

Strengthening Digital Signatures Via Randomized Hashing

Shai Halevi and Hugo Krawczyk

IBM T.J. Watson Research Center, Yorktown Heights, New York 10598
shaih@alum.mit.edu, hugo@ee.technion.ac.il

Abstract. We propose randomized hashing as a mode of operation for cryptographic hash functions intended for use with standard digital signatures and without necessitating of any changes in the internals of the underlying hash function (e.g., the SHA family) or in the signature algorithms (e.g., RSA or DSA). The goal is to free practical digital signature schemes from their current reliance on strong collision resistance by basing the security of these schemes on significantly weaker properties of the underlying hash function, thus providing a *safety net* in case the (current or future) hash functions in use turn out to be less resilient to collision search than initially thought.

We design a specific mode of operation that takes into account engineering considerations (such as simplicity, efficiency and compatibility with existing implementations) as well as analytical soundness. Specifically, the scheme consists of a regular use of the hash function with randomization applied only to the message before it is input to the hash function. We formally show the sufficiency of weaker than collision-resistance assumptions for proving the security of the scheme.

1 Introduction

Recent cryptanalytical advances in the area of collision-resistant hash functions (CRHF) [8,5,15,6,16,29,30,31,32], especially the attacks against MD5 and SHA-1, have shaken our confidence in the security of existing hash functions as well as in our ability to design secure CRHF. These attacks remind us that cryptography is founded on heuristic constructions whose security may be endangered unexpectedly, and highlight the importance of designing mechanisms that require as little as possible from their basic cryptographic building blocks. In particular, they indicate that one should move away (to the extent possible) from schemes whose security fundamentally depends on full collision resistance.

The most prominent example of schemes that are endangered by these attacks are digital signatures where collisions in the hash function directly translate into forgeries. The goal of this work is to free existing (standardized) signature schemes from their dependence on full collision-resistance of the underlying hash function, while making as small as possible modifications to the signature and hashing algorithms. Specifically, we propose a *randomized mode of operation* for hash functions such that (i) no change to the underlying hash functions

(e.g., the SHA family), the signature algorithms (e.g., RSA, DSA), or their implementations is required; (ii) changes to the signing process are the minimal required by any randomization scheme, namely, the choice of a short random string by the signer and transporting this string as part of, or in addition to, the existing signature; and (iii) *security of the resultant signature scheme does not depend on the resistance of the hash function to off-line collision attacks.*

Armoring signature schemes with this mode of operation provides a *safety net* for the security of digital signatures in the case that the collision resistance of the underlying hash function is broken or weakened. At the same time, by following important engineering considerations, as in points (i) and (ii) above, one facilitates the adoption of the resultant schemes into practice. In particular, treating the hash function as a black box allows to preserve existing implementations of these functions or their future replacements. Moreover, the many other applications of hash functions (e.g., MAC, PRFs, commitments, fingerprinting, etc.) need not be changed or adapted to a new hash design; instead, these applications may use the hash function, or any “drop-in replacement”, exactly as they do now.

Collision resistance and the Merkle-Damgård (M-D) construction. Roughly, a (length-decreasing) function H is collision resistant if given the description of H , no feasible attacker can find two messages M, M' such that $H(M) = H(M')$, except with insignificant probability.¹ Contemporary constructions of (allegedly) collision-resistant hash functions follow the so called Merkle-Damgård (M-D) iterated construction [17,9]. Such constructions start with a *compression function* h that maps input pairs (c, m) into an output c' where c and c' are of fixed length n (e.g., $n = 160$) and m is of fixed length b (e.g., $b = 512$). Given such a compression function h and a fixed n -bit initial value (IV), denoted c_0 , a hash function H on arbitrary-length inputs is defined by iterating h as follows: On input message M , the message is broken into b -bit blocks $M = (m_1, m_2, \dots, m_L)$, then one computes $c_i \leftarrow h(c_{i-1}, m_i)$ for $i = 1, 2, \dots, L$ and finally the last n -bit value c_L is output as the result of $H(M)$.

Throughout this paper we typically denote the (iterated) hash function by H and its compression function by h . For an M-D function H , we also sometimes write $H^{c_0}(M)$ when we want to explicitly name the initial value c_0 . To handle non-suffix-free message spaces, the M-D construction appends the length at the end of the input and uses some padding rules. We will initially ignore these issues and assume that the input consists of an integral number of blocks and that it is taken from a suffix-free set; we will return to the issue of length appending in Section 5.

Target Collision Resistance. Signature schemes that do not depend on full collision resistance were constructed in the influential work of Naor and Yung [20] who introduced the notion of *universal one-way hash functions, or UOWHF*. Later, Bellare and Rogaway [3] renamed them to the more descriptive (and

¹ Formalizing collision resistance requires to consider H as a family of functions rather than as a single function; we ignore this technicality for the informal discussion here.

catchy) name of *target collision resistant* (TCR) hash functions, a term that we adopt here. Roughly, a family of hash functions $\{H_r\}_{r \in R}$ (for some set R) is target collision resistant if no efficient attacker A can win the following game, except with insignificant probability: A chooses a first message M , then receives a random value $r \in_R R$, and it needs to find a second message $M' \neq M$ such that $H_r(M') = H_r(M)$. The value r is called a hashing key, or a *salt*. In [3], TCR families were used to extend signature schemes on short messages into signature schemes on arbitrary messages: To sign a long message M , the signer chooses a fresh random salt r and then uses the underlying signature scheme to sign the pair $(r, H_r(M))$.

Enhanced target collision resistance. Standard signature schemes such as RSA and DSA are defined using a deterministic hash function H and designed to only sign the hash value $H(M)$. When moving to a TCR-based scheme one replaces $H(M)$ with $H_r(M)$ but also needs to sign the salt r . Unfortunately, certain signature schemes, such as DSA, do not accommodate the signing of r in addition to $H_r(M)$. Even in cases, such as RSA, where both r and $H_r(M)$ can be included under the signed block (RSA moduli are long enough to accommodate both values) signing r requires changing the message encoding standards (such as PKCS#1) which may be impractical in some settings. To solve these difficulties, we propose a strengthened mode of operation, that provides a stronger security guarantee than TCR. The new notion, that we call *enhanced TCR* (eTCR), is sufficiently strong to ensure security of the resultant signatures even if we only apply the underlying signature to $H_r(M)$ and *do not sign the salt* r . Specifically, the TCR attack scenario is modified so that after committing to M and receiving the salt r , the attacker can supply a second message M' and a second salt r' , and it is considered successful if $(r, M) \neq (r', M')$ but $H_r(M) = H_{r'}(M')$.

1.1 Our Randomization Schemes

The main contribution of this work is in presenting simple and efficient randomization schemes that use any Merkle-Damgård iterated hash function as-is (i.e., in a black-box fashion) and have short salts, and which we prove to be TCR and/or eTCR under weak assumptions on the underlying compression function. Specifically, we relate the security of our schemes to “second-preimage resistance” properties of the compression function. Recall that a compression function h is called *second preimage resistant* (SPR) if given a uniformly random pair (c, m) it is infeasible to find a second pair $(c', m') \neq (c, m)$ such that $h(c', m') = h(c, m)$.

The first scheme that we consider, denoted H_r , simply XOR’s a random b -bit block r to each message block before inputting it into the hash function (see below for a comparison of this scheme to one in [26]). That is,

$$H_r^c(m_1, \dots, m_L) \stackrel{\text{def}}{=} H^c(m_1 \oplus r, \dots, m_L \oplus r). \quad (1)$$

We show this scheme to be TCR under SPR-like assumption on the underlying compression function h (see below). On the other hand, the H_r construction is clearly not eTCR; in order to obtain an eTCR scheme we modify H_r by

prepending the salt block r to the message (in addition to xor-ing r to every message block as before), thus computing

$$\tilde{H}_r^c(M) \stackrel{\text{def}}{=} H_r^c(0|M) = H^c(r, m_1 \oplus r, \dots, m_L \oplus r) \quad (2)$$

Remarkably, for any M-D iterated hash function H we can prove that \tilde{H}_r is enhanced-TCR under *the same relaxed conditions on the compression function* that we use to prove that H_r is TCR.

These precise conditions are presented in Section 2. Roughly, we show two properties of the compression function h that are sufficient for H_r to be TCR (resp. for \tilde{H}_r to be eTCR). Both properties are related to the second-preimage resistance of h . One property, c-SPR, is rather natural but also rather strong, and is provided mostly as a toy problem for cryptanalysts to sink their teeth in. The other property, e-SPR, is less natural but is closer to SPR and reflects more accurately the “real hardness” of the constructions H_r and \tilde{H}_r . We also note that c-SPR is related to the “hierarchy of collision resistance” of Mironov [18], and that e-SPR is a weaker property than the L -order TCR property of Hong et al. [14].

We stress that ultimately, the question of whether or not a particular compression function (such as those in the SHA family) is e-SPR can only be addressed by cryptanalysts trying to attack this property. As we explain in Section 2, breaking the compression function in the sense of e-SPR is a much harder task than just finding collisions for it. Moreover, the randomized setting represents a *fundamental shift* in the attack scenario relative to the traditional use of deterministic hash functions. In the latter case, the attacker can find collisions via (heavy) *off-line* computation and later use these collisions in a signature forgery attack at any time and with any signer. In contrast, in the randomized setting, a meaningful collision is one that utilizes a random, unpredictable r *chosen by the signer anew with each signature* and revealed to the attacker only with the signature on the message M (in particular, only after the attacker committed to the message M).

Can the signer cheat? A possible objection to the randomized setting is that it may allow a signer to find two messages with the same hash value. This, however, represents no violation of signature security (c.f. the standard definition of Goldwasser, Micali, and Rivest [12]). Specifically, the potential ability of the signer to find collisions does not allow any other party to forge signatures thus protecting the signer against malicious parties. At the same time, recipients of signatures are also protected since the signer cannot disavow a signature by presenting a collision; the simple rule is, as in the case of human signatures, that the signer is liable for any message carrying his valid signature (even if the signer can show a collision between two messages). It is also worth noting that our schemes have the property that as long as the underlying hash function is collision resistant then even the signer cannot find collisions. More generally, the signature schemes resulting from our randomization modes are never weaker (or less functional) than the schemes that use a deterministic hash function,

not even when the underlying hash function is truly collision resistant; and, of course, our schemes are much more secure once the hash function ceases to be collision resistant.

Practical Considerations. As seen, our randomized hashing schemes require no change to the underlying (compression and iterated) hash function. These schemes can be used, as described before, with algorithms such as RSA and DSA, to provide for digital signatures that remain secure even in the presence of collisions for the underlying hash functions. For this, the value $H(M)$ currently signed by the standard schemes needs to be replaced with $H_r(M)$ or with $\tilde{H}_r(M)$. However, while in the case of H_r one needs to also sign the salt r , in the \tilde{H}_r case the latter is not needed. In this sense, \tilde{H}_r provides for a more flexible and practical scheme and hence it is the one that we suggest for adoption into practice. In particular, using \tilde{H}_r , the salt r may be transported by an application as part of (or in addition to) a message, or r can be included under the signature itself (under the signed value in RSA or by re-using the random $r = g^k$ component in DSA). We further discuss these issues in Section 5. A full specification of the \tilde{H}_r mode is included in [13].

Related Schemes. Bellare and Rogaway [3] explored the problem of constructing TCR hashing for long messages from TCR functions that work on short inputs, and showed that the M-D iteration of a TCR compression function *does not necessarily* result in a long-message TCR. Instead, they presented a few other constructions of long-message TCR hash functions from short-message TCR compression functions. Shoup [26] continued this line of research and offered the following elegant scheme that comes close to the original M-D iterated construction. The salt in Shoup's scheme consists of one b -bit message block r and a set S of $\log L$ chaining variables s_0, s_1, \dots , where L is the number of message blocks. The scheme consists of a regular M-D construction except that (i) every message block is xor-ed with r before it is input to the compression function h , and (ii) the intermediate values c_i output by each iteration of h are xor-ed with one of the s_i 's to form the chaining value input into the following application of h . That is, the M-D scheme is modified as follows: for $i = 1, 2, \dots, L$, $c_i \leftarrow h(c_{i-1} \oplus s_{j_i}, m_i \oplus r)$ where j_i chooses one of the elements in the set S . Shoup proved that if the compression function is SPR then the construction yields a TCR function on messages of arbitrary length.

Note that our H_r scheme (Eqn. (1)) is similar to Shoup's in that the input blocks are xor-ed with the random r . However, the two schemes differ substantially in that: (i) H_r uses the hash function as a black box while Shoup's scheme requires the change of the internals of H to accommodate the xor-ing of the s_i 's values with the partial results of the compression function; and (ii) in H_r the amount of randomness is fixed while in Shoup's it is logarithmic in the length of the message and hence non-constant; this requires significantly more random bits and results in longer signatures. In addition, Shoup's scheme (like H_r) is not eTCR. For the latter property, we need the modified \tilde{H}_r scheme (Eqn. (2)).

HMAC. One interesting question is why not take as the randomization mode a well established construction such as HMAC. It is easy to see that for the

purpose of TCR hashing the HMAC scheme (with the salt used as a random but known key) is not stronger than just pre-pending a random value to a message, a measure of limited effect against recent cryptanalysis. Other measures, such as appending randomness at the end, or even in the middle of a message, are even worse in that they do not help against collision attacks. In particular, this shows that the simplicity of our schemes cannot be explained by the naïve view that “any randomization of the input to the hash function will work.” Yet, one analogy to the case of HMAC is illustrative here: the fact that the HMAC design (as a message authentication code) has not been endangered by the latest collision attacks is due to its deliberate non-reliance on full collision resistance. We believe that randomized hashing could have a similar significant effect on the future security of digital signatures.

Relations between various notions. In our presentation we use quite a few notions of security for hash functions, some new and others well known (see [24]). Articulating the relations between these notions is *not* the focus of this paper. In some places we comment about implication or separation between these notions, but we do not explicitly show separation examples in this extended abstract. (This will divert attention from what we see as the heart of this work, namely the analysis of the practical constructions H_r and \tilde{H}_r .)

Organization. In Section 2 we present and discuss the SPR-like notions that we use in proving the TCR and eTCR properties of our schemes. These proofs are presented in Section 3 and Section 4, respectively. In Section 5 we discuss some practical aspects related to the integration of our schemes with standard signature schemes. We end with some open problems in Section 6.

2 Variants of Second-Preimage Resistant Hashing

Our goal in this work is to analyze the randomized modes of operation defined in Equations 1 and 2, finding “as weak as possible SPR-like properties” of h that suffice to ensure that H_r is TCR and \tilde{H}_r is eTCR. Below we describe several “games” representing different forms of second-preimage resistance (SPR) of the compression function h (recall that h represents a single, unkeyed, function with two arguments, c and m , and H^c represents the M-D iteration of h with $IV=c$). In all cases the goal of the attacker (that we denote by S) is to present pairs (c, m) and (c', m') such that $(c, m) \neq (c', m')$ and yet $h(c, m) = h(c', m')$. The difference between the games is how the values c, m, c', m' are determined and by whom.

The case where S chooses all four values is known as the “pseudo-collision” search problem for h and is at the basis of the Merkle-Damgård (M-D) construction of collision resistance. However, in the games below the adversary’s task is significantly harder since some of these values are chosen randomly and are not under the control of S . More specifically, the games are variants of “second-preimage resistance” of the function h since *in all the value m is chosen independently at random* and c', m' are chosen by the attacker at will after seeing m .

The difference is in the way c is determined: at random in r-SPR, chosen (by the attacker) in c-SPR, and evaluated (as a function of m) in e-SPR:

r-SPR: S receives random m and c , and it chooses m', c' .

c-SPR: S receives random m , and it chooses c, c', m' .

e-SPR: The game is parametrized by an IV c_0 . S chooses $\ell \geq 1$ values Δ_i , $i = 1, \dots, \ell$, each of length b bits; then S receives a random $r \in \{0, 1\}^b$ and c, m are set to $m = r \oplus \Delta_\ell$ and $c = H^{c_0}(r \oplus \Delta_1, \dots, r \oplus \Delta_{\ell-1})$. Finally S chooses c', m' .²

Clearly, r-SPR corresponds to the standard notion of second-preimage resistance of h , while c-SPR and e-SPR are variants that we use in the analysis of our schemes, and we elaborate on them below.

The IV. The game e-SPR can be considered in either the “uniform setting” where c_0 is chosen at random and given to S as input, or the “non-uniform setting” where c_0 is a parameter of the game and S may depend on that parameter. The discussion below applies equally to both settings.

Attack parameters. When discussing the security (or hardness) of different games we denote by t and L the attacker’s resources (i.e., time and length of messages used in the attack) and by ε the probability of success by the attacker. In particular, if G denotes one of the SPR or TCR games discussed in the paper then we say that a function (either h or the family H_r) is $G(\varepsilon, L, t)$ if no attacker with resources t and L can win the game with probability better than ε . (In some of the games the parameter L is irrelevant and then omitted.)

2.1 The c-SPR Game

We remark that the game c-SPR is a significantly easier variant (for the attacker) than the r-SPR and e-SPR games, since the attacker gets to choose c , after seeing m , in addition to choosing c' and m' . For example, the c-SPR game is clearly vulnerable to generic birthday attacks, whereas r-SPR and e-SPR are not.

We observe that for a compression function $h(c, m)$, the c-SPR property is equivalent to the property that the family $h_r(m, c) = h(c, m \oplus r)$ is $CR_n(n + b, n)$ as defined by Mironov [18].³ A family of functions $\{h_r\}_r$ belongs to $CR_\ell(n', n)$ if each h_r maps n' bits to n bits, and in addition no feasible attacker can win the following game except with insignificant probability:

1. The attacker first commits to $n' - \ell$ bits of the first message (call this $(n' - \ell)$ -bit string M_1),
2. Then the attacker is given the salt r ,
3. The attacker wins if it can find an ℓ -bit string M_2 and a second n' -bit message M' such that for $M = M_1 | M_2$ it holds that $h_r(M) = h_r(M')$.

² We choose to define $m = r \oplus \Delta_\ell$ for notational convenience in our proofs; however, note that the xor with Δ_ℓ does not change the fact that the value of m is determined uniformly at random (and independent of the Δ values).

³ Here we reverse the order of the inputs to h_r in order to match Mironov’s formulation.

Now, consider the following equivalent formulation of the c-SPR game: the attacker first commits to a b -bit message m , then it gets the randomness r and it needs to find c, c' and m' such that $h(c, m \oplus r) = h(c', m' \oplus r)$. Clearly, under this formulation we see that $\{h_r\}_r$ is c-SPR if and only if $\{h_r\}_r$ belongs to the class $CR_n(n+b, n)$. Mironov proved that the existence of $CR_n(n+b, n)$ families implies the existence of collision-resistant families (from $n+1$ to n bits). Hence c-SPR implies the *existence* of collision-resistant hashing families (in particular, this shows that Simon's [27] separation result applies to black-box constructions of c-SPR; it does *not* mean, however, that the M-D iteration of a c-SPR compression function is necessarily collision resistant). Still, for a particular compression function h , breaking c-SPR is also significantly harder (for the attacker) than the traditional pseudo-collision search. For example, although a pseudo-collision attack on MD5 is known for many years, breaking it in the sense of c-SPR seems beyond the state of the art in cryptanalysis, and even more so for SHA-1. (Also, assuming that c-SPR functions exist, it is easy to construct compression functions that are c-SPR but are not resistant to pseudo-collisions.) We believe that the c-SPR game provides a useful toy-example for cryptanalysts to develop tools for TCR-type attack models.

2.2 The e-SPR Game

The game e-SPR is significantly harder for the attacker than c-SPR (and in particular it is not open to generic birthday attacks). It is also closer to r-SPR. There are, however, two important differences between r-SPR and e-SPR:

- The distribution from which m is chosen is uniform, but c is determined as a function of m . In particular, the joint distribution of c and m has b bits of entropy, while the pair (c, m) is $(b+n)$ -bits long. This difference can in principle make e-SPR either easier or harder than r-SPR, depending on the underlying hash function.
- The attacker S gets to influence the distribution on c via its choice of the Δ_i 's (but note that S must commit to the Δ_i 's before it gets the random value r !)

Due to the first point above, e-SPR and r-SPR are formally incomparable, and it is possible to concoct examples of compression functions where one is easy and the other is hard (assuming r-SPR/e-SPR functions exist).

Relation of e-SPR to L -order TCR. As noted earlier, the fact that a family of compression functions is TCR does not necessarily imply that the multi-block extension of that family obtained via M-D iterations is TCR. Therefore, it is natural to search for (stronger) properties of compression functions that suffice to ensure that the M-D extension is TCR. Hong, Preneel and Lee [14] identified such a property, that they called *L -order TCR* (we recall it below), and proved that if a compression family is L -order TCR then the M-D extension is TCR on inputs of (up to) L blocks. Thus, one way to prove that our H_r construction (Eqn. (1)) is TCR is to assume that the compression function family $\{h_r\}_r$ defined by

Eqn. (1) (i.e. $h_r(c, m) = h(c, m \oplus r)$) is L -order TCR. While this property may be a plausible assumption on practical compression functions (and hence can be used as evidence to support the TCR property of our H_r scheme), it turns out that our Theorem 1 provides for a stronger result. Indeed, that theorem only assumes h_r to be e-SPR and we show in Proposition 1 below that e-SPR is a weaker property (i.e., harder to break) of the h_r scheme than L -order TCR. (We note that this relative strength of e-SPR and L -order TCR may not hold for all families of compression functions but it does hold for our specific scheme).

Hong et al. [14] define a function family $\{h_r\}_r$ to be L -order TCR if no feasible adversary can win the following game, except with insignificant probability: As in the TCR game, the attacker commits to a message m , then it is shown the salt r and it needs to find another message m' such that $h_r(m) = h_r(m')$. The difference is that before committing to the first message m the attacker gets to learn some information about the salt r by means of L adaptive queries m_i to the hash function for which the attacker gets as a response the corresponding hash values $h_r(m_i)$.

Proposition 1. *Let $h : \{0, 1\}^{b+n} \rightarrow \{0, 1\}^b$ be a compression function. If the family $\{h_r\}_{r \in \{0, 1\}^b}$ defined by $h_r(c, m) = h(c, m \oplus r)$ is order- L TCR(ε, t), then $h(\cdot)$ is e-SPR($\varepsilon, L, t - O(L)$).*

The proof is presented in the full version of the paper.

3 Achieving Target Collision Resistance

Recall the definition of the construction H_r from Eqn. (1),

$$H_r^c(m_1, \dots, m_L) \stackrel{\text{def}}{=} H^c(m_1 \oplus r, \dots, m_L \oplus r).$$

Using the games from Section 2 we can establish some relations between the corresponding SPR flavors of h and the TCR property of H_r . For this we define the TCR game for a family H_r as follows:

TCR Game: The attacker T knows the fixed IV c_0 . It chooses a message M of length $L > 0$ blocks, receives a random $r \in_{\mathbb{R}} \{0, 1\}^b$, and outputs a second message M' of length L' . Attacker T wins if $M' \neq M$ and $H_r^{c_0}(M') = H_r^{c_0}(M)$.

For simplicity of presentation, we first state and prove our results in the case where the TCR attacker can only provide M, M' of length an integer multiple of the block length, and it is not allowed to output M, M' such that one is a suffix of the other. This restriction on the attacker T is not significant in practice since actual implementations pads the messages and append the length to the last block, thus forcing full-block suffix-freeness. (To address the case where such length appending is not necessarily done one needs an additional “one-wayness assumption”, as described in Section 3.1. Also, see Section 5 for a discussion on how to handle partial blocks.) Formally, we consider a modified game (denoted TCR*) where the attacker does not win the game if it violates the condition on the messages.

TCR* Game: Same as TCR but the attacker wins only if in addition to the regular TCR conditions it also holds that M, M' consists of an integral number of b -bit blocks and neither of them is suffix of the other.

Theorem 1. (*TCR Theorem, suffix-free case*)

1. If h is c -SPR(ε, t) then H_r is $\text{TCR}^*(L\varepsilon, L, t - O(L))$
2. If h is e -SPR(ε, L, t) then H_r is $\text{TCR}^*(L\varepsilon, L, t - O(L))$

This theorem is proven below. We note that these are sufficient conditions for TCR-ness *but not necessary ones*. In other words, the failure of a compression function to one of the above attacks does not necessarily mean that the induced family H_r is not TCR.

It is also worth commenting on the tightness of the reductions in the above theorem. In both cases, c -SPR and e -SPR, there is a linear degradation of security when going from SPR to TCR. Note, however, that there is a significant difference between the *generic security* of the two games c -SPR and e -SPR. Against the former there is a trivial birthday-type generic attack while against e -SPR generic attacks achieve only linear advantage. Thus, the reduction from e -SPR shows a worst-case “quadratic degradation” generic TCR attack against H_r (and not “cubic degradation” as the reduction cost in the c -SPR to TCR case could indicate). Moreover, TCR attacks with quadratic degradation (i.e., with success probability in the order of $tL/2^n$) against H_r are indeed possible since SPR attacks against H can be translated into TCR attacks against H_r and we know that such birthday-type SPR attacks against H exist [11,16].

This motivates two questions: Can we have a flavor of SPR for which there is a tight reduction to TCR? Can this SPR game be defined such that it is only vulnerable to linear generic attacks? The answer to the first question is YES: in the full version of this paper we present such a game, called m -SPR (m for multiple), and its tight reduction to TCR. The answer to the second question is obviously NO, since as said SPR birthday attacks against H_r do exist. It is possible, however, to design variants of M-D hash functions (e.g., appending a sequence number to each block or using the “dithering” technique proposed by Rivest [23]) for which the best generic attacks achieve linear degradation only. Thus, the main motivation and usefulness of the game m -SPR and its tight reduction to TCR is for the analysis of such variants.

3.1 Proof of Theorem 1

In the proof below we assume that the attacker T never outputs M, M' where one is the suffix of the other. We also make the convention that when algorithm A calls another algorithm B and B aborts then A aborts too.

Lemma 1. *If h is c -SPR(ε, t) then H_r is $\text{TCR}^*(L\varepsilon, L, t - O(L))$.*

Before proving this lemma, we comment that the exact assertion of this lemma (and thus the proof) depend on whether we view the IV c_0 as a random input to the TCR attacker or a parameter of the TCR game (cf. comment on page 47).

In the reduction below we suppress c_0 with the understanding that if it is a parameter then the same parameter is used in the reduction, and if it is an input then S chooses c_0 at random and gives it to T .

Proof. Given TCR attacker T we build c-SPR attacker S :

1. S gets input m , invokes T and gets from T the first message $M=(m_1, \dots, m_L)$.
2. S chooses $\ell \in_{\mathbb{R}} \{1, \dots, L\}$, sets $r = m \oplus m_\ell$, and returns r to T .
3. T outputs the second message $M' = (m'_1, \dots, m'_{L'})$ or \perp .
4. S checks that there exists $\ell' \in \{1, \dots, L'\}$ such that (1) either $m'_{\ell'} \neq m_\ell$ or $H_r(m'_1, \dots, m'_{\ell'-1}) \neq H_r(m_1, \dots, m_{\ell-1})$; and (2) $H_r(m'_1, \dots, m'_{\ell'}) = H_r(m_1, \dots, m_\ell)$. If no such ℓ' exists then S aborts.
5. Otherwise S outputs $c = H_r(m_1, \dots, m_{\ell-1})$, $c' = H_r(m'_1, \dots, m'_{\ell'-1})$ and $m' = m'_{\ell'} \oplus r$. (If there is more than one index with the properties from above then S chooses ℓ' arbitrarily among all these indexes.)

Turning to the analysis of S , we first observe that if T does not abort then there must exist a pair of indexes i, i' that satisfy the properties from Step 4, namely $H_r(m'_1, \dots, m'_{i'}) = H_r(m_1, \dots, m_i)$ and either $m'_{i'} \neq m_i$ or $H_r(m'_1, \dots, m'_{i'-1}) \neq H_r(m_1, \dots, m_{i-1})$.

To see that, note that since M, M' are suffix-free there exists an integer $j < \min(L, L')$ such that $m_{L-j} \neq m'_{L'-j}$. Consider the smallest such j (i.e., the last block where M, M' differ when their ends are aligned). Also, let k be smallest non-negative integer such that $H_r(m'_1, \dots, m'_{L'-j+k}) = H_r(m_1, \dots, m_{L-j+k})$ (i.e., the first “collision” after these “last differing blocks” in the computation of $H_r(M')$, $H_r(M)$). Clearly, such k exists since T did not abort and so there must be some collision because $H_r(M') = H_r(M)$.

Now setting $i = L - j + k$ and $i' = L' - j + k$ we get the pair that we need. This is because we have $H_r(m'_1, \dots, m'_{L'-j+k}) = H_r(m_1, \dots, m_{L-j+k})$ by definition, and (a) either $k = 0$ in which case $m_{L-j+k} \neq m'_{L'-j+k}$, or (b) $k > 0$ in which case $H_r(m'_1, \dots, m'_{L'-j+k-1}) \neq H_r(m_1, \dots, m_{L-j+k-1})$.

Next, we argue that the random choice of ℓ is independent of the transcript that T sees, and in particular it is independent of the values of any such pair i, i' as above. This is easy to see, since the distributions of all the variables in the execution of S remains unchanged if instead of choosing m at random and setting $r = m_\ell \oplus m$ we choose r at random and set $m = m_\ell \oplus r$. In the new game, however, we can postpone the choice of ℓ (and m) until the end, and so it is clearly independent of M, M' .

It follows that whenever T does not abort, there exists a pair i, i' as above, and in this case the probability that S chooses $\ell = i$ is (at least) $1/L$. If that happens then there exists ℓ' as needed (i.e., $\ell' = i'$) so S does not abort. It remains to show that in this case S wins the c-SPR game. This follows since by the condition on ℓ, ℓ' we have:

either $c = H_r(m_1, \dots, m_{\ell-1}) \neq H_r(m'_1, \dots, m'_{\ell'-1}) = c'$ or $m' = m'_{\ell'} \oplus (m_\ell \oplus m) \neq m$

and also (due to $m = r \oplus m_\ell$ and $m' = r \oplus m'_{\ell'}$):

$$h(c, m) = h(c, r \oplus m_\ell) = H_r(m_1, \dots, m_\ell) = H_r(m'_1, \dots, m'_{\ell'}) = h(c', r \oplus m'_{\ell'}) = h(c', m')$$

The attacker in the proof of Lemma 1 makes little use of the full freedom provided by the c -SPR game in (adaptively) choosing c . Indeed the value of c used by S is set as soon as S chooses ℓ , and this choice is *independent of m* . This motivates the e -SPR game where, by definition, S can influence c only through the choice of values Δ to which it commits before seeing r . Thus, e -SPR represents a better approximation of the TCR game, by giving the attacker in e -SPR less artificial freedom than in c -SPR, and making it much closer to SPR (in particular, by disallowing generic birthday attacks against it). This makes the following lemma much stronger than Lemma 1 though its proof follows the same lines. Note that we are still paying a linear (in L) degradation of security but this time the reduction is to a “non-birthday” problem. (See also the discussion about security degradation preceding Section 3.1.)

Lemma 2. *If h is e -SPR(ε, L, t) then H_r is $\text{TCR}^*(L\varepsilon, L, t - O(\ell))$.*

Proof. This is a simple adaptation of the proof of Lemma 1. Given TCR attacker T we build e -SPR attacker S :

1. S invokes T and gets from T the first message $M = (m_1, \dots, m_L)$.
2. S chooses $\ell \in_{\mathbb{R}} \{1, \dots, L\}$ and sets $\Delta_i = m_i$ for $i = 1, \dots, \ell$.
3. S receives a random r and forward it to T . (The values of c, m are evaluated as $c = H_r(\Delta_1, \dots, \Delta_{\ell-1}) = H_r(m_1, \dots, m_{\ell-1})$ and $m = \Delta_\ell \oplus r = m_\ell \oplus r$.)
4. T outputs the second message $M' = (m'_1, \dots, m'_{L'})$ or \perp .
5. S checks that there exists $\ell' \in \{1, \dots, L'\}$ such that (1) either $m'_{\ell'} \neq m_\ell$ or $H_r(m'_1, \dots, m'_{\ell'-1}) \neq H_r(m_1, \dots, m_{\ell-1})$; and (2) $H_r(m'_1, \dots, m'_{\ell'}) = H_r(m_1, \dots, m_\ell)$. If no such ℓ' exists then S aborts.
6. Otherwise S outputs $c' = H_r(m'_1, \dots, m'_{\ell'-1})$ and $m' = m'_{\ell'} \oplus r$. (If there is more than one index with the properties from above then S chooses ℓ' arbitrarily among all these indexes.)

The analysis of S is nearly identical to the analysis in the proof of Lemma 1. The only difference is that here it is even easier to argue that the choice of ℓ is independent of the transcript of T .

Extension to Non-suffix-free Messages. We stated Theorem 1 and its proof in terms of the TCR^* game, namely, the TCR game in which the attacker is restricted to suffix-free pairs M, M' . For the sake of generality, we extend these results here to the case where inputs to H_r may not be suffix free. For this extension we need the following one-wayness assumption.

Assumption 2 (OWH). *The assumption OWH for the compression function h asserts that given a random c , it is hard to find a non-null message M such that $H^c(M) = c$.*

We use our usual notation $\text{OWH}(\varepsilon, t)$ to denote the assumption that no t -time attacker can violate the assumption OWH with probability better than ε .

Some comments are in order here. First, note that this assumption is stated in the “uniform” setting where the IV c is chosen at random. Its “non-uniform”

counterpart (with respect to a parameter c_0) does not make formal sense, for the same reason that the assertion “SHA256 as per FIPS-180-2 is collision resistant” does not make formal sense. This means that formally, Theorem 3 below only applies to this “uniform” setting.⁴

Second, the above assumption is clearly implied by the assumption that for a random c it is hard to find c', m such that $h(c', m) = c$. Furthermore, under a very mild condition on the structure of h , this last assumption is equivalent to the standard assumption that h is a one-way function (i.e., for random c, m it is hard given $h(c, m)$ to find c', m' such that $h(c', m') = h(c, m)$). Specifically, the condition that we need is that for random c, m , the distribution over $h(c, m)$ is statistically close to (or at least indistinguishable from) the uniform distribution on $\{0, 1\}^n$.

It is well known that one-wayness is implied by second-preimage resistance, and is it easy to verify that the same holds also for our notions of c-SPR and e-SPR (the latter in its “uniform” interpretation where the IV is random). It follows that under that mild structural condition, the assumption OWH is implied by the assumptions e-SPR and c-SPR and so is redundant in Theorem 3 below. Still, we prefer to state the theorem with this assumption since (a) it makes it easier to understand, and (b) it applies also to functions h for which that structural condition does not hold.

Theorem 3. (*TCR Theorem , non-suffix-free case*)

1. If h is c-SPR(ε, t) and OWH(ε', t) then H_r is TCR($\varepsilon' + L\varepsilon, L, t - O(L)$).
2. If h is e-SPR(ε, L, t) and OWH(ε', t) then H_r is TCR($\varepsilon' + L\varepsilon, L, t - O(L)$).

The proof is a small adaptation of the proofs of lemmas 1 and 2. One just observes that in the case where one of M, M' is a proper suffix of the other, then either the prefix of the longer message violates the one-wayness assumption (i.e., we have $H^c(\text{prefix}) = c$), or else the same analysis from lemmas 1 and 2 applies. The reason that ε' is not multiplied by L in the expressions above is that in the former case we do not care about the value of ℓ that was chosen at the beginning of the reduction. (See the proof of Theorem 4 for the formal argument.)

4 Enhanced Target Collision Resistance

Recall the definition of the construction \tilde{H}_r from Eqn. (2),

$$\tilde{H}_r^c(M) \stackrel{\text{def}}{=} H_r^c(0|M) = H^c(r, m_1 \oplus r, \dots, m_L \oplus r)$$

We show that under the same c-SPR and e-SPR assumptions, the construction \tilde{H}_r is enhanced TCR (eTCR). We start by defining the eTCR game.

⁴ However, the reduction that proves Theorem 3 is meaningful also for the case of fixed IV, showing a constructive transformation from TCR attack to either an SPR attack or to an attack that violates the one-wayness assumption.

eTCR Game: Attacker T chooses a message M of length $L \geq 0$ blocks, receives a random $r \in_{\mathbb{R}} \{0, 1\}^b$, and outputs another $r' \in \{0, 1\}^b$ and a message M' of length L' blocks. Attacker T wins if $(M', r') \neq (M, r)$ and $H_r^{c_0}(M') = H_r^{c_0}(M)$.

Note that the attacker can win this game even when $M = M'$, so long as $(M', r') \neq (M, r)$. We remark that as opposed to the case of the construction H_r and the standard TCR game, here assuming suffix-freeness does not help us, because the attacker can specify $r' \neq r$ and so even if M, M' are suffix free the messages $M \oplus r, M' \oplus r'$ perhaps are not. (However, we can go back to the suffix-free case by using length padding outside of the randomness, as discussed in Section 5.) Below we therefore use the OWH assumption as in Theorem 3, thus getting:

Theorem 4 (eTCR security)

1. If h is c-SPR(ε, t) and OWH(ε', t) then \tilde{H}_r is eTCR($\varepsilon' + (L+1)\varepsilon, L, t - O(L)$).
2. If h is e-SPR($\varepsilon, L+1, t$) and OWH(ε', t) then \tilde{H}_r is eTCR($\varepsilon' + (L+1)\varepsilon, L, t - O(L)$).

Proof. Below we only prove part 2. Part 1 can be obtained as a corollary, since an e-SPR attacker can be trivially converted to a c-SPR attacker with the same success probability. The proof is a small adaptation of the proof of Lemma 2. The idea is that S sets $\Delta_0 = 0$ (and the other Δ_i 's as before) and then uses r' instead of r to compute c' and m' . In more details, given an enhanced-TCR attacker T we build e-SPR attacker S :

1. On input an IV c , S invokes $T(c)$ and gets from T the first message $M = (m_1, \dots, m_L)$.
2. S sets $m_0 = 0$, chooses $\ell \in_{\mathbb{R}} \{0, \dots, L\}$ and sets $\Delta_i = m_i$ for $i = 0, \dots, \ell$.
3. S receives a random r and forward it to T . (Also c, m are evaluated as $c = H_r^c(\Delta_0, \Delta_1, \dots, \Delta_{\ell-1}) = H_r^c(0, m_1, \dots, m_{\ell-1})$ and $m = \Delta_{\ell} \oplus r = m_{\ell} \oplus r$.)
4. T outputs either the second salt r' and message $M' = (m'_1, \dots, m'_{L'})$ or \perp .
5. S sets $m'_0 = 0$, and searches for an index $\ell' \in \{0, \dots, L'\}$ such that (1) either $m'_{\ell'} \oplus r' \neq m_{\ell} \oplus r$ or $H_{r'}^c(0, \dots, m'_{\ell'-1}) \neq H_r^c(0, m_1, \dots, m_{\ell-1})$; and (2) $H_{r'}^c(0, m'_1, \dots, m'_{\ell'}) = H_r^c(0, m_1, \dots, m_{\ell})$. If no such ℓ' exists then S aborts.
6. Otherwise S outputs $c' = H_{r'}^c(0, m'_1, \dots, m'_{\ell'-1})$ and $m' = m'_{\ell'} \oplus r'$. (If there is more than one index with the properties from above then S chooses ℓ' arbitrarily among all these indexes.)

To analyze S , let SUFF denote the event in which one of the messages

$$X = (r, m_1 \oplus r, \dots, m_L \oplus r) \text{ and } X' = (r', m'_1 \oplus r', \dots, m'_{L'} \oplus r')$$

in the game above is a suffix of the other, and in addition hashing only the prefix of the longer message yields back the IV c . In other words, the event SUFF occurs whenever we have either $X' = Y|X$ or $X = Y|X'$ with some $Y \neq \Lambda$ for which $H^c(Y) = c$.

Let $\varepsilon_{\text{suff}}$ be the probability of the event SUFF, and let ε^* be the probability of the event in which T succeeds but the event SUFF does not happen. In the latter case we can apply the same analysis as in Lemma 2, showing that there must exist a pair of indexes i, i' that satisfy the conditions in step 5, that with probability at least $1/(L+1)$ we have $\ell = i$, and if that happens then S succeeds.

The only part of the analysis that is slightly different than in the proof of Lemma 2 is proving the existence of the indexes i, i' in the case where one of X, X' is a prefix of the other but $H^c(Y) \neq c$. In this case, denote by $X|_i$ the $(i+1)$ -block prefix of X (i.e., blocks 0 through i), and similarly denote $X'|_i$. Let j be the smallest integer such that $H^c(X|_{L-j}) \neq H^c(X'|_{L-j})$ (i.e., the last place where the chaining values do not agree). Since $H^c(Y) \neq c$ then there exists such index $j \leq \min(L+1, L'+1)$, and on the other hand $j > 0$ since the success of T implies that $H^c(X_L) = H^c(X'_L)$. Then setting $i = L - j + 1$ and $i' = L' - j + 1$ we get the pair of indexes that we need.

We thus conclude that the success probability of S is at least $\varepsilon^*/(L+1)$. Since (by definition) the success probability of T is at most $\varepsilon^* + \varepsilon_{\text{suff}}$, we can complete the proof by showing a OWH-attacker with success probability $\varepsilon_{\text{suff}}$. Such attacker S' is obvious: It gets c as input and runs the same execution as the e-SPR attacker S from above (choosing itself the random r), and at the end if the event SUFF occurs then it outputs the prefix Y .

5 Using Randomized Hashing in Signatures

The randomized hashing modes presented in this paper are intended to replace the deterministic hashing used by standardized signature schemes such as RSA [21] and DSS [10]. As shown throughout the paper, this replacement frees these schemes from their current essential dependency on full collision resistance by basing their security on much harder to break properties of the underlying hash functions. Of the two schemes analyzed here, the \tilde{H}_r mode (Eqn. (2)) is best suited for this task as it does not require the explicit signing of the salt r and hence allows for more implementation flexibility. We present a full specification of this mode and its use in digital signature in [13]. Here we provide a short (and partial) discussion of possible approaches.

The two main approaches depend on whether an application can accommodate the sending of the salt r in addition to the message and signature or whether any increase on the size of the information is not possible. The first case is simpler and requires the least changes to standards. Most applications will be able to send the the extra salt, especially that r needs not be too long, say in the range of 128-256 bits (see below). Examples of such applications are digital certificates, code signing, XML signatures, etc. In this case, upon the need to sign a message M the modified application will: (i) choose a random salt r ; (ii) transform the message $M = (m_1, \dots, m_L)$ into the message $M' = (r, m_1 \oplus r, \dots, m_L \oplus r)$; (iii) invoke the *existing* signature operation (including the existing hash operation) on M' ; (iv) when the message M and its signature are to be transmitted, transmit r , M and the computed signature (on M'). The verification side computes M' from

the received r and M and applies its *existing* signature verification operation on M' . This approach allows for preserving the existing processing order (including a one-pass over the signed message) and the possible pre-computation (ahead of signature generation) of the pair $(r, \tilde{H}_r(M))$.

Note that in this case the application can use its existing signature and verification processing (whether in software, hardware or both) *without any change*. In this case, signature standards [21,10] need no change except for adding the “front end” message randomization (to generate M'). The details of implementation of the required changes is application and implementation dependent. The extra operations may be placed at the application itself or in a “signing module” invoked by the application and which will be responsible for the signature generation as well as for the generation and transmission of r .

Applications where an increase in the size of messages or signatures is impractical will need to resort to a different approach: either re-use the existing randomness (in the case of probabilistic signatures) or encode the salt r “under the signature”. The former is the case for the probabilistic schemes DSS [10] and RSA-PSS [21] where the already existing random values r (used internally by the PSS operation in the case of RSA-PSS and as the signature component $r = g^k$ in the case of DSS) can be re-used as the hashing salt and thus require no additional transmission. The latter case, i.e., encoding r “under the signature”, is the case of the traditional deterministic RSA encoding of PKCS#1 v1.5 [21]. In this case, instead of padding the (deterministic) hash value to a full modulus size as done today, one will pad the concatenation of r and $\tilde{H}_r(M)$ to a full modulus size and then apply to it the private RSA operation. In this case, the salt r can be retrieved by the verifying party using the RSA public operation and hence no extra transmission is required.

Evaluating the complexity of performing the above changes depends in particular settings and applications. Clearly, no change to existing applications, signature modules or standards is “too simple” and one cannot ignore the cost of engineering the above changes. However, considering that applications that compute digital signatures will be upgraded in any case to use new hash functions, one should use this opportunity to also upgrade to randomized hashing. In particular, it is worth noting that many difficult issues for applications such as preserving backwards compatibility, the signaling/negotiation of algorithms and capabilities, version rolling, etc. are not made worse by our proposal than what is already needed to support a simple hash function upgrade [4]. Considering the simplicity and minimalistic nature of our randomization scheme, we believe that the extra changes (transmission or re-use of r) are well worth the substantial security gain they provide to existing and future signature schemes.

5.1 Specification Details

As said, a detailed specification of the \tilde{H}_r mode is presented in [13]. Here we mention two elements that we omitted so far and that need to be included in a practical instantiation.

Shorter salt. The construction \tilde{H}_r as defined in Eqn. (2) uses salt of length b , i.e. one message block. In practical instantiations, we propose to choose a salt string r' of length between 128 and b , and create the b -bit salt r by concatenating r' with itself as needed, possibly with the last r' repetition being truncated (e.g., in the case of $b = 512$ and a 160-bit r' , one would have $r = r' || r' || r' || r''$ where r'' represents the first 32 bits of r'). The analysis in this paper applies to this modified salt, except that now the assumptions c-SPR, e-SPR are stated in terms of random messages distributed according to the distribution induced by r' . Of course, the plausibility of these assumptions needs to be evaluated; however, as long as r' is not too short, these modified SPR-type assumptions seem very reasonable (it is even possible that the added “structure” in these repetitions makes the cryptanalytical problem more difficult).

Last block padding. The Merkle-Damgård construction pads the input message to an integral number of blocks and includes an encoding of the length of the original message in the last block (thus ensuring a suffix-free message space). The analysis in this work assumes that the salt r is xor-ed to the padded message, but in practice it is likely that the xor will happen first (as part of the randomized mode-of-operation) and the padding will then be applied to the randomized message (as part of the underlying hash function). Simply switching the order and using randomize-then-pad may mean that the last block is only randomized in its first few bits. A more robust alternative is to use two levels of padding. Namely, we first use suffix-free padding of the last block to a full block minus 64 bits (i.e., 448 bits), then xor the salt to this padded message, and finally apply the original hashing scheme (which will include the length of the padded message in the remaining 64 bits). See [13].

6 Open Problems

In light of the results in this paper, we feel that more focus should be placed on evaluating current and future hash functions against the c-SPR and e-SPR attack scenarios. Specifically, we would like to offer the c-SPR scenario as a “toy example” for developing cryptanalytical tools that may prove useful in assessing randomized hashing, and then study the e-SPR scenario as a stronger foundation for our scheme. Other open problems that arise from this work are briefly discussed next.

A different construction for eTCR. Another natural proposal for obtaining enhanced-TCR is to define $\overline{H}_r^c(M) = H^c(r | H_r^c(M))$. It may be interesting to study this variant and see what advantages (and disadvantages) it offers vis-a-vis the construction \tilde{H} from Eqn. (2).

The random-oracle model and weak hash functions. One inherent difficulty in formally proving the security of TCR and eTCR schemes in the context of DSS and RSA as specified in the official standards [21,10] is that these schemes do not have a proof of security (not even in the case that the underlying hash function is fully collision resistant). The “closest relatives”, namely, the DSA-II

variant of Brickel, Pointcheval, Vaudenay and Yung [7] or the RSA-PSS scheme of Bellare and Rogaway [2], have a proof of security in the random oracle model. Interestingly, these proofs use in an essential way the randomization of the hash function, not unlike the TCR or eTCR constructions. In some sense one can think, for example, of the use of DSS with a eTCR scheme as an “instantiation” of DSS-II. The obvious question is: how can we state anything related to random oracles when we are dealing with relatively weak (at least not CR) hash functions (clearly, random oracles are “very strong hash functions”). See some results related to this question in the full version of this paper and in [19].

We also point out that the random-oracle proofs for some of the above (randomized) signature schemes do not essentially use the collision resistance property of the random-oracle (as evidenced, for example, by the fact that some of these proofs, such as those following [22], remain meaningful even when you have a random-oracle with relatively short output, e.g., 80 bits). This brings up the interesting question of exhibiting variants of the random-oracle model where one can argue about functions that “behave randomly but are not collision-resistant” (e.g., a random oracle H with an associated oracle that outputs random H collisions upon request).

Acknowledgments. We thank Ran Canetti, Orr Dunkelman, Charanjit Jutla, Burt Kaliski, Tom Shrimpton and Lisa Yin for helpful discussions.

References

1. Mihir Bellare, Ran Canetti, Hugo Krawczyk, “Keying Hash Functions for Message Authentication”. CRYPTO 1996. 1-15
2. Mihir Bellare and Phillip Rogaway, “The Exact Security of Digital Signatures – How to Sign with RSA and Rabin”, EUROCRYPT 96.
3. Mihir Bellare and Phillip Rogaway, “Collision-Resistant Hashing: Towards Making UOWHFs Practical”, CRYPTO 97, LNCS 1294, 1997
4. Steven M. Bellovin and Eric K. Rescorla, “Deploying a New Hash Algorithm”, NDSS’06. <http://www.cs.columbia.edu/~smb/papers/new-hash.pdf>
5. Eli Biham and Rafi Chen, “Near-Collisions of SHA-0”, CRYPTO 2004.
6. Eli Biham, Rafi Chen, Antoine Joux, Patrick Carribault, Christophe Lemuet and William Jalby, “Collisions of SHA-0 and Reduced SHA-1”, EUROCRYPT 2005.
7. Ernest Brickell, David Pointcheval, Serge Vaudenay, and Moti Yung, “Design and Validations for Discrete Logarithm Based Signature Schemes”, PKC’2000.
8. Florent Chabaud and Antoine Joux, “Differential Collisions in SHA-0”, CRYPTO 98.
9. Ivan Damgård, “A design principle for hash functions”, CRYPTO 1989.
10. Digital Signature Standard (DSS), FIPS 186, May 1994.
11. Richard Drews Dean, “Formal Aspects of Mobile Code Security”, Ph.D Dissertation, Princeton University, January 1999.
12. Shafi Goldwasser, Silvio Micali and Ronald L. Rivest, “A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks”, *SIAM J. Comput.* 17(2): 281-308 (1988)
13. Shai Halevi and Hugo Krawczyk, “Randomized Hashing: Specification and Deployment”, in preparation.

14. Deukjo Hong, Bart Preneel and Sangjin Lee, "Higher Order Universal One-Way Hash Functions", ASIACRYPT 2004.
15. Antoine Joux, "Multicollisions in Iterated Hash Functions. Application to Cascaded Constructions", CRYPTO 2004.
16. John Kelsey and Bruce Schneier, "Second Preimages on n-Bit Hash Functions for Much Less than 2^n Work", EUROCRYPT 2005.
17. Ralph Merkle, "One way hash functions and DES", CRYPTO 1989.
18. Ilya Mironov, "Hash Functions: From Merkle-Damgård to Shoup", EUROCRYPT 2001: 166-181.
19. Ilya Mironov, "Collision-Resistant No More: Hash-and-Sign Paradigm Revisited", Public Key Cryptography 2006: 140-156.
20. Moni Naor and Moti Yung, "Universal One-Way Hash Functions and their Cryptographic Applications", STOC 1989.
21. PKCS #1 v2.1: RSA Cryptography Standard, RSA Laboratories, June 14, 2002
22. David Pointcheval and Jacques Stern, "Security Arguments for Digital Signatures and Blind Signatures", *J.Cryptology* (2000) 13:361-396.
23. Ron Rivest, "Abelian square-free dithering for iterated hash functions", Presented at ECrypt Hash Function Workshop, June 21, 2005, Cracow.
24. Phillip Rogaway, Thomas Shrimpton, "Cryptographic Hash-Function Basics: Definitions, Implications, and Separations for Preimage Resistance, Second-Preimage Resistance, and Collision Resistance". FSE 2004, 371-388.
25. John Rompel, "One-way functions are necessary and sufficient for secure signatures", STOC 1990, pp. 387-394.
26. Victor Shoup, "A Composition Theorem for Universal One-Way Hash Functions", EUROCRYPT 2000.
27. Dan Simon, "Finding collisions on a one-way street: Can secure hash functions be based on general assumptions?", EUROCRYPT 98, pp. 334-345.
28. Michael Szydlo and Yiqun Lisa Yin, "Collision-Resistant usage of MD5 and SHA-1 via Message Preprocessing", Cryptology ePrint Archive, Report 2005/248.
29. Xiaoyun Wang, Xuejia Lai, Dengguo Feng, Hui Chen, and Xiuyuan Yu, "Cryptanalysis of the Hash Functions MD4 and RIPEMD", EUROCRYPT 2005.
30. Xiaoyun Wang and Hongbo Yu, "How to Break MD5 and Other Hash Functions", EUROCRYPT 2005.
31. Xiaoyun Wang, Hongbo Yu, and Yiqun Lisa Yin, "Efficient Collision Search Attacks on SHA-0", CRYPTO 2005.
32. Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu "Finding Collisions in the Full SHA-1", CRYPTO 2005.