363 standard. Alternatively, there would have also been a chosen-message key-recovery attack if IEEE P1363 had simply used the BR93 RO-instantiation [7], due to the preimage attack.

Some of the problems we discuss are related to the issue of how to derandomize a signature scheme. As a side-remark, we show that the derandomization technique proposed by Granboulan [27] to fix ESIGN during the NESSIE European project is not completely sound: when applied to ESIGN or DSA, it may leak the secret key. Finally, while this paper focuses on the ROM, we believe that in general, not just in the ROM, it is also interesting to identify necessary (possibly minimal) assumptions for security, and to assess the security impact when assumptions do not hold. This is useful when comparing cryptographic schemes, and can complement provable security results.

ROAD MAP. We assume the reader is familiar with hash functions, the ROM [7] and provable security for signatures [26]. In Sect. 2, we recall and analyze random-oracle instantiations for large output size. Next, we study the implications of hash function defects for ROM signatures based on a trapdoor one-way function and a padding: we recall such schemes in Sect. 3. In Sect. 4, we study

the robustness of derandomized Rabin/Rabin-Williams signatures, which are used in [10] and the ID-based cryptosystem of [13]. In Sect. 5, we compare the robustness of RSA and Rabin/Rabin-Williams. App. A gives more information on the ID-based cryptosystem [13]. App. B gives a converse to the security proof of PRab [9]. In App. C, we study the robustness of a different kind of scheme: the GPV lattice-based signature scheme [23].

2 Random-Oracle Instantiations for Large Output

In this section, we describe and analyze random-oracle instantiations supporting arbitrary output bit-length that have been proposed in the literature, namely [7, 9, 18] and the instantiations implicit in the PKCS#1 v2.1 [51] and IEEE P1363 [28] standards. For completeness, we also briefly discuss collision-resistant hash functions such as [15, 40]. While some of the instantiations make use of MD5, that alone is insufficient to discard them. Indeed, though the collision-resistance of MD5 is seriously broken [59, 35], many usages of MD5 are not threatened yet: for instance, there is still no practical attack on HMAC-MD5.

2.1 The 1993 Instantiation by Bellare and Rogaway

Description. In their seminal paper on the ROM [7], Bellare and Rogaway gave guidelines to instantiate a random oracle (see [7, Sect. 6]), but the only explicit construction is the following one, which we call BR93:

- Let $h_4 : \{0, 1\}^{512} \rightarrow \{0, 1\}^{128}$ be the first compression function of MD5, that is the compression function evaluated with the initial IV of MD5.
- Let $h' : \{0, 1\}^{256} \rightarrow \{0, 1\}^{64}$ defined by $h'(x)$ being the first 64 bits of $h_4((xx) \oplus C)$, for a randomly chosen 512-bit constant C . The function h' defines a pseudo-random number generator $h''(x) : \{0, 1\}^{192} \rightarrow \{0, 1\}^*$ by counter as follows¹: $h''(x) = h'(x\langle 0 \rangle) || h'(x\langle 1 \rangle) || h'(x\langle 2 \rangle) \dots$ where $\langle i \rangle$ is the encoding of i into 64 bits.
- Finally, the BR93 instantiation of the random oracle is the truncation (prefix) of $h(x) : \{0, 1\}^* \rightarrow \{0, 1\}^*$ defined as follows. First, one applies a padding to x by adding a bit 1 and enough bits 0 to obtain a bitstring x' whose bit-length is a multiple of 128. Then, if we divide x' into 128-bit blocks as $x' = x'_0 \dots x'_{n-1}$, then $h(x) = h''(x'_0\langle 0 \rangle) \oplus h''(x'_1\langle 1 \rangle) \oplus \dots \oplus h''(x'_{n-1}\langle n-1 \rangle)$, that is, $h(x)$ is the XOR of the n streams produced by each of the x'_i .

Weaknesses. We claim that BR93 is much weaker than a random oracle with respect to collision and preimage resistance, independently of the choice of the underlying compression function (MD5 here). In the proceedings version [39] of this paper, we gave a preimage attack based on Wagner's generalized birthday [57]. But Bellare [3] recently pointed out to us that there is actually a much more efficient (and simpler) preimage attack, which was presented by Bellare and Micciancio in [5] for different purposes, and which we will now describe. Indeed, BR93 turns out to be an instantiation of the so-called XHASH [5], where each message-block produces an independent output, and the outputs are combined by a simple XOR sum. Finding preimages in such a function is also known as the k -sum problem: given some lists L_1, L_2, \dots, L_k of r -bit values, we want to find $l_1 \in L_1, \dots, l_k \in L_k$ such that $\bigoplus_{i=1}^k l_i = 0$. This problem can be solved by Wagner's generalized

¹The paper [7] actually says 224 instead of 192 for the input size of h' , but that would be incompatible with the definition of h' as $224 + 64 = 288 > 256$.

birthday [57], which gives a collision attack with complexity 2^{67} and preimage attack with complexity 2^{68} (see [39]). However, the parameters of this k -sum instance can be quite degenerate: we can have k as big as we like, if we use a message of appropriate length. In this case, we can use a very efficient attack based on linear algebra when $k \geq r$.

Concretely, we give a preimage attack on n bits with polynomial complexity n^3 using messages of n blocks: by applying this preimage attack twice, one can compute two random preimages of the same output, and therefore obtain collisions. For instance, a preimage attack on 1024 bits with messages of 1024 blocks only costs 2^{30} elementary operations, which is very practical.

Linearization attack. We describe a preimage attack on BR93 truncated to n bits, where the main part of the attack is the resolution of a linear system of size n . Let us choose two distinct 128-bit blocks a and b , and let \tilde{h} be the target value of our preimage attack. We build the following vector v and matrix Δ , where all the calls to h'' are truncated to n bits, and x'_n is the padding of a message of exactly n blocks:

$$\begin{aligned} v &= h''(a\langle 0 \rangle) \oplus h''(a\langle 1 \rangle) \oplus \cdots \oplus h''(a\langle n-1 \rangle) \oplus h''(x'_n\langle n \rangle) \oplus \tilde{h} \\ \Delta &= [h''(a\langle i \rangle) \oplus h''(b\langle i \rangle)]_{i=0}^{n-1} \end{aligned}$$

We solve the equation $\Delta \times u = v$ by linear algebra. Since the matrix Δ has size $n \times n$, there is a good probability to find a solution. Then we build the message x as follows:

$$x_i = \begin{cases} a & \text{if } u^{[i]} = 0 \\ b & \text{if } u^{[i]} = 1 \end{cases}$$

This will give a preimage because

$$\begin{aligned} h(x) &= h''(x'_n\langle n \rangle) \oplus \bigoplus_{i=0}^{n-1} h''(x'_i\langle i \rangle) \\ &= h''(x'_n\langle n \rangle) \oplus \bigoplus_{u^{[i]}=0} h''(a\langle i \rangle) \oplus \bigoplus_{u^{[i]}=1} h''(b\langle i \rangle) \\ &= h''(x'_n\langle n \rangle) \oplus \bigoplus_{i=0}^{n-1} h''(a\langle i \rangle) \oplus \bigoplus_{u^{[i]}=1} (h''(a\langle i \rangle) \oplus h''(b\langle i \rangle)) \\ &= \tilde{h} \oplus v \oplus \Delta \times u = \tilde{h} \end{aligned}$$

This attack was first described by Bellare and Micciancio in [5, Appendix A], and is cited in Wagner's paper [57]. This linearization attack has also been used to break IFSB, a hash function that can be described as a k -sum instance [52].

2.2 The 1996 Instantiation by Bellare and Rogaway

Description. In their paper [9] on PSS, Bellare and Rogaway proposed another instantiation [9, App. A], which we call BR96: let $H = \text{MD5}$ or SHA-1 , and define $h_{\text{BR96}}(x)$ as the appropriate truncation (prefix) of:

$$H(\text{const}\langle 0 \rangle x) || H(\text{const}\langle 1 \rangle x) || H(\text{const}\langle 2 \rangle x) || \dots$$

where the constant `const` should be unique to h . If another instantiation is needed, one should change `const`. BR96 is very close to a previous construction described in the OAEP paper [8] where H above is replaced by the 80-bit truncation of SHA-1, but that construction was not explicitly recommended to instantiate a random oracle, though it was used as a building block in an implementation of RSA-OAEP.

Weaknesses. We note that since H is a MD function, h_{BR96} can be viewed as the concatenation of MD functions, where each $H_i : x \mapsto H(\text{const}\langle i \rangle x)$ is a distinct iterative hash function, which implies:

- h_{BR96} can be distinguished from a random oracle. More precisely, BR96 suffers from the same extension problems as any MD function: if the output size of h_{BR96} is an exact multiple of that of H , and if m_1 has appropriate size, then $h_{BR96}(m_1||m_2)$ can be computed from $h_{BR96}(m_1)$ and m_2 .
- h_{BR96} is weaker than a random oracle with respect to collision and preimage resistance, independently of the choice of the underlying hash function (except its output size), thanks to Joux’s multicollision technique [30].

Recall that Joux [30] showed that the concatenation of two or more MD-iterated hash functions has roughly the same security as a single hash function. More precisely, if one concatenates k iterated hash functions with an internal size of n bits each, [30] finds a collision in the concatenation for a workload of $n^{k-1} \times 2^{n/2}$, and preimages for an imprecise workload of $\text{poly}(n^k)2^n$.

A closer analysis of the collision attack shows that the cost $n^{k-1} \times 2^{n/2}$ can actually be reduced to $[(n/2)^{k-1} + (n/2)^{k-2} + \dots + 1] \times 2^{n/2} \leq (n/2)^{k-1} \times (n/2)/(n/2-1) \times 2^{n/2} \approx (n/2)^{k-1} \times 2^{n/2}$. And there seems to be a more efficient preimage attack by generalizing the basic preimage attack against two hash functions as follows. First, build a $2^{n^{k-1}/2^{k-2}}$ -multicollision on the first hash function F_1 , and look for an extra block that maps this multicollision to the target value of F_1 . Then build a multicollision in F_2 using the messages of the first multicollision: each collision in F_2 requires a set of $2^{n/2}$ messages, which will be built from $n/2$ colliding pairs in F_1 . Thus we should get a $2^{n^{k-2}}/2^{k-3}$ -multicollision in F_1 . We will also use the last n colliding pairs for a preimage search on F_1 . This gives us a $2^{n^{k-2}}/2^{k-3}$ -multicollision in $F_1||F_2$ which is also a preimage. We apply the technique iteratively to build a 2^n -multicollision for $F_1||F_2||\dots||F_{k-1}$ which is also a preimage. If we compute F_k on the set of 2^n colliding messages, we expect to find one preimage against the full concatenation. The most expensive steps of this attack are the preimage search, because the collision finding steps all have complexity $O(n^k \times 2^{n/2})$. The preimage step on F_i requires to compute F_i on 2^n messages, which are made of n block pairs of length $n^{i-2}/2^{i-2}$ and one block of length $n^{i-2}/2i - 3$. If we do an amortized analysis, each computation requires to hash 2 blocks from message pairs, and the final block, which gives a cost of $n^{i-2}/2^{i-4} \times 2^n$. The cost of the full preimage search is roughly equivalent to the cost of the last preimage search, which is $n^{k-2}/2^{k-4} \times 2^n$.

We now apply this to BR96. For instance, if H is MD5, we can find collisions in 1024 bits of the output with a workload of essentially $64^7 \cdot 2^{64} = 2^{106}$, where the colliding messages will be of length $64^7 = 2^{42}$ blocks; and we can find preimages of 1024 bits of the output with a workload of $128^6/2^4 \cdot 2^{128} = 2^{166}$. These complexities are clearly impractical and do not threaten [9], but they are much lower than the theoretical security of a 1024-bit random oracle.

For the same reason, BR96 is also malleable. For instance, we can create pairs of messages x_0, x_1 such that $H(\text{const}\langle i \rangle x_0) = H(\text{const}\langle i \rangle x_1)$ for all i ’s except the last one. We will build a multicollision set of $2^{n/4}$ such messages, and we expect to find one quadruplet such that $H(x_0) \oplus H(x_1) \oplus H(x_2) \oplus H(x_3) = 0$. In Appendix C, we show how this kind of malleability can be exploited

to attack GPV [23].

As another example, consider the previous instantiation of BR96 for 1024-bit digests, using MD5. We can find two near-collisions where the most significant 384 bits are equal, with of a workload of essentially $64^2 \cdot 2^{64} = 2^{76}$. Such near-collisions give existential forgeries for the historical version [46] of ESIGN.

2.3 Recent Instantiations by Coron *et al.* (CDMP)

Description. Coron *et al.* [18] (CDMP) proposed several variations of Merkle-Damgård to build a random oracle from an (ideal) compression function or an (ideal) block-cipher using the Davies-Meyer mode. They proposed four variants of MD for *input* domain extensions (namely, *Prefix-Free Encoding*, *Dropping Some Output Bits*, *Using NMAC*, and *Using HMAC*) and one scheme (only in the full version of [18]) for *output* domain extension. The output extension scheme is similar to BR96, but the counter is included after the message (which is reminiscent of the MGF1 pseudo-random number generator used in several standards [28, 51]):

$$h_{CDMP}(x) = H(x\langle 0 \rangle) || H(x\langle 1 \rangle) || H(x\langle 2 \rangle) || \dots$$

where H is one of the four input extension schemes. This choice is due to efficiency considerations, but we will see that it has a strong security impact. The main advantage of [18] is its security proof: all the constructions are proved indifferentiable from a random oracle (in the sense of Maurer *et al.* [41]), if the underlying compression function is a random oracle, or if it uses the Davies-Meyer mode with an ideal block cipher. However, no recommendation is given in [18] for the choice of the underlying compression function (or the underlying block cipher for Davies-Meyer). So strictly speaking, unlike [7, 9], there was no fully concrete proposal of a random-oracle instantiation: still, one may want to apply the constructions to usual compression functions.

Weaknesses. One should be careful not to overestimate the significance of indifferentiability security proofs: in practice, there is no ideal compression function. It was shown by [6] that none of the CDMP constructions necessarily preserve collision-resistance: they give (theoretical) examples of collision-resistant compression functions for which the resulting hash function is not collision-resistant.

While [18] was presented as a fix to the MD construction, we show that if one applies these fixes to MD5 or SHA-1, one can still find collisions in the new hash function (independently of the chosen output length) with the same cost as the original MD5 or SHA-1. This means that [18] does not address collision attacks on MD5 and SHA-1. To see this, we first show that the four input extensions are not collision resistant if applied to the compression functions of MD5 or SHA-1. This is trivial for *Dropping Some Output Bits*, *Using NMAC*, and *Using HMAC*, because these constructions are nested: an inner collision becomes an outer collision. So the only potentially tricky case is *Prefix-Free Encoding*, for which [18] proposed only two instantiations:

- prepend the message size as the first block. It turns out that MD5/SHA-1 collision attacks [59, 58] can be extended to this case, because the number of blocks of colliding messages produced is equal and already known in advance, and it is well-known that existing MD5/SHA-1 collision attacks can be extended to any given IV.
- use the first bit of each message block as a flag to distinguish the last message block. Since the number of blocks in MD5/SHA-1 colliding messages is very small, and the first bit of each block is random looking, we can simply produce random collisions until one has the required form.

Now, because of the iterated structure of the four input extensions, these collisions give rise to collisions in the output extension h_{CDMP} . More generally, while h_{CDMP} is indiffereniable from a random oracle if H is also indiffereniable, any collision in H becomes a collision in h_{CDMP} if H has an iterative structure like MD or the four input extensions: namely, $H(x_0) = H(x_1)$ implies $H(x_0\langle i \rangle) = H(x_1\langle i \rangle)$ and therefore $h_{CDMP}(x_0) = h_{CDMP}(x_1)$.

Hence, we have shown that if the CDMP constructions are applied to the compression functions of MD5 or SHA-1 for an arbitrary output size, the cost of producing collisions remains essentially the same as for MD5 or SHA-1 [31]. Of course, one could try to use different compression functions, but no concrete recommendation is given in [18].

2.4 Instantiations in PKCS and IEEE Standards

Description. No cryptographic standard currently specifies a random-oracle instantiation for arbitrary size. However, several instantiations are implicit in PKCS #1 v2.1 [51] and IEEE P1363 [28], because RSA-OAEP [8] and RSA-PSS [9] are standardized:

- RSA-OAEP requires two random oracles G and H with small input size (less than the RSA modulus), which are both instantiated in PKCS by the MGF1 pseudo-random number generator [51]. Recall that MGF1 is simply a hash function in counter mode like h_{CDMP} , except that the counter is over four bytes: $MGF1(x) = h(x\langle 0 \rangle) || h(x\langle 1 \rangle) || h(x\langle 2 \rangle) || \dots$, where h is either SHA-1 or a SHA-2.
- RSA-PSS also requires two random oracles G and H , but while G still has small input size, H has a small output size but possibly large inputs. In PKCS, H is instantiated by SHA-1 or SHA-2, and G is instantiated by MGF1.

Thus, none of the oracles required by RSA-OAEP and RSA-PSS have both a large input and output as would be required by RSA-FDH. Still, MGF1 is a potential random-oracle instantiation, because it supports arbitrarily large input and output.

There is another implicit instantiation in IEEE P1363 [28]. Indeed, it includes a Rabin-Williams signature using a variant of the PSS encoding [9] (as described in [29]) called EMSA-PSS in [51] and EMSA4 in [28]: the main difference between EMSA-PSS and PSS [9] (described in Sect. 3.2) is that the message is first hashed before going through the PSS encoding. But it is specified in [28] that the salt can optionally be set to zero, in which case “*the signature scheme is deterministic, similar to Full-Domain Hashing*”. Thus, one can view EMSA-PSS with zero salt as an instantiation of a FDH: since the padding constants are zero, this amounts to essentially hash the message twice in a row, then apply MGF1; concatenate the output and the input of MGF1, and append the “BC” byte.

Weaknesses. The case of MGF1 has already been analyzed with the CDMP case in the previous subsection: using SHA-1 or any MD-iterated hash function, the cost of producing collisions in MGF1 remains as low as for the underlying hash function. And EMSA-PSS with zero salt is clearly no more collision-resistant than the underlying hash function. Note also that the “BC” byte makes it differentiable from a random oracle. Hence, independently of the output size chosen, finding collisions on the PKCS/IEEE instantiations costs as low as for the underlying hash function MD5 or SHA-1.

2.5 Provably collision-resistant hash functions

To conclude this section, we briefly discuss the case of hash functions which are provably collision-resistant under appropriate computational assumptions. Though not designed nor recommended to instantiate random oracles, they might be potential candidates since they usually support large output size. But it is folklore that none should be viewed nor used as a random oracle, because they have special properties which are not satisfied by a random oracle, typically malleability. Consider for instance two recent collision-resistant hash functions:

- VSH [15], which is collision-resistant provided that a certain problem related to factorization is hard. The output set is \mathbb{Z}_N^\times , where N is hard to factor.
- SWIFFT [40], which is (asymptotically) collision-resistant and one-way, provided that certain lattice approximation problems are hard. The smallest output size is 528 bits, but larger sizes are possible.

These functions are malleable in the following sense. In [40], it is noted that for any two inputs x_1 and x_2 such that $x_1 + x_2$ is a valid input, $\text{SWIFFT}(x_1) + \text{SWIFFT}(x_2) = \text{SWIFFT}(x_1 + x_2)$. By definition of VSH [15], it is easy to generate $M_0 \neq M_1$ such that $4\text{VSH}(M_0) \equiv \text{VSH}(M_1) \pmod{N}$ where N is the public modulus. More generally, for any product $s > 1$ of very small distinct primes (chosen among the primes used by the VSH compression function), it is easy to generate $M_0 \neq M_1$ such that $s^2\text{VSH}(M_0) \equiv \text{VSH}(M_1) \pmod{N}$.

We will see that such malleability relationships can be exploited to attack certain signature schemes. The malleability of SWIFFT can be exploited to attack the GPV signature [23] (see Appendix C), and the malleability of VSH can be exploited to attack Rabin and Rabin-Williams signatures. But we stress that neither VSH or SWIFFT were recommended to be used with these signatures.

3 Padding-based Signatures in the Random-Oracle Model

We study the impact of hash function defects for the class of ROM-secure signatures obtained by combining a trapdoor one-way function/permutation and a padding. More precisely, we consider secure versions of RSA, Rabin, Rabin-Williams and ESIGN with appropriate paddings: FDH [7], PSS [9], PFDH [16] and KW [32]. This section briefly recalls these components.

3.1 Signatures from Trapdoor One-Way Functions

Let σ denote the raw signature algorithm which, given as input a message $m \in \mathcal{M}$ outputs a signature $\sigma(m) \in \mathcal{S}$: the algorithm can be either deterministic or probabilistic. We consider signatures based on a trapdoor one-way function $f : \mathcal{S} \rightarrow \mathcal{M}$:

- If f is 1-to-1, we obtain a deterministic scheme with $\sigma(m) = f^{-1}(m)$.
- If f is many-to-1, we obtain a probabilistic scheme with $\sigma(m)$ selected uniformly at random in the set $f^{-1}(m)$ of preimages: we require that the trapdoor enables such preimage sampling.

Verification checks that a given signature s belongs to \mathcal{S} and that $f(s) = m$.

RSA. Let (N, e, d) be the usual RSA keys where the exponent d is secret. We have $\mathcal{M} = \mathcal{S} = \mathbb{Z}_N$, and take the RSA trapdoor permutation $f(x) = x^e$ defined over \mathcal{M} , whose inverse is $f^{-1}(x) = x^d$.

Rabin [50]. Let $N = pq$ be a usual RSA modulus. Then we take the squaring function $f(x) = x^2$, which is a 4-to-1 mapping from $\mathcal{S} = \mathbb{Z}_N^\times$ to the subgroup \mathcal{M} of quadratic residues mod N . Inverting f is equivalent to factoring N .

Rabin-Williams [60]. This is a variation of Rabin signatures based on tweaks, using a modulus $N = pq$ such that $p \equiv 3 \pmod{8}$ and $q \equiv 7 \pmod{8}$. This has two notable features: one can take $\mathcal{M} = \mathbb{Z}_N$ (rather than having to deal with quadratic residues), and one can obtain mappings onto \mathcal{M} which are either 4-to-1 or 1-to-1, and whose inversion is equivalent to factoring. For any $m \in \mathcal{M}$, there are exactly four triplets $(e, f, s) \in \mathcal{S} = \{-1, 1\} \times \{1, 2\} \times \{0, \dots, N-1\}$ such that $m \equiv efs^2$, and these so-called tweaked square roots can all be efficiently computed using p and q . Furthermore, the *principal* tweaked square root is the only square root such that e is 1 if m is a square modulo q , otherwise -1; f is 1 if em is a square modulo p , otherwise 2; and s is a square modulo $N = pq$. We thus obtain two trapdoor one-way functions, depending on the choice of \mathcal{S} :

- By taking the 4-to-1 mapping, we obtain the probabilistic signature scheme PRW.
- By taking the 1-to-1 mapping where \mathcal{S} is confined to principal tweaked square roots, we obtain the deterministic signature scheme DRW. Ignoring technical details, this is essentially the Rabin-Williams used in IEEE P1363 [28].

ESIGN [46, 47]. We only give an informal description. ESIGN uses an RSA modulus of the form $N = p^2q$ such that p, q have bit-length k , and N has bit-length $3k$. There is a small public exponent $e \geq 8$ not necessarily coprime with $\phi(N)$. The one-way function is the truncation of the RSA permutation $f(x) = x^e$ to its k most significant bits. This is a many-to-one mapping from $\mathcal{S} = \mathbb{Z}_N$ to $\mathcal{M} = \{0, \dots, \lfloor N/2^{2k} \rfloor\}$, whose inversion problem is called AER [47].

3.2 Paddings

A padding Π specifies how to sign arbitrary messages $m \in \{0, 1\}^*$: it may be deterministic or randomized. Then the signature is $\sigma(\Pi(m))$, with additional data in the case of PFDH. With the exception of PSS, all the following paddings use a full-domain hash h from $\{0, 1\}^*$ to \mathcal{M} .

FDH [7]. This deterministic padding is simply $\Pi(m) = h(m)$.

PFDH [16]. One selects a salt $r \leftarrow_{\mathcal{R}} \{0, 1\}^k$, and let $\Pi(m) = h(m||r)$ where h is a random oracle from $\{0, 1\}^*$ to \mathcal{M} . The signature must include the salt r .

KW [32]. If m has never been signed, select a one-bit salt $r \leftarrow_{\mathcal{R}} \{0, 1\}$, and let $\Pi(m) = h(r||m)$.

PSS [9]. Assume to simplify that $\mathcal{M} = \{0, 1\}^k$ where $k \geq k_0 + k_1$ for some integers k_0 and k_1 . Two random oracles $h : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and $g : \{0, 1\}^{k_1} \rightarrow \{0, 1\}^{k-k_1}$ are used. We let g_1 be the function which on input $w \in \{0, 1\}^{k_1}$ returns the first k_0 bits of $g(w)$, and let g_2 be the function which on input $w \in \{0, 1\}^{k_1}$ returns the remaining $k - k_0 - k_1$ bits of $g(w)$. For any message $m \in \{0, 1\}^*$, one selects $r \in_R \{0, 1\}^{k_0}$ and let $w = h(m||r)$ and $r^* = g_1(w) \oplus r$. Finally, $\Pi(m) = w||r^*||g_2(w)$.

Standards. The PKCS [51] and IEEE P1363 [28] standards actually implement a slightly different version of PSS, called EMSA-PSS [29]: the main difference is that the message is first hashed before going through the PSS encoding. As mentioned earlier in Sect. 2.4, Rabin-Williams in IEEE P1363 is implemented as essentially DRW-PSS, but the salt can optionally be zero, in which case it becomes DRW-FDH with a specific RO-instantiation.

3.3 Derandomization

Several schemes [32, 10, 13] crucially require to derandomize the signature or the padding:

- The main scheme analyzed by Bernstein [10] is derandomized PRW-FDH with the requirement that if ever the same message is submitted twice, the same signature should be output.

- The ID-based cryptosystem of Boneh *et al.* [13] uses derandomized Rabin-FDH for mapping identities to secret keys: see Appendix A.
- To implement RSA-KW, Katz and Wang [32] suggested to select the one-bit salt deterministically from a secret key and the message.

We will see that how the derandomization is performed has a big impact on the security. Bernstein [10] did not specify how this derandomization should be performed: he only mentioned that if ever the same message is submitted twice, the same signature should be output, otherwise an attacker would be able to compute two random tweaked square roots of the same element (by signing twice the same message), which discloses the secret key. But Katz and Wang discussed [32, Sect. 4.1] that issue in the context of RSA signatures, and proposed the following methods:

- KW1: select the nonce as $r = h'(K||m)$, where h' is an independent random oracle, K is an additional secret key and m is the message.
- KW2: select $r = h'(K||h(m))$ with the same notations as KW1.
- KW3: select $r = F_K(m)$ where F is a PRF and K is an additional secret key. This is the method used in the ID-based cryptosystem [13], which is shown to preserve the ROM security proof.

In order to derandomize ESIGN to fix its security proof (see [55, 48]), Granboulan [27] earlier proposed a slightly different method: select $r = \phi(h(m)||K||c)$, where ϕ is a one-way function with uniformly distributed output, K is an additional secret key, and c is an optional counter. The counter is necessary in ESIGN, because the signature process may actually fail for certain nonces. This derandomization technique was later adopted with a specific choice of ϕ in the revised submission [22] of ESIGN to the NESSIE European project.

3.4 Security Results

It is well-known that for any trapdoor permutation f , the signatures FDH, PFDH and PSS are all provably secure in the ROM under the hardness of inverting the permutation, but the security proof is loose [7, 20]. In the particular case of RSA, all these security proofs can be improved, by exploiting the multiplicativity of RSA: RSA-FDH remains loose [16, 7, 9], but RSA-PFDH and RSA-PSS have a tight reduction provided that the salt is sufficiently large (see [16, 20]). Surprisingly, RSA-KW also has a tight reduction [32]. These security results also apply to DRW (under the factoring assumption), because DRW uses a 1-to-1 mapping which is homomorphic (see [29]).

The picture is a bit different with Rabin and PRW, since they use 4-to-1 mappings, but squaring is homomorphic. The derandomized versions of Rabin-FDH and PRW-FDH have a tight reduction (see [13] for Rabin-FDH and [10] for PRW-FDH). Rabin-PSS has a tight reduction (see [9]). For a complete picture of the ROM security of all Rabin-Williams variants, see [10].

ESIGN-FDH [47] and ESIGN-PSS [36] both have a loose security proof under the AER assumption, but the proof of ESIGN-FDH requires a more restricted security model than usual (see [48, 55]). There is no tightness because the one-way function is not multiplicative.

4 Robustness of Derandomized Signatures

We now study the robustness of derandomized signatures described in Sect. 3, for the derandomization methods proposed in [27, 32, 13], which we described in Sect. 3.3. This derandomization is crucial for the tightness of certain security proofs. We focus on derandomized Rabin-FDH and

PRW-FDH.

4.1 Soundness

First of all, one should make sure that the derandomization technique does not affect the security proof of the randomized scheme. For KW3, this was proved in [13, App. B] using of course the PRF assumption. And a similar argument can be proved for KW1 and KW2, but both require another random oracle, which is debatable if the goal is to understand what are the minimal assumptions.

For the fourth randomization method however, Granboulan [27] only gave an informal argument. In fact, we note that his method is not completely sound. More precisely, if the informal argument was correct, it would also apply to the choice $r = \phi(m||K||c)$. Now, assume that we take for ϕ an MD-iterated function for which it is easy to find collisions, but still, the function is one-way with uniformly distributed output: one potential example is SHA-1. Then an adversary could create a collision (m, m') where m and m' have the same length, which would imply that $\phi(m||K||c) = \phi(m'||K||c)$ for all K and c . This means that by querying the signature of m and m' , the adversary would obtain two pairs message-signature which both used the same nonce r : in the case of ESIGN, this discloses the secret factorization of the modulus, and in the case of DSA, this clearly discloses the secret key. This means that [27] did not give the right assumption: instead of a one-way function with uniform output, one should consider a PRF such as in KW3.

4.2 Robustness to Collisions

We now look at the security if the full-domain hash h has defects, namely collisions. We show that any hash collision suffices to disclose the master key in the ID-based cryptosystem of Boneh *et al.* [13], and the secret key in the Rabin-Williams signature scheme for which Bernstein proved tight security [10], because of the way derandomization is performed. But we also show that these attacks can be prevented by slightly modifying the derandomization.

The idea is, of course, to obtain two random preimages of a known element $m \in \mathcal{M}$. Doing so for the trapdoor one-way function f of Rabin and PRW discloses the secret factorization of the modulus with probability $1/2$. Let $h : \{0, 1\}^* \rightarrow \mathcal{M}$ be the random-oracle instantiation, to be paired with derandomized versions of Rabin or PRW. We have the following chosen-message key-recovery attack:

- Assume that the attacker is able to generate a collision (M_0, M_1) on h . Then $H(M_0) = H(M_1)$ with $M_0 \neq M_1$.
- The attacker queries the signing oracle on M_0 and M_1 , and obtains the signature s_0 and s_1 .
- Depending on how the derandomization is performed, we claim that s_0 and s_1 will be two random preimages of the same element $h(M_0) = h(M_1) \in \mathbb{Z}_N$, in which case it is easy to obtain the factorization of N with probability $1/2$. This is true in either of the following cases:
 - If one follows the informal method of Bernstein [10]: since M_0 and M_1 are different, the signer is not required to output the same preimage, so each preimage will be chosen at random.
 - If one follows KW1 [32], because $h(M_0) = h(M_1)$ is independent from $h'(K||M_0) = h'(K||M_1)$ where K denotes the secret key and h' is another random oracle.
 - If one follows KW3 [32] as in the ID-based cryptosystem [13], because $h(M_0) = h(M_1)$ is independent from $F_K(M_0) = F_K(M_1)$.

But if one follows KW2 [32], the same preimage will be output, and the attack will fail. An alternative method to prevent the attack is to use the following variant of KW3: $r = F_K(h(m))$ so that collisions on h gives collisions on r . For both variants, the previous key-recovery attacks no longer work, but that does not mean that the schemes are immune to collisions: any hash collision gives rise to an existential forgery. The interest of KW2 and the KW3 variant is that they both decrease (but not remove) the security impact of collisions.

Independently of the random-oracle instantiation and the choice of the derandomization, these weaknesses also appear if one considers fault attacks. Indeed, a similar attack works if the adversary is able to submit twice the same message, and perturbate the calculation of the nonce, using fault attacks.

By contrast, the deterministic version DRW-FDH (and the one implemented in IEEE P1363 [28]) is immune to such attacks. It might help to see a concrete example. Assume that we plug the compression function of MD5 into the CDMP random-oracle construction [18] (see Sect. 2.4), and that we use this instantiation as a full-domain hash for DRW-FDH and derandomized PRW-FDH. Then, because of the indistinguishability framework [41], both signature schemes become provably secure under the factoring assumption, in the ideal cipher model (with respect to the MD5 block cipher) or in the random oracle model (with respect to the MD5 compression function). But in practice, there is an instant chosen-message key-recovery attack on the PRW scheme, which fails on the DRW one.

4.3 Robustness to malleability

The previous key-recovery attack can be adapted to malleability variants of collisions on the hash function. To simplify, consider first the case of derandomized PRW-FDH: a similar attack works for derandomized Rabin-FDH. Assume that the attacker is able to generate a pair (M_0, M_1) of distinct messages such that:

$$4h(M_0) \equiv h(M_1) \pmod{N}. \quad (1)$$

From Sect. 2, we know that this is easy if ever h is VSH [15] using the same modulus N , even though it might be hard to find collisions on VSH, but we stress that it was never suggested to use VSH for Rabin/Rabin-Williams that way: we give this example to show that solving (1) is not necessarily harder than finding collisions. And we saw that solving (1) is also easy if h is BR93 [7]: select any M_0 then apply the preimage attack to find M_1 satisfying (1). Again, the attacker queries the signing oracle on M_0 and M_1 , which gives rise to tweaked square roots $(e_i, f_i, s_i) \in \{-1, 1\} \times \{1, 2\} \times \{0, \dots, N-1\}$ of $h(M_i)$. Note though that there is a one-to-one correspondance between the four tweaked square roots of $h(M_0)$ and the four tweaked square roots of $h(M_1)$, thanks to (1). More precisely, if (e, f, s) is a tweaked square root of $h(M_0)$, then $(e, f, 2s \bmod N)$ is a tweaked square root of $h(M_1)$. This implies that $(e_0, f_0, 2s_0 \bmod N)$ and (e_1, f_1, s_1) are two “independent” random tweaked square roots of $h(M_1)$, which means that one can factor N with probability $1/2$.

Obviously, this attack is independent of the way derandomization is performed, and can be adapted to other malleability properties. For instance, similar attacks apply if one is able to find a pair (M_0, M_1) of distinct messages such that $h(M_0) \equiv -h(M_1) \pmod{N}$, or $k^2h(M_0) \equiv h(M_1) \pmod{N}$ for some known $k \in \mathbb{Z}_N^\times$.

Furthermore, as opposed to collision attacks, the previous attack can be adapted to DRW-FDH. Starting again from (1), let $(e_i, f_i, s_i) \in \{-1, 1\} \times \{1, 2\} \times \{0, \dots, N-1\}$ be the principal

tweaked square root of $h(M_i)$. Because 4 is a square mod p and q , (1) implies that $e_0 = e_1$ and $f_0 = f_1$. Since $(e_0, f_0, 2s_0 \bmod N)$ is a tweaked square root of $h(M_1)$, we have $4s_0^2 \equiv s_1^2 \pmod{N}$, and therefore $s_1 \times (2s_0)^{-1} \bmod N$ is a square root of 1 mod N . But it must be a non-trivial square root because it has different Legendre symbols mod p and q : indeed, both s_0 and s_1 are squares mod N , while $\left(\frac{2}{p}\right) = 1$ and $\left(\frac{2}{q}\right) = -1$. Hence, this discloses the factorization of N . This attack can be generalized if congruence (1) is replaced by $k^2 h(M_0) \equiv h(M_1) \pmod{N}$ for any known $k \in \mathbb{Z}_N^\times$ such that $\left(\frac{k}{p}\right) \neq \left(\frac{k}{q}\right)$, which is slightly more restrictive than the PRW case.

There are similar attacks for the Rabin case.

4.4 Robustness to preimages

The previous attacks also show that both derandomized PRW-FDH and DRW-FDH become strongly insecure if the full-domain hash function h is not one-way, like BR93 [7]. Alternatively, one can simply select $(e, f, s) \in \{-1, 1\} \times \{1, 2\} \times \{0, \dots, N-1\}$ uniformly at random, and compute $m = efs^2 \pmod{N}$. By inverting h , one obtains a message M such that $m = h(M)$. Finally, by signing the message M with either DRW-FDH or derandomized PRW-FDH, one will obtain another tweaked square root of m (principal or not), which will disclose the factorization of N with probability at least 1/2 because (e, f, s) is a random tweaked square root. A similar attack works for the Rabin case.

5 Compared Robustness of Signatures

5.1 RSA Signatures

The PKCS#1 v2.1 standard [51] uses RSA-PSS since Sept. 1999 (or more precisely, the variant RSA-EMSA-PSS [29] of RSA-PSS), and it has been reported that one of the main reasons why RSA-PSS was selected over RSA-FDH was the tightness of the security proof. If tightness was the main factor, one might now be tempted to select RSA-KW over RSA-PSS, because the salt in RSA-KW is reduced to one bit (which can be deterministically derived from the secret key and the message). However, by comparing the robustness of RSA signatures with respect to potential defects in the random-oracle instantiation, a different picture emerges.

Robustness to collisions. Because RSA-FDH and RSA-EMSA-PSS are hash-and-sign schemes, they do not tolerate collisions: any collision obviously leads to a chosen-message existential forgery. Similarly, any collision leads to a chosen-message existential forgery on RSA-KW, with probability 1/2 because of the one-bit salt. One may think that the probabilistic schemes RSA-PFDH and RSA-PSS are more robust. In this direction, Numayama *et al.* [45] showed that RSA-PFDH tolerates collisions in a weakened ROM, but their model does not take into account MD-iterated hash functions. We observe that if h is a MD-iterated hash function, then any collision in h with the same number of blocks gives rise to a chosen-message existential forgery on RSA-PFDH and RSA-PSS. This is because RSA-PFDH and RSA-PSS both use $h(m||r)$. And if $h(m_1) = h(m_2)$ where m_1 and m_2 have the same number of blocks, then $h(m_1||r) = h(m_2||r)$ for any r . This implies that for both RSA-PFDH and RSA-PSS, any signature of m_1 is also valid for m_2 . It can be noted that if RSA-PFDH and RSA-PSS had used $h(r||m)$ instead of $h(m||r)$, then the ROM security proofs would remain valid, but the previous attack would fail.

Robustness to preimages. It is easy to prove that if the full-domain hash is not one-way, then there are chosen-message universal forgery attacks on RSA-FDH, RSA-PFDH and RSA-KW. On the other hand, preimages in h do not seem to provide stronger attacks than chosen-message existential forgeries on RSA-PSS.

Conclusion. While RSA-KW has a much better security reduction than RSA-FDH, there are essentially the same attacks on both RSA-KW and RSA-FDH as soon as there are defects in the full-domain hash. On the other hand, RSA-PSS with a large salt seems more robust than all other paddings, especially if one uses $h(r||m)$ instead of $h(m||r)$. This differs from the conclusion of [38], where it was argued that RSA-FDH was the best method known to sign with RSA.

5.2 Rabin and Rabin-Williams

Based on Sect. 4, the advantages of Rabin over Rabin-Williams are unclear from a security point of view. There are two benefits with Rabin-Williams: one can select $\mathcal{M} = \mathbb{Z}_N$, and one can use a 1-to-1 mapping instead of a 4-to-1 mapping. This 1-to-1 mapping avoids collision attacks or the fault attacks on derandomization: This suggests that DRW is preferable to both PRW and Rabin. And for the same reason as for RSA, a PSS padding with large salt seems more robust than all other paddings. Furthermore, when using Rabin or PRW, the size of the salt is extremely important to avoid key-recovery replay attacks. By submitting many times the same message, an adversary would obtain two random preimages of the same element, hence the factorization. We illustrate this phenomenon in Appendix B, with a converse to the security proof of [9].

Acknowledgements. We thank Mihir Bellare, Dan Boneh, Charles Bouillaguet, Louis Granboulan, Eike Kiltz, Kaoru Kurosawa, Alfred Menezes, Martijn Stam and the anonymous reviewers for their comments. Part of this work was supported by the Commission of the European Communities through the ICT program under contract ICT-2007-216676 ECRYPT II.

A Rabin Signatures in the Boneh-Gentry-Hamburg ID-based scheme

Boneh *et al.* [13] recently presented an identity-based encryption without pairings, which uses Rabin signatures in the key generation process as follows.

Let $N = pq$ be an RSA modulus, and denote by $J(N)$ the subgroup of \mathbb{Z}_N^\times of elements with Jacobi symbol equal to 1. Denote by $\text{QR}(N)$ the subgroup of $J(N)$ of quadratic residues. The integer ℓ defines the bit-length of messages, and \mathcal{ID} denotes the set of identities. The public parameters are (N, u, H) where u is chosen with uniform distribution in $J(N) \setminus \text{QR}(N)$, and $H : \mathcal{ID} \times [1, \ell] \rightarrow J(N)$ is a hash function. The master key is the factorization of N and a random key K for a pseudorandom function $F_K : \mathcal{ID} \times [1, \ell] \rightarrow \{0, 1, 2, 3\}$.

The key generation is as follows. Given as input the master key and an identity ID, for each $j = 1, \dots, \ell$ do: $R_j \leftarrow H(\text{ID}, j) \in J(N)$ and $w \leftarrow F_K(\text{ID}, j)$; let $a \in \{0, 1\}$ be such that $u^a R_j \in \text{QR}(N)$; set $r_j \leftarrow z_w$ where $\{z_0, \dots, z_3\}$ are the four square roots of $u^a R_j$ in \mathbb{Z}_N . The decryption key corresponding to ID is (r_1, \dots, r_ℓ) together with the public parameters. The PRF F ensures that the key generator always outputs the same square roots for a given ID, but an adversary cannot tell ahead of time which of the four square roots will be output.

B A Converse to the Security Proof of P_{Rab}

We consider P_{Rab} [9], which is a PSS version of Rabin signatures by Bellare and Rogaway: we give a converse to the security proof of [9], namely that if the parameters are selected in such a way that the security reduction is not tight, then there is an efficient key-recovery attack on P_{Rab}. And we show how to improve this attack if there are defects in the oracle instantiation. Thus, in this case, it is extremely important to select parameters carefully so that the security proof is tight.

Let $N = pq$ be the Rabin public key of bit-length $k \geq k_0 + k_1$. Two hash functions $h : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and $g : \{0, 1\}^{k_1} \rightarrow \{0, 1\}^{k-k_1}$ are used. We let g_1 be the function which on input $w \in \{0, 1\}^{k_1}$ returns the first k_0 bits of $g(w)$, and let g_2 be the function which on input $w \in \{0, 1\}^{k_1}$ returns the remaining $k - k_0 - k_1$ bits of $g(w)$. To sign a message $M \in \{0, 1\}^*$, one first repeats the following process until y is a quadratic residue mod N :

1. Choose $r \in_R \{0, 1\}^{k_0}$ and let $w = h(M||r)$ and $r^* = g_1(w) \oplus r$;
2. Let $y = w||r^*||g_2(w)$.

Then the signature s is chosen uniformly at random among the four distinct square roots of $y \in \mathbb{Z}_N^\times$. Bellare and Rogaway [9] showed the security of P_{Rab} in the ROM, under the hardness of factoring N , and the tightness of their security proof was similar to that of their proof for RSA-PSS. More precisely, if ε and ε' denote the security proof advantages for respectively P_{Rab} and factoring, then:

$$\varepsilon = 2\varepsilon' + [4(q_{sig} + q_{hash})^2 + 2] \cdot (2^{-k_0} + 2^{-k_1}), \quad (2)$$

where q_{sig} and q_{hash} denote respectively the numbers of signature queries and hash queries. Under appropriate choices of k_0 and k_1 , the right-hand term becomes negligible with respect to ε' , in which case the security proof is said to be tight.

We now show that it is extremely important to select k_0 to make the security proof tight (see [37] for a discussion on the importance of tightness). Consider first the following elementary known-message key-recovery attack:

- Collect n signatures s_1, \dots, s_n of random messages M_1, \dots, M_n . Denote by r_i the (unknown) random number used by the signer to generate s_i .
- Since the output of h is $\{0, 1\}^{k_1}$ and $r \in_R \{0, 1\}^{k_0}$, by the birthday paradox, as soon as $n \geq \Omega(2^{(k_0+k_1)/2})$, there should be $i \neq j$ such that $h(M_i||r_i) = h(M_j||r_j)$ and $r_i = r_j$. Such a collision can be detected by looking at all $s_i^2 \pmod N$.
- Then s_i and s_j are two random square roots of the same $y \in \mathbb{Z}_N^\times$. This is because the signing process of P_{Rab} does not check whether the same y has been “signed” before. Hence $\gcd(s_i - s_j, N)$ is a non-trivial factor of N with probability $1/2$.

This gives a generic attack whose cost is essentially $2^{(k_0+k_1)/2}$ signatures of random messages.

This attack can be improved in the chosen-message setting. If we submit $\Theta(2^{k_0/2})$ times the same message M to the signing oracle, we will obtain a collision on the r_i 's, and the attack still applies. This chosen-message attack costs $O(2^{k_0/2})$ signature queries (on the same message), which proves that the security reduction of [9] summarized by Eq. (2) is essentially tight in a different sense: as soon as $q_{sig}^2 \geq \Omega(2^{k_0})$, an adversary can actually recover the factorization of N . In particular, if one wrongly applies Coron's security result [16] for RSA-PSS to P_{Rab}, one might believe that $k_0 = 30$ is a good choice, in which case there is an attack requiring only 2^{15} signature queries on the same message.

Note that if h is a Merkle-Damgård iterated hash function, the adversary does not even have to submit the same message to the signing oracle: he may even restrict to distinct messages.

Indeed, by using Joux’s multicollision technique [30], the adversary can generate n distinct messages M_1, \dots, M_n such that $h(M_i) = h(M_j)$ and all the M_i ’s have the same number of blocks. Then, by definition of Merkle-Damgård, if $r_i = r_j$, we have $h(M_i || r_i) = h(M_j || r_j)$, so we can still apply the attack. The cost of generating n multicollisions is $O(2^{k_1/2} \log n)$ (see [30]), which can become as low as $O(\log n)$ if it is easy to generate chosen-IV collisions (such as is the case for MD5 [59]). However, we still need to make $O(2^{k_0/2})$ signature queries. The attack becomes particularly efficient if h is MD5, and the size k_0 of the random salt is small.

C Security Robustness of the GPV Lattice-based Signature Scheme

C.1 Description of GPV

Gentry, Peikert and Vaikuntanathan (GPV) [23] recently proposed a provably-secure variant of the GGH lattice-based signature scheme [24]. Nguyen and Regev [42] obtained a polynomial-time key-recovery known-message attack against GGH signatures, but their attack exploited the deterministic choices made by the signature algorithm. In contrast, GPV is probabilistic, based on Klein’s randomized variant [34] of Babai’s nearest plane algorithm [2]. The tight security proof of [23] is strikingly similar to Bernstein’s tight security proof [10] for the TRW scheme. This is because both signature schemes are actually based on the same primitive: a trapdoor collision-resistant hash function with preimage sampling. In [10], collision resistance is based on factoring, while in [23], it is based on the hardness of finding very short vectors in certain lattices.

We now recall the GPV signature scheme [23], using as less lattice terminology as possible: much more details can be found in [23]. There is a main integer parameter n , and two additional integers m and q such that $q = \text{poly}(n)$ and $m = \Theta(n \log q)$. The signer selects a matrix $A \in \mathbb{Z}_q^{n \times m}$ in such a way that: the distribution of A is statistically close to the uniform distribution, and the signer knows a secret set S of m very short linearly independent vectors in the m -dimensional lattice L formed by all vectors $\mathbf{x} \in \mathbb{Z}^m$ such that $A\mathbf{x} \equiv 0 \pmod{q}$.

There is a full-domain hash function $h : \{0, 1\}^* \rightarrow \mathbb{Z}_q^n$. To sign a message $M \in \{0, 1\}^*$, the signer applies the following process:

1. Compute the hash $\mathbf{h} = h(M) \in \mathbb{Z}_q^n$.
2. Thanks to S , the signer is able to compute in a probabilistic way (with overwhelming probability) a vector $\mathbf{e} \in \mathbb{Z}^m$ such that $A\mathbf{e} \equiv \mathbf{h} \pmod{q}$ and \mathbf{e} is very short, namely $\|\mathbf{e}\| \leq s\sqrt{m}$ where the parameter $s \in \mathbb{R}^+$ is related to S and the lattice L . There are actually many possible \mathbf{e} ’s, and the one selected by the signer has a Gaussian distribution, which is the main reason why one can obtain a security proof in the ROM. Again, like in DRW, it is assumed that if ever M has already been signed, then the same \mathbf{e} is chosen.
3. Output $\mathbf{e} \in \mathbb{Z}^m$ as the signature of M .

To check that $\mathbf{e} \in \mathbb{Z}^m$ is a valid signature of M , one verifies that $A\mathbf{e} \equiv h(M) \pmod{q}$ and that \mathbf{e} is sufficiently short, namely $\|\mathbf{e}\| \leq s\sqrt{m}$.

The security proof of [23] in the ROM is with respect to the following hard lattice problem: finding non-zero lattice vectors $\mathbf{x} \in L$ such that $\|\mathbf{x}\| \leq 2s\sqrt{m}$. In the terminology of Sect. 3.1, GPV roughly corresponds to the following situation: $\mathcal{M} = \mathbb{Z}_q^n$, $\mathcal{S} = \{\mathbf{e} \in \mathbb{Z}_q^n, \|\mathbf{e}\| \leq s\sqrt{m}\}$ and $f(\mathbf{e}) = A\mathbf{e} \pmod{q}$. But there are differences: the preimage sampling follows a Gaussian distribution, rather than the uniform distribution over \mathcal{S} .

We now present chosen-message attacks on the GPV signature scheme, when the underlying

hash function is weak. We first present attacks that break the underlying computational assumption, by finding very short vectors in the lattice L , then we discuss why such attacks might lead to forgery attacks and key-recovery attacks. In the Rabin-Williams case, breaking the underlying computational assumption (that is, factoring) directly leads to key recovery. This may not be the case with GPV [23].

C.2 Robustness with respect to collisions

Let $h : \{0, 1\}^* \rightarrow \mathbb{Z}_q^n$ be the random-oracle instantiation. Assume that the attacker is able to generate a collision (M_0, M_1) on h . The attacker queries the signing oracle on M_0 and M_1 , and obtains two signatures $\mathbf{e}_0, \mathbf{e}_1 \in \mathbb{Z}^m$. Then:

$$A\mathbf{e}_0 \equiv A\mathbf{e}_1 \equiv h(M_0) \equiv h(M_1) \pmod{q},$$

which implies that $A(\mathbf{e}_0 - \mathbf{e}_1) \equiv 0 \pmod{q}$. Thus, $\mathbf{e}_0 - \mathbf{e}_1$ belongs to the m -dimensional lattice L , of norm $\leq 2s\sqrt{m}$, but could it be zero? In the exact description of GPV scheme [23, Section 5], it is written that if a message M has already been signed, the same signature should be output, but similarly to TRW, it does not say that this should also be the case if the digest $h(M)$ has already been “signed”. Thus, we may assume that both \mathbf{e}_0 and \mathbf{e}_1 have been generated independently, as $M_0 \neq M_1$: since the distribution of both \mathbf{e}_0 and \mathbf{e}_1 is Gaussian, the difference $\mathbf{e}_0 - \mathbf{e}_1$ is nonzero with overwhelming probability. In other words, using only two signature queries, we are likely to have obtained a non-zero lattice vector of L of norm $\leq 2s\sqrt{m}$, which is exactly the underlying hard problem of the GPV security proof.

C.3 Robustness with respect to malleability

Similarly to the Rabin-Williams case, the previous attack can be adapted to malleability properties of the hash function. Assume indeed that the attacker is able to find two distinct messages M_0 and M_1 such that:

$$h(M_0) \equiv h(M_1) + A\mathbf{x} \pmod{q},$$

where $\mathbf{x} \in \mathbb{Z}^m$ is non-zero and very short, say $\|\mathbf{x}\| \ll s\sqrt{m}$. Again, let $\mathbf{e}_0, \mathbf{e}_1 \in \mathbb{Z}^m$ be signatures corresponding to M_0 and M_1 . Then:

$$A\mathbf{e}_0 \equiv A\mathbf{e}_1 + A\mathbf{x} \pmod{q},$$

which implies that $\mathbf{e}_0 - \mathbf{e}_1 - \mathbf{x}$ is a very short vector of L , whose norm might be nearly as short as $2s\sqrt{m}$.

In a similar vein, assume that the attacker is able to generate two distinct pairs of distinct messages (M_0, M_1) and (M'_0, M'_1) such that:

$$h(M_0) - h(M_1) \equiv h(M'_0) - h(M'_1) \pmod{q}.$$

In Section 2, we showed that this was easy with SWIFFT [40] and one can easily build instantiations of h based on MD5/SHA-1 with similar properties, thanks to near-collisions techniques. Then, if we denote by \mathbf{e}_i and \mathbf{e}'_i the respective signatures of M_i and M'_i , we obtain that $\mathbf{e}_0 - \mathbf{e}_1 - \mathbf{e}'_0 + \mathbf{e}'_1$ is a short vector of L , of norm $\leq 4s\sqrt{m}$.

C.4 Forgery and Key recovery

The previous attacks made it possible to find very short nonzero vectors in the lattice L , possibly sufficiently short to break the computational assumption of GPV. However, this alone may not be sufficient to recover the secret key S , nor to generate new signatures. We now explain why the attacker may still be hopeful.

First of all, the attacker may try to use his collection of newly obtained short lattice vectors to compute even shorter lattice vectors. This is exactly the principle of sieve algorithms [1, 43], which can be heuristically improved by performing strong lattice reduction of low-dimensional sublattices spanned by subsets of vectors. The ability of the attacker to forge signatures depends on his ability to approximate the closest vector problem (CVP) in the lattice L . This approximation is usually obtained by Babai's algorithm [2] or its variants [34], which in turn depends on how short and how orthogonal are the lattice vectors known by the attacker. However, the performances of Babai's algorithm can be improved by exhaustive search: more precisely, Babai's algorithm is based on a repeated use of a one-dimensional CVP subroutine, which can heuristically be improved by using higher-dimensional CVP subroutines. The success of these methods will depend on the exact choice of GPV parameters. Since there is currently no concrete proposal for such parameters, it is difficult to assess the difficulty of forgery or key recovery, based on the previous attacks.

References

- [1] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC '01*, pages 601–610. ACM, 2001.
- [2] L. Babai. On Lovász lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.
- [3] M. Bellare, 2009. Private communication.
- [4] M. Bellare, A. Boldyreva, and A. Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In *EUROCRYPT '04*, volume 3027 of *LNCS*. Springer, 2004.
- [5] M. Bellare and D. Micciancio. A new paradigm for collision-free hashing: Incrementality at reduced cost. In *EUROCRYPT '97*, volume 1233 of *LNCS*, pages 163–192. Springer, 1997.
- [6] M. Bellare and T. Ristenpart. Multi-property-preserving hash domain extension and the EMD transform. In *ASIACRYPT '06*, volume 4284 of *LNCS*. Springer, 2006.
- [7] M. Bellare and P. Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *CCS '93*, pages 62–73. ACM Press, 1993.
- [8] M. Bellare and P. Rogaway. Optimal asymmetric encryption. In *Proc. of Eurocrypt '94*, volume 950 of *LNCS*, pages 92–111. IACR, Springer, 1995.
- [9] M. Bellare and P. Rogaway. The exact security of digital signatures - how to sign with RSA and Rabin. In *EUROCRYPT '96*, volume 1070 of *LNCS*. Springer, 1996.
- [10] D. J. Bernstein. Proving tight security for Rabin-Williams signatures. In *EUROCRYPT '08*, *LNCS*. Springer, 2008.

- [11] J. Black, P. Rogaway, and T. Shrimpton. Black-box analysis of the block-cipher-based hash-function constructions from PGV. In *CRYPTO '02*, volume 2442 of *LNCS*. Springer, 2002.
- [12] D. Boneh and X. Boyen. Efficient selective-ID secure identity based encryption without random oracles. In *EUROCRYPT '04*, volume 3027 of *LNCS*, pages 223–238. Springer, 2004.
- [13] D. Boneh, C. Gentry, and M. Hamburg. Space-efficient identity based encryption without pairings. In *FOCS '07*, pages 647–657. IEEE Computer Society, 2007.
- [14] R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594 (electronic), 2004. Preliminary version at STOC '98.
- [15] S. Contini, A. K. Lenstra, and R. Steinfeld. VSH, an efficient and provable collision-resistant hash function. In *EUROCRYPT '06*, volume 4004 of *LNCS*, pages 165–182. Springer, 2006.
- [16] J.-S. Coron. Optimal security proofs for PSS and other signature schemes. In *EUROCRYPT '02*, volume 2332 of *LNCS*, pages 272–287. Springer, 2002.
- [17] J.-S. Coron. Security proof for partial-domain hash signature schemes. In *CRYPTO '02*, volume 2442 of *LNCS*, pages 613–626. Springer, 2002.
- [18] J.-S. Coron, Y. Dodis, C. Malinaud, and P. Puniya. Merkle-Damgård revisited: How to construct a hash function. In *CRYPTO '05*, volume 3621 of *LNCS*. Springer, 2005.
- [19] Y. Dodis, R. Oliveira, and K. Pietrzak. On the generic insecurity of the full domain hash. In *CRYPTO '05*, volume 3621 of *LNCS*, pages 449–466. Springer, 2005.
- [20] Y. Dodis and L. Reyzin. On the power of claw-free permutations. In *SCN '02*, volume 2576 of *LNCS*, pages 55–73. Springer, 2003.
- [21] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO '86*, volume 263 of *LNCS*, pages 186–194. Springer, 1987.
- [22] E. Fujisaki, T. Kobayashi, H. Morita, H. Oguro, T. Okamoto, and S. Okazaki. Esign-d specification. Submission to the NESSIE European Project, available on the NESSIE webpage, 2002.
- [23] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. Cryptology ePrint Archive, Report 2007/432, 2007. Also in Proc. STOC '08.
- [24] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *CRYPTO '97*, volume 1294 of *LNCS*. Springer, 1997.
- [25] S. Goldwasser and Y. T. Kalai. On the (in)security of the Fiat-Shamir paradigm. In *FOCS '03*. IEEE Computer Society, 2003.
- [26] 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. Special issue on cryptography.
- [27] L. Granboulan. How to repair esign. In *SCN '02*, volume 2576 of *LNCS*, pages 234–240. Springer, 2003.

- [28] IEEE. P1363: Standard specifications for public-key cryptography. Available at <http://grouper.ieee.org/groups/1363/>.
- [29] J. Jonsson. Security proofs for the RSA-PSS signature scheme and its variants. In *Proc. 2nd NESSIE Workshop*, 2001. Full version available as Report 2001/053 of the Cryptology ePrint Archive.
- [30] A. Joux. Multicollisions in iterated hash functions. application to cascaded constructions. In *CRYPTO '04*, volume 3152 of *LNCS*, pages 306–316. Springer, 2004.
- [31] A. Joux and T. Peyrin. Hash functions and the (amplified) boomerang attack. In *CRYPTO '07*, volume 4622 of *LNCS*, pages 244–263. Springer, 2007.
- [32] J. Katz and N. Wang. Efficiency improvements for signature schemes with tight security reductions. In *CCS '03*. ACM, 2003.
- [33] E. Kiltz and K. Pietrzak. On the security of padding-based encryption schemes (or: Why we cannot prove OAEP secure in the standard model). In *EUROCRYPT '09*, LNCS. Springer, 2009.
- [34] P. N. Klein. Finding the closest lattice vector when it's unusually close. In *Proc. SODA*, pages 937–941, 2000.
- [35] V. Klima. Tunnels in Hash Functions: MD5 Collisions Within a Minute. Cryptology ePrint Archive, Report 2006/105, 2006. <http://eprint.iacr.org/>.
- [36] T. Kobayashi and E. Fujisaki. Security of esign-pss. *IEICE Transactions*, 90-A(7):1395–1405, 2007.
- [37] N. Kobitz and A. Menezes. Another look at "provable security". II. In *INDOCRYPT '06*, volume 4329 of *LNCS*. Springer, 2006. Also available as ePrint 2006/229.
- [38] N. Kobitz and A. J. Menezes. Another look at "provable security". *J. Cryptology*, 20(1):3–37, 2007.
- [39] G. Leurent and P. Q. Nguyen. How risky is the random-oracle model? In *CRYPTO '09*, LNCS. Springer, 2009.
- [40] V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE '08*, volume 5086 of *LNCS*. Springer, 2008.
- [41] U. M. Maurer, R. Renner, and C. Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *TCC '04*, volume 2951 of *LNCS*, pages 21–39. Springer, 2004.
- [42] P. Q. Nguyen and O. Regev. Learning a Parallelepiped: Cryptanalysis of GGH and NTRU Signatures. In *EUROCRYPT '06*, volume 4004 of *LNCS*. Springer, 2006.
- [43] P. Q. Nguyen and T. Vidick. Sieve algorithms for the shortest vector problem are practical. *J. of Mathematical Cryptology*, 2(2), 2008.

- [44] J. B. Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In *CRYPTO '02*, volume 2442 of *LNCS*, pages 111–126. Springer, 2002.
- [45] A. Numayama, T. Isshiki, and K. Tanaka. Security of digital signature schemes in weakened random oracle models. In *PKC '08*, volume 4939 of *LNCS*, pages 268–287. Springer, 2008.
- [46] T. Okamoto. A fast signature scheme based on congruential polynomial operations. *IEEE Transactions on Information Theory*, 36(1):47–53, 1990.
- [47] T. Okamoto, E. Fujisaki, and H. Morita. TSH-ESIGN: Efficient digital signature scheme using trisection size hash. Submission to IEEE P1363a, 1998.
- [48] T. Okamoto and J. Stern. Almost uniform density of power residues and the provable security of esign. In *ASIACRYPT '03*, volume 2894 of *LNCS*, pages 287–301. Springer, 2003.
- [49] P. Paillier and D. Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In *ASIACRYPT '05*, volume 3788 of *LNCS*, pages 1–20. Springer, 2005.
- [50] M. Rabin. Digital signatures and public key functions as intractable as factorization. Technical report, MIT Laboratory for Computer Science, 1979. Report TR-212.
- [51] RSA Laboratories. PKCS #1 v2.1: RSA cryptography standard. June 14, 2002.
- [52] M.-J. O. Saarinen. Linearization attacks against syndrome based hashes. In *INDOCRYPT 2007*, volume 4859 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2007.
- [53] V. Shoup. Using hash functions as a hedge against chosen ciphertext attack. In *EUROCRYPT '00*, *LNCS*, pages 275–288. Springer, 2000.
- [54] A. Sotirov, M. Stevens, J. Appelbaum, A. Lenstra, D. Molnar, D. A. Osvik, and B. de Weger. Short chosen-prefix collisions for MD5 and the creation of a rogue CA certificate. In *CRYPTO '09*, *LNCS*. Springer, 2009.
- [55] J. Stern, D. Pointcheval, J. Malone-Lee, and N. P. Smart. Flaws in applying proof methodologies to signature schemes. In *CRYPTO '02*, volume 2442 of *LNCS*. Springer, 2002.
- [56] M. Stevens, A. K. Lenstra, and B. de Weger. Chosen-prefix collisions for MD5 and colliding X.509 certificates for different identities. In *EUROCRYPT '07*, volume 4515 of *LNCS*. Springer, 2007.
- [57] D. Wagner. A generalized birthday problem. In *CRYPTO '02*, volume 2442 of *LNCS*, pages 288–303. Springer, 2002.
- [58] X. Wang, Y. L. Yin, and H. Yu. Finding collisions in the full SHA-1. In *Advances in cryptology— Proc. CRYPTO '05*, volume 3621 of *Lecture Notes in Comput. Sci.*, pages 17–36. Springer, Berlin, 2005.
- [59] X. Wang and H. Yu. How to break MD5 and other hash functions. In *Advances in Cryptology – Proc. EUROCRYPT '05*, volume 3494 of *LNCS*, pages 19–35. Springer, 2005.

- [60] H. C. Williams. A modification of the RSA public-key encryption procedure. *IEEE Trans. Inform. Theory*, 26(6):726–729, 1980.
- [61] R. S. Winternitz. A secure one-way hash function built from DES. In *Proc. IEEE Symposium on Security and Privacy*, pages 88–90, 1984.