

Improving the Exact Security of Digital Signature Schemes*

Silvio Micali and Leonid Reyzin

MIT LCS, 545 Technology Square,
Cambridge, MA 02139, U.S.A.
reyzin@theory.lcs.mit.edu

Communicated by Jaques Stern

Received September 1999 and revised September 2000

Online publication 9 April 2001

Abstract. We put forward a new method of constructing Fiat–Shamir-like signature schemes that yields better “exact security” than the original Fiat–Shamir method. (We also point out, however, that such tight security does not make our modified schemes always preferable to the original ones. Indeed, there exist particularly efficient Fiat–Shamir-like schemes that, though only enjoying “loose security,” by using longer keys may *provably* provide more security at a lower computational cost than their “tight-security” counterparts.)

Key words. Digital signatures, ID schemes, Random oracles, Exact security.

1. Introduction

1.1. *Exact Security of Signature Schemes*

Goldwasser et al.’s [GMR] classical notion of security for a digital signature scheme is asymptotic in nature. In essence, a proof of security amounts to a reduction from forging a signature to solving a computationally hard problem: if a polynomial-time forger exists, then we can use it to solve the hard problem in polynomial time.

It has been often pointed out that this asymptotic approach, which uses notions such as “polynomial time” and “sufficiently large,” is too coarse for practical security recommendations. Knowing that no polynomial-time adversary has a better than exponentially small chance of forgery for a sufficiently large security parameter does not provide one with an answer to the practical problem of finding the appropriate security parameters to ensure security against adversaries with certain concrete capabilities.

* This material is based upon work supported in part under a National Science Foundation Graduate Fellowship. A preliminary extended abstract of this work appears in [MR]. A version with more complete proofs is available from <http://theory.lcs.mit.edu/~reyzin>.

Bellare and Rogaway [BR2] argue that, in order to be able to deduce concrete security recommendations, it is important to be precise in the reduction from a forger to the algorithm that solves the hard problem. For example, if one knows that factoring integers of length l is no more than 100 times harder than breaking a certain signature scheme with security parameter l , then one could pick l so that even 1% of the work required to factor integers of length l is considered infeasible.

A reduction in which the difficulty of forging and the difficulty of solving the underlying hard problem are close is called *tight*; otherwise, it is called *loose*. (Naturally, “close,” “tight,” and “loose” are imprecise terms and make more sense when used in the comparative.) A scheme whose exact security is tightly related to the difficulty of factoring is also proposed in [BR2].

1.2. Loose Security of Fiat–Shamir-Like Signature Schemes

A fruitful method for constructing signature schemes was introduced by Fiat and Shamir [FS]. Although claimed for a specific ID scheme, the method works with a general *commit–challenge–respond* ID scheme. The method consists of replacing the verifier’s random challenge by a publicly known “random” function H computed on the prover’s commitment and the message being signed. This removes interaction and adds the message into the picture, thus changing an ID scheme into a signature scheme.

Many of such signature schemes have been proven secure when the “random” function is modeled as a random oracle ([BR1] provides a formal treatment of this model). However, the reductions in these proofs are quite loose, thus necessitating larger key sizes. Unless a tighter reduction has been overlooked, the only way to improve the security of such signature schemes is to modify them to allow for tighter reductions.

1.3. Tightening the Security of Fiat–Shamir-Like Schemes

This paper’s main contribution is a modification to the factoring-based Fiat–Shamir-like schemes that makes their security very tightly related to the problem of integer factorization. Our modification is quite general and can be applied, in particular, to the schemes from [FS], [FFS], [OO1], [MS], [OS], [Ok], [Mi], [Sh], and [Sc2].

To exemplify our method and make the description concrete, we picked one of the simpler and more efficient schemes from the above list, the one in [Mi]. We henceforth call it “MSA” (for “Micali’s signature algorithm”). We first present an exact analysis of the loose security of MSA, then propose the modification (called the “swap method”) and present an exact analysis of the tight security of the modified scheme (called “MSA-swap”). Both MSA and MSA-swap are quite practical, with the performance comparable with that of the schemes currently used in practice.

1.4. When to Use Tight Security?

As Bellare and Rogaway explain in [BR2] using the schemes Full-Domain-Hash-RSA and PSS as examples, tighter exact security results in smaller security parameters and, hence, higher efficiency. Because Full-Domain-Hash-RSA and PSS have about the same running time for a given security parameter, PSS, with its tighter security, is better than Full-Domain-Hash-RSA.

Since then, folklore has often regarded schemes with tight security as superior to those with loose security, misunderstanding the examples in [BR2]. Indeed, an obvious (but often forgotten) observation is that

a more efficient scheme with loose security may be better than a less efficient one with tight security,

because the efficiency of the loosely secure scheme can compensate for the higher security parameter.

To demonstrate this with a quantitative example, in Section 5 we analyze MSA against MSA-swap and PRab (a variant of PSS, also from [BR2], based on squaring rather than RSA). We find that the loosely secure but more efficient MSA sometimes delivers more security for less cost than tightly secure schemes MSA-swap and PRab.

To our knowledge, despite the fact that our analysis is quite straightforward, no one has carried it out before (the paper [BR2] compares Full-Domain-Hash RSA and PSS, but the analysis is simpler because the two schemes have the same efficiency). The fact that a Fiat–Shamir-like scheme may be preferable to Rabin-based PRab for *both* security and efficiency reasons has not been pointed out before, either.

Our technique can be used to choose from among several signature schemes for a particular application. In particular, it can help one decide whether to use a Fiat–Shamir-like scheme directly or to apply our modification.

2. Definitions

Notation. We denote by $A^?$ a (probabilistic) oracle-calling algorithm; A^O denotes the same algorithm making calls to the specific oracle O , where $O : \{0, 1\}^* \rightarrow \{0, 1\}$. Because oracles returning more than one bit per question, as well as multiple oracles, can all be easily simulated by a single oracle returning one bit per question, we slightly abuse this notation and speak of, for example, $A^{?,?}$ and $A^{G,H}$, for some $G : \{0, 1\}^* \rightarrow \{0, 1\}^k$ and $H : \{0, 1\}^* \rightarrow \{0, 1\}^l$. Additionally, in some cases we need to allow one probabilistic oracle algorithm $A^?$ access to another probabilistic oracle algorithm $B^?$ with the oracle O ; we denote it by A^{B^O} . Note that because $B^?$ is probabilistic, it is not, strictly speaking, an oracle; rather, this notation implies that, for each query, B^O gives a single answer selected, with the appropriate probability, from the set of all possible answers. When convenient and clear, we omit the superscript when speaking about $A^?$ or A^O .

We denote by $x \leftarrow A$ the fact that the probabilistic algorithm A output x .

Signature schemes. Our definition of a signature scheme follows the one found in [GMR] and refined in [BR1] to allow for random oracles. As is common, we decouple the notions of a signature scheme and of its security.

Definition 1. A signature scheme with an oracle is a triple of probabilistic oracle algorithms $\Pi = (Gen^?, Sign^?, Ver^?)$. Gen is a key generation algorithm that, given the security parameter k , outputs a key pair (pk, sk) .¹ Commonly, pk is called a *public key* and

¹ Usually, we are interested in the running time of Gen as a function of k rather than $\log k$. Therefore, technically, we need to think of Gen as being given k in unary notation. This is denoted by 1^k .

sk a *secret key* (sk is sometimes also called a *private key*). *Sign* is a signing algorithm: given a message M and a secret key sk it returns a signature x . *Ver* is a verification algorithm: given a public key pk , a message M , and a purported signature x , it outputs “accept” or “reject.”

A signature x on a message M is called *valid* with respect to a public key pk and an oracle H if $Ver^H(pk, M, x) = \text{“accept.”}$ The only required relationship between the three algorithms is that the algorithm *Sign* output valid signatures: that is, for any oracle H and message M , if $(pk, sk) \leftarrow Gen^H(1^k)$ and $x \leftarrow Sign^H(sk, M)$, then x is a valid signature on M with respect to pk and H .

Signature schemes with more than one security parameter can be defined similarly.

Note that in the above definition, we do not specify what the oracle H is or what it means for a signature scheme to be secure. This is discussed below. Our definition of security is a modified version of that in [BR2], which is based on [BR1] and [GMR]. This definition concerns itself with exact, rather than asymptotic, security.

Intuitively, we want to capture the following in our definition of security: there is no algorithm (called “forger”) that, for a random oracle H , is able to produce new valid signatures with reasonable probability in reasonable time without knowing sk . Moreover, we should assume that an attacker can coerce the signer into signing some number of messages of the attacker’s choice—to carry out the so-called “adaptive chosen-message attack” [GMR]. We model this by giving the forger oracle access to the oracle H and to the algorithm $Sign^H(sk, \cdot)$.

Definition 2. A forger $F^{?,?}$ is a probabilistic two-oracle algorithm that is given a security parameter k and a public key pk as input. The first oracle of F is called a *hashing oracle* and the second oracle is called a *signature oracle*. Let H be a hashing oracle, and let $(pk, sk) = Gen^H(1^k)$ for some k . We say that the forger *succeeds* if $(M, x) \leftarrow F^{H, Sign^H(sk, \cdot)}(1^k, pk)$ and x is a valid signature on M with respect to pk and H , and F did not query its signature oracle on M .

We say that a forger $(t, q_{sig}, q_{hash}, \varepsilon, \delta)$ -breaks the signature scheme if, for a security parameter k , the following hold:

- its running time (plus the size of its description) does not exceed $t(k)$,
- the number of its queries to the signature oracle does not exceed $q_{sig}(k)$,
- the number of its queries to the hashing oracle does not exceed $q_{hash}(k)$,
- with probability at least $\delta(k)$, $Gen^H(1^k)$ generates such a key (pk, sk) that the probability of the forger’s success on input $(1^k, pk)$ is at least $\varepsilon(k)$ (here, the probability of the forger’s success is taken over a random choice of the oracle H , the random tape of the forger, the random tape of the signer to whom the forger addresses the chosen-message queries, but not the choice of pk).

Finally, we say that a signature scheme is $(t, q_{sig}, q_{hash}, \varepsilon, \delta)$ -secure if no forger $(t, q_{sig}, q_{hash}, \varepsilon, \delta)$ -breaks it.

(As an aside for the reader familiar with the definition of [BR2], we point out that if a scheme is $(t, q_{sig}, q_{hash}, \varepsilon\delta)$ -secure in the sense of the [BR2], then it is $(t, q_{sig}, q_{hash}, \varepsilon, \delta)$ -secure in the sense of the above definition. We simply separate the component of the

probability that is due to the selection of the public key. This separation allows us to correct some minor errors in the security analysis present in the prior literature. See the proof of Theorem 1 for details.)

Measuring signature scheme security. Now that we have defined what it means for a signature scheme to be secure, how do we actually prove anything about security? To do so, we relate the security of a signature scheme to the difficulty of some problem; in our case, the difficulty of factoring.

Definition 3. Let $Gen(1^l)$ be an algorithm generating l -bit products of two primes. We say that an algorithm A (t, ε, δ) -factors integers generated by Gen if, for a given parameter l ,

- A 's running time (plus the size of its description) does not exceed $t(l)$,
- with probability at least $\delta(l)$, $Gen(1^l)$ generates such an integer n that A has at least $\varepsilon(l)$ probability (taken over only the random choices of the algorithm, not the choice of n) of producing the correct factors of n on input n .

We say that factoring integers generated by Gen is (t, ε, δ) -secure if no such A exists.

Given this definition of the difficulty of a problem, we can then explain the security of a signature scheme Π in the following terms, as suggested in [BR2]: if some problem is $(t', \varepsilon', \delta')$ -secure, then the Π scheme is $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -secure. If t is not much smaller than t' and ε and δ are not much larger than ε' and δ' , even for a reasonably large q_{sig} and q_{hash} , then the reduction proving the security is called *tight*.

3. Micali's Signature Algorithm (MSA)

3.1. Signature and Verification Algorithms

We describe the following ID and signature scheme from [Mi], with similarities to the Ong–Schnorr [OS] and the Guillou–Quisquater [GQ] schemes.

Number theory. Let k and l be two security parameters. Let $p_1 \equiv 3 \pmod{8}$ and $p_2 \equiv 7 \pmod{8}$ be two primes of approximately equal size and let $n = p_1 p_2$ be an l -bit integer (such an n is called a *Williams integer* [Wi]). To simplify further computations, we assume not only that $n > 2^{l-1}$, but also that $|Z_n^*| = n - p_1 - p_2 + 1 > 2^{l-1}$, and that $p_1 + p_2 - 1 < 2^{l/2+1}$. Let Q denote the set of nonzero quadratic residues modulo n . Note that $|Q| > 2^{l-3}$. Note also that for $x \in Q$, exactly one of its four square roots is also in Q (this follows from the fact that -1 is a nonsquare modulo p_1 and p_2 and the Chinese remainder theorem). Thus, squaring is a permutation over Q . From now on, when we speak of “the square root of x ,” we mean the single square root in Q ; by $x^{2^{-k}}$ we denote the single $y \in Q$ such that $x = y^{2^k}$. Also note that 2 is a nonsquare modulo p_1 and a square modulo p_2 (because $(2/p) = (-1)^{(p^2-1)/8}$), so $1 \in Q$ and $-1, 2, -2 \notin Q$. In general, for any $x \in Z_n^*$, exactly one of $x, -x, 2x, -2x$ is in Q .

Following [GMR], define $F_0(x) = x^2 \bmod n$, $F_1(x) = 4x^2 \bmod n$, and, for an m -bit binary string $\sigma = b_1 \cdots b_m$, define $F_\sigma : Q \rightarrow Q$ as $F_\sigma(x) = F_{b_m}(\cdots (F_{b_2}(F_{b_1}(x))) \cdots)$

$= x^{2^m} 4^\sigma \bmod n$ (note that 4^σ is a slight abuse of notation, because σ is a binary string, rather than an integer; what is really meant here is 4 raised to the power of the integer represented in binary by σ). Because squaring is a permutation over Q and $4 \in Q$, F_σ is a permutation over Q .

Note that $F_\sigma(x)$ can be efficiently computed by anybody who knows n . Also, if one knows p_1 and p_2 , one can efficiently compute $x = F_\sigma^{-1}(y)$ (as shown by Goldreich in [Gol]) by computing $s = 1/4^{2^{-|\sigma|}} \bmod n$ and then letting $x = y^{2^{-|\sigma|}} s^\sigma \bmod n$ (these calculations can be done modulo p_1 and p_2 separately, and the results combined using the Chinese remainder theorem). However, if one does not know p_1 and p_2 , then F_σ^{-1} is hard to compute, as shown in the lemma below.

Lemma 1. *If one can compute, for a given $y \in Q$ and two different strings σ and τ of equal length, $x_1 = F_\sigma^{-1}(y)$ and $x_2 = F_\tau^{-1}(y)$, then one can factor n .*

Proof. The proof is by induction on the length of the strings σ and τ .

If $|\sigma| = |\tau| = 1$, then assume, without loss of generality, that $\sigma = 0$ and $\tau = 1$. Then $F_0(x_1) \equiv F_1(x_2) \equiv y \bmod n$, i.e., $x_1^2 \equiv 4x_2^2 \bmod n$, i.e., $n \mid (x_1 - 2x_2)(x_1 + 2x_2)$. Note that $x_1, x_2 \in Q$ and $\pm 2 \notin Q$, so $\pm 2x_2 \notin Q$, so $x_1 \not\equiv \pm 2x_2 \pmod{n}$, so n does not divide either $x_1 - 2x_2$ or $x_1 + 2x_2$. Thus, by computing the gcd of $x_1 + 2x_2$ and n , we can get either p_1 or p_2 .

For the inductive case, let σ and τ be two strings of length $m + 1$. Let σ' and τ' be their m -bit prefixes, respectively. If $F_{\sigma'}(x_1) \equiv F_{\tau'}(x_2) \pmod{n}$, we are done by the inductive hypothesis. Otherwise, the last bit of σ must be different from the last bit of τ , so, without loss of generality, assume the last bit of σ is 0 and the last bit of τ is 1. Then $F_0(F_{\sigma'}(x_1)) \equiv F_1(F_{\tau'}(x_2)) \pmod{n}$, and the same proof as for the base case works here. \square

The ID scheme. The above lemma naturally suggests the following ID scheme. A user has n as the public key and p_1, p_2 as the secret key. To prove his identity (i.e., that he knows p_1 and p_2) to a verifier, he commits to random $X \in Q$ and sends it to the verifier. The verifier produces a random k -bit challenge σ and sends it to the prover (the user). The prover responds with $z = F_{0\sigma}^{-1}(X)$ (note that σ here is prefixed with a single 0 bit, whose use will be explained shortly). The verifier checks that $X = F_{0\sigma}(z) = F_\sigma(z^2)$ and that $z \not\equiv 0 \pmod{n}$. Informally, the security of this protocol is based on the fact that if the prover is able to respond to two different challenges σ and τ , then, by Lemma 1, he knows p_1 and p_2 . The 0 bit in front of σ is to save the verifier from having to check that the prover's response is in Q (which is a hard problem in itself)—instead, she just squares the prover's response and thus puts it in Q .

We say no more about the security of this ID scheme because we are not concerned with it in this paper. We, however, point out an efficiency improvement for the prover. First, as part of key generation, the prover computes, using the Chinese remainder theorem, $s = 1/4^{2^{k-1}} \bmod n$. Then, when committing to a random $X \in Q$, the prover randomly selects an $x \in Z_n^*$ and sets $X = x^{2^{k+1}} \bmod n$ (note that X gets selected with uniform distribution as long as x is selected thus). Now, to respond to a challenge σ , the prover simply computes $z = xs^\sigma \bmod n$.

MSA. The standard way to change the above ID scheme into a signature scheme is to replace the verifier with a random function $H : \{0, 1\}^* \rightarrow \{0, 1\}^k$. The exact steps of the algorithms *Gen*, *Sign*, and *Ver* follow.

Key Generation

1. Generate two random primes $p_1 \equiv 3 \pmod{8}$ and $p_2 \equiv 7 \pmod{8}$ and $n = p_1 p_2$ so that $n < 2^l$, $n - p_1 - p_2 + 1 > 2^l + 1$ and $p_1 + p_2 - 1 < 2^{l/2+1}$.
2. Generate coefficient $c = p_2^{-1} \pmod{p_1}$ for use in the Chinese remainder theorem.
3. Compute $u_i = ((p_i + 1)/4)^{k+1} \pmod{(p_i - 1)/2}$ for $i = 1, 2$ (note that u_i is such that raising a square to the power u_i modulo p_i will compute its 2^{k+1} root).
4. Compute $s_i = ((p_i + 1)/4)^{u_i} \pmod{p_i}$ for $i = 1, 2$ (this computes $1/4^{2^{-(k+1)}} \pmod{p_i}$).
5. Compute $v = (s_1 - s_2)c \pmod{p_1}$ and $s = s_2 + vp_2$ to get $s = 1/4^{2^{-(k+1)}} \pmod{n}$.
6. Output n as the public key and (n, s) as the secret key (to make signing more efficient, p_1, p_2, s_1, s_2 and c should be added to the secret key, in which case s is not needed and the previous step can be omitted).

Signing

1. Generate X by picking a random $x \in \mathbb{Z}_n^*$ and computing $X = x^{2^{k+1}} \pmod{n}$ (note that this step can be done off-line, before the message is known).
2. Compute $\sigma = H(X, M)$, $z = F_{0\sigma}^{-1}(X)$ via $t = s^\sigma \pmod{n}$ (this can be done via $t_i = s_i^\sigma \pmod{p_i}$ for $i = 1, 2$, $v = (t_1 - t_2)c \pmod{p_1}$, $t = t_2 + vp_2$), and $z = xt \pmod{n}$.
3. Output (z, σ) .

Verifying

1. Verify that $z \not\equiv 0 \pmod{n}$ and compute $X = F_\sigma(z^2)$ via $t_1 = z^{2^{k+1}} \pmod{n}$, and $t_2 = 2^\sigma \pmod{n}$, and $X = t_1 t_2 \pmod{n}$.
2. Verify if $\sigma = H(X, M)$.

3.2. Security of MSA

Theorem 1. *If there exists a forger that $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -breaks MSA with security parameters l and k , such that $\varepsilon > 2^{-k+1}(q_{\text{hash}} + 1) + \gamma$, then there exists an algorithm that $(t', \varepsilon', \delta)$ -factors integers generated by Gen for*

$$t' = \frac{(2q_{\text{hash}} + 3)(t + q_{\text{sig}}T_1)}{\varepsilon - \gamma - 2^{-k+1}(q_{\text{hash}} + 1)} + T_2,$$

$$\varepsilon' = \frac{1}{2} \left(1 - \frac{1}{e}\right)^2 > 0.199,$$

where T_1 is the time required to perform an MSA signature verification, T_2 is the time required to factor n given the conditions of Lemma 1 (essentially, a gcd computation), and $\gamma = 2^{-l+3}q_{\text{sig}}(q_{\text{hash}} + 1)$ (note that γ is close to 0 for a large enough l).

Proof. Let F be a forger that $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -breaks MSA. We will construct a factoring algorithm A that uses F to produce $y, z \in Z_n^*$ and $\sigma \neq \tau \in \{0, 1\}^k$ such that $F_\sigma(z^2) = F_\tau(y^2)$. This will allow A to factor n by using the method given in the proof of Lemma 1.

The main idea of this proof is given by the “forking lemma” of [PS1]. It is to allow F to run once to produce one forgery—a signature (z, σ) on a message M such that $\sigma = H(X, M)$ where $X = F_\sigma(z^2)$. Note that F had to ask a hashing-oracle query on (X, M) —otherwise its probability of success is at most 2^{-k} . Then run F the second time, giving the same answers to all the oracle queries before the query (X, M) . For (X, M) give a new answer τ . Then, if F again forges a signature (y, τ) using X and M , we will have achieved our goal.

Assuming n is such that F has probability at least ε of success, the probability that A will factor n using this approach is roughly $\varepsilon^2/q_{\text{hash}}$, because F needs to succeed twice and we have no guarantee that F will choose to use (X, M) for its second forgery and not any of its other q_{hash} oracle queries.

To increase the probability of success to a constant, we iterate A sufficiently many times, as suggested in [OO2] and [PS2]. More specifically, A will run F about $1/\varepsilon$ times the first time, to achieve a constant probability of a successful forgery, and about $2q_{\text{hash}}/\varepsilon$ times the second time, to achieve a constant probability of a successful forgery that uses the pair (X, M) .

We do not provide the details of our construction or the probability analysis, because they are quite similar to those of [OO2] and [PS2]. We wish to point out one crucial difference, however: both [OO2] and [PS2] assume, in their probability analyses, that the $1/\varepsilon$ first runs of F are entirely independent. This is incorrect, because the public key input to F is always the same. Indeed, if F is such that it never succeeds on a certain public key, then running it repeatedly will not increase the success probability. Therefore, it is necessary to separate the component of the probability that is due to the selection of the public key, as we do in our definition of security. \square

Corollary 1. *If factoring l -bit integers generated by Gen is $(t', 0.199, \delta)$ -secure, then MSA is $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -secure for*

$$t = \frac{(t' - T_2)\varepsilon}{4q_{\text{hash}} + 6} - q_{\text{sig}}T_1$$

as long as

$$\varepsilon \geq 2(2^{-k+1}(q_{\text{hash}} + 1) + 2^{-l+3}q_{\text{sig}}(q_{\text{hash}} + 1)).$$

Proof. Solving for t the equation for t' of Theorem 1, we obtain

$$t = \frac{(t' - T_2)(\varepsilon - \gamma - 2^{-k+1}(q_{\text{hash}} + 1))}{2q_{\text{hash}} + 3} - q_{\text{sig}}T_1.$$

The condition on ε ensures that $\varepsilon - \gamma - 2^{-k+1}(q_{\text{hash}} + 1) \geq \varepsilon/2$. Observe that we are allowed to decrease t , as this will only weaken the result, so we are allowed to replace $\varepsilon - \gamma - 2^{-k+1}(q_{\text{hash}} + 1)$ by $\varepsilon/2$.

Note also that Theorem 1 requires that $\varepsilon \geq 2^{-k+1}(q_{\text{hash}} + 1) + \gamma$, which is also assured here by the conditions on ε . \square

4. The Swap Method

4.1. Motivation

As exemplified by the proof of Theorem 1 above, all known results for the security of Fiat–Shamir-like signature schemes involve losing a factor of q_{hash} (in either time or success probability) in the reduction from a forger to an algorithm that breaks the underlying hard problem (see, for example, [FS], [Sc2], [PS1], [Sh], [OO2], and [PS2]). While no proof exists that the loss of this factor is necessary, the problem seems inherent in the way signature schemes are constructed from ID schemes, as explained below.

The security of an ID scheme usually relies on the fact that a prover would be unable to answer two different challenges for the same commitment without knowing the private key. Therefore, in the proof of security of the corresponding signature scheme, we need to use the forger to get two signatures on the same commitment, as we did in the proof of Theorem 1. The forger, however, has any of its q_{hash} queries to pick for the commitment for the second signature—hence, our loss of the factor of q_{hash} . We want to point out that q_{hash} is a significant factor, and its loss definitely makes a reduction quite loose. This is because a reasonable bound on the number of possible hash queries of committed adversaries is about $q_{\text{hash}} = 2^{80}$ (see Section 4.4).

We therefore devise a new method of constructing signature schemes from ID schemes so that any one signature from the forger is enough to break the underlying hard problem.

4.2. Method

Recall that in Fiat–Shamir-like signature schemes, the signer comes up with the commitment and then uses H applied to the commitment and the message to produce the challenge. We propose that instead the signer *first come up with the challenge and then use H applied to the challenge and the message to produce the commitment*. In a way, we swap the challenge and the commitment.

This method applies whenever the signer can compute the response given only the challenge and the commitment. It does not apply when information used during the generation of the commitment is necessary to compute the response. For example, it does not apply to discrete-logarithm-based ID schemes (such as the Schnorr scheme [Sc1]) in which the prover needs to know the discrete logarithm of the commitment in order to provide the response.

Additionally, in order to use this method, one needs to get around the problem that the commitment is selected from some structured set (such as Q in the case of MSA), while H returns a random binary string. This problem can usually be easily solved. The only case known to us when it seems to present a real obstacle is in the scheme of Ohta and Okamoto [OO2] in the case when an exponent L is used such that $\gcd(L, (p_1 - 1)(p_2 - 1)) > 2$.

The key-generation algorithm and the private key may need to be modified slightly in order to provide the signer with the additional information needed to compute the response from a random commitment, rather than from a commitment that it generated. The verification algorithm remains vastly unchanged.

In the next section we exemplify our proposed method and explain why it results in a tighter security reduction.

4.3. MSA-swap

4.3.1. Description

The scheme depends on two security parameters: k and l . Let $H : \{0, 1\}^* \rightarrow \{0, 1\}^{l-1}$ be a random function.

Key Generation. The key generation is the same as in MSA, except for one additional step (step 6) and extra information in the private key:

1. Generate two random primes $p_1 \equiv 3 \pmod{8}$ and $p_2 \equiv 7 \pmod{8}$ and $n = p_1 p_2$ so that $n < 2^l$, $n - p_1 - p_2 + 1 > 2^l + 1$, and $p_1 + p_2 - 1 < 2^{l/2+1}$.
2. Generate coefficient $c = p_2^{-1} \pmod{p_1}$ for use in the Chinese remainder theorem.
3. Compute $u_i = ((p_i + 1)/4)^{k+1} \pmod{(p_i - 1)/2}$ for $i = 1, 2$ (note that u_i is such that raising a square to the power u_i modulo p_i will compute its 2^{k+1} root).
4. Compute $s_i = ((p_i + 1)/4)^{u_i} \pmod{p_i}$ for $i = 1, 2$ (this computes $1/4^{2^{-(k+1)}} \pmod{p_i}$).
5. Compute $v = (s_1 - s_2)c \pmod{p_1}$ and $s = s_2 + vp_2$ to get $s = 1/4^{2^{-(k+1)}} \pmod{n}$.
6. If u_i is odd, make it even by setting $u_i = u_i + (p_i - 1)/2$ for $i = 1, 2$ (note that now u_i is such that raising a square or its negative to the power u_i modulo p_i will compute its 2^{k+1} root).
7. Output n as the public key and $(n, s, u_1, u_2, p_1, p_2, c)$ as the secret key.

Signing

1. Generate a random σ and compute $t = s^\sigma \pmod{n}$ (note that this step can be done off-line, before the message is known).
2. Compute $X = H(\sigma, M)$. If $X \notin Z_n^*$, then try another σ (the probability of that is at most $2^{-l/2+2}$). If the Jacobi symbol $(X/n) = -1$, set $X = 2X \pmod{n}$. Now either X or $n - X$ is in Q . Compute $z = F_0 \sigma^{-1}(\pm X)$ via $x_i = X^{u_i} \pmod{p_i}$ for $i = 1, 2$, $v = (x_1 - x_2)c \pmod{p_1}$, $x = x_2 + vp_2$, and $z = xt \pmod{n}$.
3. Output (z, σ) .

Verifying

1. Verify that $z \not\equiv 0 \pmod{n}$ and compute $X = F_\sigma(z^2)$ via $t_1 = z^{2^{k+1}} \pmod{n}$, $t_2 = 2^\sigma \pmod{n}$, and $X = t_1 t_2 \pmod{n}$ (this step is the same as for MSA).
2. Let $X' = H(\sigma, M)$. If $X \equiv \pm X' \pmod{n}$ or $X \equiv \pm 2X' \pmod{n}$, accept the signature (this step differs slightly from MSA).

4.3.2. Security of MSA-Swap

Proposition 1. *If there exists a forger that $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -breaks MSA-swap with security parameters l and k , then there exists an algorithm that $(t', \varepsilon', \delta)$ -factors integers generated by Gen for*

$$\begin{aligned} t' &= t + 2(q_{\text{sig}} + q_{\text{hash}} + 1)T_1 + T_2, \\ \varepsilon' &= \varepsilon - \gamma, \end{aligned}$$

where T_1 is the time required to perform an MSA-swap signature verification, T_2 is the time required to factor n given the conditions of Lemma 1 (essentially, a gcd computation), and $\gamma = 2^{-k}q_{\text{sig}}(q_{\text{hash}} + 1) + 2^{-l/2+2}(q_{\text{hash}} + q_{\text{sig}} + 1)$ (note that γ is close to 0 for large enough k and l).

Proof.

Main idea. Let F be a forger that $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -breaks MSA-swap. Similarly to the proof of Theorem 1, we will construct an algorithm that, after interacting with F , will produce $y, z \in Z_n^*$ and $\sigma \neq \tau \in \{0, 1\}^k$ such that $F_\sigma(z^2) = F_\tau(y^2)$.

The main idea is to answer each hash query on (σ, M) with an X computed via $X = F_\tau(y^2)$ for a random $y \in Z_n^*$ and arbitrary τ that is different from σ . Then if F forges a signature (z, σ) on M , we will have $F_\sigma(z^2) = F_\tau(y^2)$ and will be able to factor n .

Details. We assume that F performs the necessary bookkeeping and does not ask the same hash query twice.² Note that F may ask the same signature query twice, because the answers will most likely be different. We also modify F so that it will not output a signature (z, σ) on a message M unless it knows that the signature is correct; in particular, that means that F first has to ask, in a hash query, for the value of $H(\sigma, M)$. This may increase the number of hash queries by one and the running time of F by the cost of one signature verification.

A maintains two tables, for signature query answers and for hash query answers. To answer the j th signature query on a message M'_j , A picks a random $\tau'_j \in \{0, 1\}^k$ and checks its signature query table to see if a signature query on M'_j has already been asked and if τ'_j was used in answering it. If so, it returns the answer from the table. If not, it picks a random $y'_j \in Q$ and a random $\mu'_j \in \{1, -1, 2, -2\}$, computes $X'_j = F_{\tau'_j}(y'^2_j)/\mu'_j \bmod n$, and checks if $X'_j < 2^{l-1}$. If not, it restarts with new random y'_j and μ'_j . The expected number of trials is less than two, because X'_j is a random element in Z_n^* (because $F_{\tau'_j}(y'^2_j)$ is a random element in Q and μ'_j is also random) and $|Z_n^*| < n < 2^l$. Once X'_j satisfies the desired condition, A outputs (y'_j, τ'_j) as the signature and records $(j, y'_j, \tau'_j, X'_j, M'_j)$ in its signature query table.

To answer the i th hash query on (σ_i, M_i) , A first checks if an entry $(j, y'_j, \tau'_j, X'_j, M'_j)$ such that $\sigma_i = \tau'_j$ and $M_i = M'_j$ already exists in its signature table. If so, it returns X'_j .

² This may slightly increase the running time of F , but we ignore costs of simple table look-up for the purposes of this analysis.

Otherwise, it picks $\tau_i \neq \sigma_i$, computes y_i, μ_i , and X_i in the same way as when answering a signature query, outputs X_i , and adds $(i, y_i, \tau_i, \sigma_i, M_i)$ to its hash query table.

Assume F outputs a forgery (z, σ) on a message M . A then computes $X = F_\sigma(z^2)$ and searches its hash query table for such an i that $\sigma_i = \sigma$ and $M_i = M$. (Such an i has to exist, because we modified F to verify the validity of its forgery, so F had to ask a hash query on (σ, M) . Note that there would be no entry in the hash query table if the value of $H(\sigma, M)$ was first obtained by a signature query; however, this is impossible, because F is prohibited from outputting a forgery on a message for which it asked a signature query.) Then $F_\sigma(z^2) = X = \mu_i X_i = F_{\tau_i}(y_i^2)$, and $\sigma \neq \tau_i$ because of the way τ_i was picked, so A can easily factor n by the method given in the proof of Lemma 1.

Probability analysis. For this analysis, we assume that n is such that F has probability at least ε of success. By assumption, such an n is generated with probability at least δ .

We need to compare the distribution of A 's answers to F 's queries to the distribution of answers given by the true oracles. The distributions differ as follows: if, for some signature query j , the hash value of τ'_j, M'_j has already been defined through a previous hash query, then A will fail, whereas the true oracles would not. The probability of this event for a given signature query is at most $(q_{\text{hash}} + 1)2^{-k}$, because τ'_j is picked at random from $\{0, 1\}^k$. Thus, the probability of F 's success is reduced as a result of this by at most $q_{\text{sig}}(q_{\text{hash}} + 1)2^{-k}$. Additionally, X in the answers to hash and signature queries is distributed uniformly over $Z_n^* \cap \{0, 1\}^{l-1}$, thus excluding the multiples of p_1 and p_2 that a real hash function could possibly hit. Since there are at most $p_1 + p_2 - 1 < 2^{l/2+1}$ such multiples in Z_n , this will reduce the forger's probability of success by at most $(2^{l/2+1})/(2^{l-1}) = 2^{-l/2+2}$ per query.

So the probability of successfully factoring n is

$$\varepsilon' \geq \varepsilon - q_{\text{sig}}(q_{\text{hash}} + 1)2^{-k} - (q_{\text{hash}} + q_{\text{sig}} + 1)2^{-l/2+2}. \quad \square$$

Theorem 2. *If there exists a forger that $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -breaks MSA-swap with security parameters l and k , such that $\varepsilon > \gamma$, then there exists an algorithm that $(t', \varepsilon', \delta)$ -factors integers generated by Gen for*

$$t' = \frac{t + 2(q_{\text{sig}} + q_{\text{hash}} + 1)T_1}{\varepsilon - \gamma} + T_2,$$

$$\varepsilon' = \left(1 - \frac{1}{e}\right) > 0.632,$$

where T_1, T_2 , and γ are as in Proposition 1.

Proof. Let

$$\alpha = \varepsilon - \gamma.$$

By assumption, $\alpha > 0$. So if we repeat the algorithm constructed in the proof of Proposition 1 Theorem 1 up to $1/\alpha$ times (except for the final gcd computation, which need only be done once), we get the desired ε' , similarly to the proof of Theorem 1. \square

Similarly to MSA, we have the following immediate corollary.

Corollary 2. *If factoring l -bit integers generated by Gen is $(t', 0.632, \delta)$ -secure, then MSA-swap is $(t, q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta)$ -secure for*

$$t = \frac{(t' - T_2)\varepsilon}{2} - 2(q_{\text{sig}} + q_{\text{hash}} + 1)T_1$$

as long as

$$\varepsilon \geq 2 \left(2^{-l/2+2} (q_{\text{hash}} + q_{\text{sig}} + 1) + 2^{-k} q_{\text{sig}} (q_{\text{hash}} + 1) \right).$$

Proof. The condition on ε ensures that $\varepsilon - \gamma \geq \varepsilon/2$. The rest follows, similarly to the proof of Corollary 1, from solving the equations of Proposition 1 for t and ε . \square

4.4. Parameter Choice

It is immediately clear from the formulas in Corollaries 1 and 2 that MSA-swap loses no factor of q_{hash} , neither in time nor in probability. This is a big advantage for MSA-swap because q_{hash} can be quite big.

A fuller comparison, provided in the next section, depends on the actual values of the parameters q_{sig} , q_{hash} , k , and l . We deal here, however, with the preliminary problem of assigning reasonable values to these parameters.

We believe it reasonable to set $q_{\text{sig}} = 2^{30}$ and $q_{\text{hash}} = 2^{80} - 1$. This is so because signature queries have to be answered by the honest signer (who may not be willing or able to sign more than a billion messages), while hash queries can be computed by the adversary alone (who may be willing to invest extraordinary resources). Notice that we recommend a higher value for q_{hash} than suggested in [BR2].

We recommend setting $k = 100$ for MSA, and $k = 130$ for MSA-swap. For MSA, this is so because, from Corollary 2, we see that $2^{-k}(q_{\text{hash}} + 1)$ has to be small (the value of $2^{-k}(q_{\text{hash}} + 1)$ is essentially the success probability of the simple attack that relies on correctly guessing one hash value among $q_{\text{hash}} + 1$ hash queries). Therefore, we need 2^{-k+80} to be small, and by setting $k = 100$ we make it less than 10^{-6} . For MSA-swap, this is so because $2^{-k}q_{\text{sig}}(q_{\text{hash}} + 1)$ has to be small, as seen in Corollary 2.

As for l , notice that both MSA and MSA-swap are immediately broken if the adversary succeeds in factoring the l -bit modulus. Therefore, l ought to be at least 1000. Given the above choices for the other parameters, such a minimum value for l is large enough to make all the constraints involving l in Corollaries 1 and 2 satisfied (for any reasonable ε in the case of Corollary 2). Thus, the value of l depends on the presumed security of factoring, as discussed in the next section.

5. Comparing Signature Schemes

Having developed the swap method and demonstrated that it tightens the security of Fiat-Shamir-like signature schemes, we wish to determine precisely when the method should be used. It would be tempting to say “whenever tight security is desired,” but this answer is meaningless. We submit that meaningful comparisons of signature schemes

should take both security and efficiency into account, and that consideration of either factor alone is insufficient.

We also use our methods to compare MSA with PRab from [BR2] and find, perhaps surprisingly, that the former may sometimes provide more security for less cost.

5.1. *The Costs of Security*

The desired level of security is usually dictated by the specific application. It is after settling on the desired amount of security that choosing among the various secure schemes becomes crucial. Indeed, when choosing a signature scheme, the goal is *to maintain the desired level of security at the lowest possible cost*. In a sense, picking a signature scheme is similar to shopping for an insurance policy for the desired face value.

The costs of a signature scheme, however, are quite varied. They may include the sizes of keys and signatures, the efficiencies of generating keys, signing and verifying, the amounts of code required, and even “external” considerations—such as the availability of inexpensive implementations or off-the-shelf hardware components. In this paper we focus on the efficiencies of signing and verifying. These are particularly important when signing or verifying is performed by a low-power device, such as a smart card, or when signing or verifying needs to be performed in bulk quantities, as on a secure server.

It is for these costs, then, that below we compare MSA and MSA-swap. We also provide a comparison of MSA with the PRab scheme from [BR2], arguably the most practical among those *tightly* related to factoring. (The reason for choosing PRab rather than its PSS variant is that the latter is tightly related to RSA, and thus potentially less secure than factoring.)

5.2. *Comparison of MSA and MSA-Swap*

The efficiency of signature verification in MSA is about the same as in MSA-swap. The security of MSA-swap is generally higher than the security of MSA for the same security parameters. Therefore, if the efficiency of verifying is the only significant component in the cost, MSA-swap will be able to provide the same amount of security for less cost than MSA.

A more difficult case to analyze is the case when the efficiency of signing is of main concern. We limit our analysis to the case when we are only concerned with the on-line part of signing. In both cases, this involves mainly a modular exponentiation. Therefore, a variety of sophisticated algebraic methods can be used here, but these methods apply equally to MSA and MSA-swap. We thus find it simpler to compare the two under “standard” implementations using the Chinese remainder theorem (CRT). MSA then involves two exponentiations of a fixed $l/2$ -bit base to a k -bit power, modulo two $l/2$ -bit primes. One modular multiplication of two $l/2$ -bit numbers takes about $l^2/4$ steps; about $3k$ such multiplications are required ($1.5k$ for each of the primes), so the total amount of time required for on-line signing in MSA is about $3kl^2/4$. Similar analysis applies for MSA-swap, except that about $1.5l$ multiplications ($1.5l/2$ for each of the two primes) are required, so the total amount of time required for on-line signing in MSA-swap is about $3l^3/8$, not counting the (relatively small) cost of computing the Jacobi symbol. (In

sum, on-line signing is $l/(2k)$ times faster for MSA than for MSA-swap *if one used the same value of l for both.*³)

We now see how the security of the two schemes compares assuming the on-line signing costs are the same. Let l_M and k_M be the security parameters for MSA, and let l_{MS} and k_{MS} be the security parameters for MSA-swap. The on-line signing costs for MSA and MSA-swap are the same if

$$l_{MS} = (2k_M l_M^2)^{1/3}. \quad (1)$$

The best known factoring algorithms take time about

$$T(l) = C \exp\left(\left(\frac{64}{9}\right)^{1/3} l^{1/3} (\ln l)^{2/3}\right)$$

for some constant C [LL]. Therefore, we assume that factoring l -bit integers generated by Gen is $(C'T(l), 0.199, \delta)$ -secure for some δ and some constant C' . This means that, for a large enough ε and small enough q_{sig} , MSA is about

$$\left(\frac{(C'T(l_M) - T_2(l_M, k_M))\varepsilon}{4q_{\text{hash}} + 6} - q_{\text{sig}} T_1(l_M, k_M), q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta\right)\text{-secure}$$

and MSA-swap is about

$$\left(\frac{(C'T(l_{MS}) - T_2(l_{MS}, k_{MS}))\varepsilon}{2} - 2(q_{\text{sig}} + q_{\text{hash}} + 1)T_1(l_{MS}, k_{MS}), q_{\text{sig}}, q_{\text{hash}}, \varepsilon, \delta\right)\text{-secure}$$

(by Corollaries 1 and 2).

Keeping the signing costs equal, we now find out when MSA becomes more secure than MSA-swap, that is, when

$$\begin{aligned} & \frac{(C'T(l_M) - T_2(l_M, k_M))\varepsilon}{4q_{\text{hash}} + 6} - q_{\text{sig}} T_1(l_M, k_M) \\ & > \frac{(C'T(l_{MS}) - T_2(l_{MS}, k_{MS}))\varepsilon}{2} - 2(q_{\text{sig}} + q_{\text{hash}} + 1)T_1(l_{MS}, k_{MS}). \end{aligned}$$

Given the discussion of Section 4.4 and (1), we now plug in $q_{\text{sig}} = 2^{30}$, $q_{\text{hash}} = 2^{80} - 1$, $k_M = 100$, $k_{MS} = 130$, and $l_{MS} = (2k_M l_M^2)^{1/3}$ to the above inequality. These values allow us to omit the terms in T_1 and T_2 , noting that they decrease the security of MSA-swap much more than they decrease the security of MSA. Thus, we are now interested in when

$$\frac{C'T(l_M)}{2^{82} + 2} > \frac{C'T((200l_M^2)^{1/3})}{2}.$$

³ Moreover, an optimization available to MSA but not to MSA-swap is precomputing some powers of the fixed base; this requires additional memory, so we assume it is not implemented for the purposes of this analysis.

After taking the natural logarithm of both sides, we are left with the inequality

$$\left(\frac{64}{9}\right)^{1/3} l_M^{1/3} (\ln l_M)^{2/3} - \ln(2^{81} + 1) > \left(\frac{64}{9}\right)^{1/3} (200l_M^2)^{1/9} \left(\frac{1}{3} \ln 200 + \frac{2}{3} \ln l_M\right)^{2/3}, \quad (2)$$

which holds as long as long as $l_M \geq 6109$.

Thus, at $l_M = 6109$, $l_{MS} = 1954$, $k_M = 100$, $k_{MS} = 130$, MSA and MSA-swap provide about the same security and the same performance for on-line signing.

In sum. The on-line portion of the signing algorithm of MSA is so fast that

Provable security and signing efficiency are the same when MSA uses 6109-bit moduli and MSA-swap 1954-bit moduli.

In both cases, the security is that of factoring a 1954-bit integer generated by *Gen*. (MSA may actually be even more secure, but we cannot prove it.)

It just so happens that this computed level of security is currently considered adequate for many applications. (Therefore, for these applications MSA-swap is preferable: MSA-swap has faster verification for the same level of security, as well as shorter keys and, therefore, shorter signatures.)

However, whenever the application calls for a *higher* level of security, and the dominant cost is that of signing, then it is the loosely secure MSA that becomes preferable. In fact, inequality (2) implies that the security gap between MSA and MSA-swap, given the same performance, increases exponentially. In essence, this is so because, in order to guarantee the same performance, l_{MS} is a fixed function of l_M : i.e., by (1), $l_{MS} = O(l_M^{2/3})$. Thus security for both MSA and MSA-swap is measured by inequality (2) in terms of l_M , and the dominant factor in the exponent of the adversary's running time is $l_M^{1/3}$ for MSA and $l_M^{2/9}$ for MSA-swap.

5.3. Comparison of MSA with Bellare–Rogaway's PRab

The security of PRab is tightly related to that of modular square roots, rather than factoring. A factor of 2 in probability is lost (compared with MSA-swap) when one relates the security of PRab to that of factoring. PRab's performance for on-line signing is about the same as MSA-swap's (PRab requires a few more Jacobi symbol computations, but no separate modular multiplication).⁴ A vastly similar analysis leads to the following conclusion:

Provable security and signing efficiency are the same when MSA uses 5989-bit moduli, and PRab 1929-bit moduli.

Also here this is a “cross-over” point: the gap in security for the same performance increases exponentially in favor of MSA. As we can see, this cross-over point is just slightly more in favor of MSA than the cross-over point of MSA and MSA-swap. This is because of the factor of 2 difference in the security of MSA-swap and PRab.

⁴ Unlike MSA-swap, however, PRab has no off-line component in signing. It also has very efficient verification.

Acknowledgments

We thank Salil Vadhan for pointing out an error in an earlier version of this work and Mihir Bellare for discussions about the value of the parameter k for MSA.

References

- [BR1] Mihir Bellare and Phillip Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communication Security*, November 1993, pages 62–73. Revised version appears in <http://www-cse.ucsd.edu/users/mihir/papers/crypto-papers.html>.
- [BR2] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: how to sign with RSA and Rabin. In [Ma], pages 399–416. Revised version appears in <http://www-cse.ucsd.edu/users/mihir/paperscrypto-papers.html>.
- [FFS] Uriel Feige, Amos Fiat, and Adi Shamir. Zero-knowledge proofs of identity. *Journal of Cryptology*, 1(2):77–94, 1988.
- [FS] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In [Od], pages 186–194.
- [GMR] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.
- [Gol] Oded Goldreich. Two remarks concerning the Goldwasser–Micali–Rivest signature scheme. In [Od], pages 104–110.
- [Go2] Shafi Goldwasser, editor. *Advances in Cryptology—CRYPTO ’88*, volume 403 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, 1990.
- [GQ] Louis Claude Guillou and Jean-Jacques Quisquater. A “paradoxical” identity-based signature scheme resulting from zero-knowledge. In [Go2], pages 216–231.
- [LL] Arjen K. Lenstra and Hendrik W. Lenstra, editors. *The Development of the Number Field Sieve*, volume 1554 of Lecture notes in Mathematics. Springer-Verlag, Berlin, 1993.
- [Ma] Ueli Maurer, editor. *Advances in Cryptology—EUROCRYPT 96*, volume 1070 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, 1996.
- [Mi] Silvio Micali. A secure and efficient digital signature algorithm. Technical Report MIT/LCS/TM-501, Massachusetts Institute of Technology, Cambridge, MA, March 1994.
- [MR] Silvio Micali and Leonid Reyzin. Improving the exact security of Fiat–Shamir signature schemes. In R. Baumgart, editor, *Secure Networking — CQRE [Secure] ’99*, volume 1740 of Lecture Notes in Computer Science, pages 167–182. Springer-Verlag, Berlin, 1999.
- [MS] Silvio Micali and Adi Shamir. An improvement of the Fiat–Shamir identification and signature scheme. In [Go2], pages 244–247.
- [Od] Andrew M. Odlyzko, editor. *Advances in Cryptology—CRYPTO ’86*, volume 263 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, 1987.
- [Ok] Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In E. F. Brickell, editor, *Advances in Cryptology—CRYPTO ’92*, volume 740 of Lecture Notes in Computer Science, pages 31–53. Springer-Verlag, Berlin, 1993.
- [OO1] Kazuo Ohta and Tatsuaki Okamoto. A modification of the Fiat–Shamir scheme. In [Go2], pages 232–243.
- [OO2] Kazuo Ohta and Tatsuaki Okamoto. On concrete security treatment of signatures derived from identification. In H. Krawczyk, editor, *Advances in Cryptology—CRYPTO ’98*, volume 1462 of Lecture Notes in Computer Science, pages 354–369. Springer-Verlag, Berlin, 1998.
- [OS] H. Ong and Claus P. Schnorr. Fast signature generation with a Fiat–Shamir-like scheme. In I. B. Damgård, editor, *Advances in Cryptology—EUROCRYPT 90*, volume 473 of Lecture Notes in Computer Science, pages 432–440. Springer-Verlag, Berlin, 1991.
- [PS1] David Pointcheval and Jacques Stern. Security proofs for signature schemes. In [Ma], pages 387–398.
- [PS2] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.

- [Sc1] Claus P. Schnorr. Efficient identification and signatures for smart cards. In J.-J. Quisquater and J. Vandewalle, editors, *Advances in Cryptology—EUROCRYPT 89*, volume 434 of Lecture Notes in Computer Science, pages 688–689. Springer-Verlag, Berlin, 1990.
- [Sc2] Claus P. Schnorr. Security of 2^t -root identification and signatures. In N. Koblitz, editor, *Advances in Cryptology—CRYPTO '96*, volume 1109 of Lecture Notes in Computer Science, pages 143–156. Springer-Verlag, Berlin, 1996.
- [Sh] Victor Shoup. On the security of a practical identification scheme. In [Ma], pages 344–353.
- [Wi] Hugh C. Williams. A modification of the RSA public-key encryption procedure. *IEEE Transactions on Information Theory*, IT-26(6):726–729, November 1980.