

Almost uniform density of power residues and the provable security of ESIGN

Tatsuaki Okamoto¹ and Jacques Stern²

¹ NTT Labs, 1-1 Hikarinooka, Yokosuka-shi 239-0847 Japan.

E-mail: okamoto@isl.ntt.co.jp.

² Dépt d'informatique, ENS, 45 rue d'Ulm, 75230 Paris Cedex 05, France.

E-mail: Jacques.Stern@ens.fr

URL: <http://www.di.ens.fr/~stern>.

September 15, 2003

Abstract. ESIGN is an efficient signature scheme that has been proposed in the early nineties (see [14]). Recently, an effort was made to lay ESIGN on firm foundations, using the methodology of provable security. A security proof [15] in the random oracle model, along the lines of [2], appeared in support for ESIGN. However, several unexpected difficulties were found. Firstly, it was observed in [20], that the proof from [15] holds in a more restricted model of security than claimed. Even if it is quite easy to restore the usual security level, as suggested in [9], this shows that the methodology of security proofs is more subtle than it at first appears. Secondly, it was found that the proof needs the additional assumption that e is prime to $\varphi(n)$, thus excluding the case where e is a small power of two, a very attractive parameter choice. The difficulty here lies in the simulation of the random oracle, since it relies on the distribution of e -th powers, which is not completely understood from a mathematical point of view, at least when e is not prime to $\varphi(n)$. In this paper, we prove that the set of e -th power modulo an RSA modulus n , which is a product of two equal size integers p, q , is almost uniformly distributed on any large enough interval. This property allows to complete the security proof of ESIGN. We actually offer two proofs of our result: one is based on two-dimensional lattice reduction, and the other uses Dirichlet characters. Besides yielding better bounds, the latter is one new example of the use of analytic number theory in cryptography.

1 Introduction

Since the appearance of the celebrated RSA cryptosystem [18], a lot of effort has been devoted to finding alternative schemes. In the area of signature, a major challenge is to reduce the computing effort needed from the signer, since it is well known that RSA requires a full-size modular exponentiation. Among the potential candidates to answer this challenge is the ESIGN signature scheme, that has been proposed in the early nineties (see [14]). While RSA generates signatures by computing an e -th root of a hash value, ESIGN only requests to

find an element whose e -th power is close enough to the hash value. Thus, the mathematical assumption underlying ESIGN is that, given an element y of \mathbb{Z}_n^* , it is hard to find x with e -th power lying in an interval with lower endpoint y and length say $n^{2/3}$. This is called the approximate e -th root problem, in short AERP. Combining this relaxed assumption with the use of an RSA modulus of the form $n = p^2q$ allows a very efficient way to sign, with a computing time essentially equivalent to a single exponentiation to the e -th power. This is especially attractive when e is small, and in particular a small power of two.

As most newly proposed cryptosystems, ESIGN has attracted cryptanalytic effort. Papers [3, 21] described several attacks against the underlying problem, for $e = 2, 3$. Still, It is fair to say that there is no known attack against AERP when e is ≥ 4 . Recently, in connection with several standardization efforts such as IEEE P1363, Cryptrec and NESSIE, an effort was made to lay ESIGN on firm foundations, using the methodology of provable security. A security proof in the random oracle model, falong the lines of [2], formally relating the security of ESIGN with the AERP problem, appeared in [15]. However, several unexpected difficulties were found. Firstly, it was observed in [20] that the proof from [15] holds in a more restricted model of security than claimed: this model, termed single occurrence chosen message attack **SO-CMA** is very similar to the usual chosen message attack scenario but does not allow the adversary to submit the same message twice for signature. This observation does not endanger the scheme in any way, and furthermore, it is quite easy to restore the usual **CMA** security, as suggested in [9]. Still, it shows that the methodology of security proofs is more subtle than it at first appears, a fact already pointed out by Shoup [19], in the context of public key encryption. Secondly, it was found that the proof needs the additional assumption that e is prime to $\varphi(n)$, thus excluding some very attractive parameter choices, notably powers of two. The difficulty here lies in the simulation of the random oracle, since it relies on the distribution of e -th powers, which is not completely understood from a mathematical point of view. In this paper, we prove that the set of e -th power modulo an RSA modulus n , which is a product of two equal size integers p, q , is almost uniformly distributed on any large enough interval. In other words, the number of e -th powers modulo n in any interval of large enough length n^δ is close to $\frac{n^\delta}{d} \frac{\varphi(n)}{n}$, where d is the number of e -th roots of unity modulo n . We actually offer two proofs of our result. The first proof relies on methods from the geometry of numbers and uses two-dimensional lattices. The second proof borrows from analytic number theory and uses Dirichlet characters and the Polya-Vinogradov inequality. Both proofs yield concrete estimates, which are enough to complete the security proof of ESIGN. Although the estimates in the second proof are sharper, we have found interesting to include the two methods, which are of independent interest.

Removing the restriction that e is prime to $\varphi(n)$ may appear a side issue. However, we believe that it is important both for practical and methodological reasons. As already noted, ESIGN has has a very fast algorithm for signature generation, since its main step is a single exponentiation to the e -th power. Making e a power of two is the best way to take advantage of this feature

and should be allowed by the security proof. Also, as shown by various results, notably [19, 20], provable security has many subtleties. In the present paper, the subtlety lies in the simulation of the random oracle. As far as we know, this is the only example where this part is not straightforward, and the underlying difficulty may easily be overlooked. In other words, it may appear obvious that picking x at random and suitably truncating $x^e \bmod n$ simulates a random oracle, which is the main result of our paper. However, it is not, at least when e is not prime to $\varphi(n)$ and it is actually related with deep mathematical questions of analytical number theory.

Our paper is organized as follows: we first recall some preliminaries from number theory. Next, we present the two proofs. Finally, we produce a proof of security for ESIGN, not using the assumption that e is prime to $\varphi(n)$. In this proof, we focus on the simulation of the random oracle, and explain where our result on power residues is needed.

2 Number theoretic preliminaries

2.1 Lattices

Let n be an RSA modulus. For any integer α , we consider the lattice

$$L(\alpha) = \{(x, y) \in \mathbb{Z}^2 \mid x - \alpha y = 0 \pmod{n}\}.$$

We note that $L(\alpha)$ is a two-dimensional lattice with determinant n . Thus, its shortest vector should be of euclidean norm of the order \sqrt{n} . It can be obtained by applying the Gaussian reduction algorithm. This algorithm outputs within time $\mathcal{O}((\log n)^3)$ a basis of $L(\alpha)$ consisting of two non-zero vectors $U(\alpha)$ and $V(\alpha)$ such that

$$\|U\| \leq \|V\| \text{ and } |(U, V)| \leq \|U\|^2/2,$$

where we have omitted α for clarity. From a geometrical point of view, the inequalities imply that the angle θ of U and V is such that $|\cos \theta| \leq 1/2$, hence $|\sin \theta| \geq \sqrt{3}/2$, and therefore

$$|U \wedge V| = n \geq \frac{\sqrt{3}|U||V|}{2}$$

We say that $L(\alpha)$ is an ε -good lattice if $|U|$ is bounded from below by $n^{1/2-\varepsilon}$. Note that, for such a lattice, we have

$$|V| \leq \frac{2}{\sqrt{3}}n^{1/2+\varepsilon}.$$

Lemma 1. *The number of elements α in \mathbb{Z}_n such that $L(\alpha)$ is not an ε -good lattice is at most $4n^{1-2\varepsilon}$.*

Proof. This follows from the fact that the shortest non zero vector of a lattice $L(\alpha)$ which is not ε -good lies in the disk centered at the origin, with radius $n^{1/2-\varepsilon}$. This number of integers in this disk is bounded by $4n^{1-2\varepsilon}$. To conclude, it is enough to observe that an element (x, y) of the disk other than $(0, 0)$ cannot belong to two distinct $L(\alpha)$ lattices, unless y is not in \mathbb{Z}_n^* , which cannot happen since n is an RSA integer, i.e. has two prime factors of almost equal size.

We let P be the parallelepiped spanned by U and V .

Lemma 2. *Let $L(\alpha)$ be ε -good. The width of P is at most $2n^{1/2+\varepsilon}$.*

Proof. The square of the width is indeed bounded by

$$|U|^2 + |V|^2 + 2|(U, V)| \leq 2|V|^2 + |V|^2 \leq 3|V|^2,$$

which is bounded by $4n^{1+2\varepsilon}$. The lemma follows.

Lemma 3. *Let $L(\alpha)$ be ε -good. Let I be an interval of length n^δ , with $\delta > 1/2$. The square $I \times I$ has at most $(n^{\delta-1/2} + 2n^\varepsilon)^2$ elements in $L(\alpha)$.*

Proof. let \tilde{P} be obtained by translating P by $-\frac{u+v}{2}$. We consider the set X of lattice points M such that the parallelepiped $M + \tilde{P}$ meets $I \times I$. The number of such points is clearly an upper bound for the number of lattice points inside $I \times I$. Now, the various parallelepiped $M + \tilde{P}$ are pairwise disjoint and, by lemma 2, they are contained in the square $J \times J$, obtained by enlarging I by $n^{1/2+\varepsilon}$ on each side. Summing up the areas of the individual cells, we get:

$$n|X| \leq (n^\delta + 2n^{1/2+\varepsilon})^2.$$

which provides the desired bound on the number $|X|$ of elements of X .

When $L(\alpha)$ is not ε -good, we can show a weaker bound:

Lemma 4. *Let α be any integer. Let I be an interval of length n^δ , with $\delta < 1$. The square $I \times I$ has at most $n^\delta + 1$ elements in $L(\alpha)$.*

Proof. For fixed y , there is at most one pair (x, y) such that $x - \alpha y = 0 \pmod n$ in any interval of length $< n$, such as I . This provides the requested bound $n^\delta + 1$.

2.2 Dirichlet characters

Let G be a finite (multiplicative) abelian group. A character χ over G is a multiplicative homomorphism from G into the multiplicative group of complex numbers. The set of characters over G is a group, called the dual of G and denoted \hat{G} . Its unit χ_0 is the *principal character*, defined by $\chi_0(g) = 1$, for any $g \in G$.

The following is well-known (see [6], chapter 7):

Theorem 1. i) There are exactly $|G|$ characters over G .
ii) For any $g \neq 1$, the following holds:

$$\sum_{\chi \in \hat{G}} \chi(g) = 0$$

iii) For any $\chi \neq \chi_0$, the following holds:

$$\sum_{g \in G} \chi(g) = 0$$

A Dirichlet character χ is a character over \mathbb{Z}_n^* , for some integer n . The characters can be extended to all integers by using the value 0 at integers not invertible mod n . In the sequel, we will need a bound on the sum of such characters over large intervals. This is given by the Polya-Vinogradov inequality (see [5] or [6], chapter 9):

Theorem 2. For any non principal Dirichlet character χ over \mathbb{Z}_n^* and any integer h , the following holds:

$$\left| \sum_{x=1}^h \chi(x) \right| \leq 2\sqrt{n} \ln n.$$

Remark. When n is a prime number p , and, more generally when χ is a so-called primitive character, the multiplicative constant 2 in the above can be replaced by 1. We will not need such refinement.

3 Almost uniform density of e -th powers

We now turn to our main result. We first review the standard situation of an RSA exponent.

3.1 The case where e is prime to $\varphi(n)$

Lemma 5. Let n be an RSA modulus and e be an integer prime to $\varphi(n)$. Let I be an interval of length n^δ , with $\delta < 1$. The number of integers from I which are e -th powers of an element of \mathbb{Z}_n^* differs from $n^\delta \frac{\varphi(n)}{n}$ by at most 4.

Proof. Since exponentiation to the e -th power is one-to-one, we have to count the number of elements in $I \cap \mathbb{Z}_n^*$. The number of multiples of p in I differs from $\frac{n^\delta}{p}$ by at most one. Similarly for q . Since there may be one multiple of pq , the final count is almost K , where

$$K = n^\delta \frac{\varphi(n)}{n}$$

and the difference with K is bounded by $3 + \frac{n^\delta}{n} \leq 4$.

We now turn to the general case. Observe that the set of e -th powers is a subgroup of \mathbb{Z}_n^* . Accordingly, we will adopt this group-theoretic setting.

3.2 A proof based on lattices

We prove the following:

Theorem 3. *Let n be an RSA modulus. Let I be an interval of length n^δ , with $2/3 < \delta < 1$. Let G be any subgroup of \mathbb{Z}_n^* and let d be the number of elements of the quotient group \mathbb{Z}_n^*/G . Then, for some constant M , the number of elements of $I \cap G$ is $K(1 + \lambda(I))$, where*

$$K = \frac{n^\delta}{d} \frac{\varphi(n)}{n}$$

and $|\lambda(I)|$ is bounded by $Mn^{1/3-\delta/2}$. Furthermore, M has the explicit bound $M \leq 5d$.

Remark. Observe that the case where $G = \mathbb{Z}_n^*$ is an easy consequence of lemma 5.

Proof. We number the elements of \mathbb{Z}_n^*/G as g_1, \dots, g_d (with g_1 the unit of G), and we let a_i be the number of elements of $\mathbb{Z}_n^* \cap I$ which equal g_i modulo G . We first show an upper bound for

$$A = \sum_{i=1}^d a_i^2$$

For any pair (x, y) in $I \times I$, we define $\sigma(x, y)$ as $xy^{-1} \bmod n$, when x, y both belong to \mathbb{Z}_n^* and set $\sigma(x, y) = \infty$ otherwise. Observe that A can be interpreted as the number of elements (x, y) of $\mathbb{Z}_n^* \cap I$ such that $\sigma(x, y) \in G$. Indeed, $xy^{-1} \bmod n$ is in G if and only if x and y are equal modulo G . We now use a counting argument to estimate the size of $\sigma^{-1}(\alpha)$, when α ranges over G . We distinguish two cases

1. When $L(\alpha)$ is an ε -good lattice, then, by lemma 3, $\sigma^{-1}(\alpha)$ has at most $(n^{\delta-1/2} + 2n^\varepsilon)^2$ elements.
2. Otherwise, we use lemma 4 to get that $\sigma^{-1}(\alpha)$ has at most $n^\delta + 1$ elements, which we replace by the (crude) bound $2n^\delta$.

Since there are at most $4n^{1-2\varepsilon}$ values of α which give rise to a lattice $L(\alpha)$ which is not ε -good, we get

$$A \leq \frac{\varphi(n)}{d} (n^{\delta-1/2} + 2n^\varepsilon)^2 + 8n^{1-2\varepsilon+\delta}.$$

Upperbounding $\varphi(n)$ by n , we get:

$$A \leq \frac{n^{2\delta}}{d} (1 + 2n^{1/2+\varepsilon-\delta})^2 + 8n^{1-2\varepsilon+\delta}.$$

We now set $\varepsilon = 1/6$. This yields the bound

$$A \leq \frac{n^{2\delta}}{d} (1 + 2n^{2/3-\delta})^2 + 8n^{2\delta} n^{2/3-\delta}.$$

Since δ is $> 2/3$, $n^{2/3-\delta}$ is < 1 and its square is bounded by $n^{2/3-\delta}$. We finally get:

$$A \leq \frac{n^{2\delta}}{d}(1 + (8 + 8d)n^{2/3-\delta}).$$

We now use the fact that the sum $B = \sum_{i=1}^d a_i$ is essentially known. Referring to the proof of lemma 5 above, we see that it differs from

$$n^\delta \frac{\varphi(n)}{n}$$

by at most 4. Now, the vector (a_1, \dots, a_d) lies on the d -dimensional hyperplane H defined by $B = \sum_{i=1}^d x_i$. Let (b_1, \dots, b_d) be the orthogonal projection of the origin on H . It is easily seen that $b_i = B/d$. The square of the euclidean distance between (a_1, \dots, a_d) and (b_1, \dots, b_d) is $\sum_{i=1}^d a_i^2 + \sum_{i=1}^d b_i^2 - 2 \sum_{i=1}^d a_i b_i$. This is $A - \frac{B^2}{d}$. we are thus led to find a lower bound for $\frac{B^2}{d}$. Using the same estimate as for the proof of lemma 5, we write

$$\frac{B^2}{d} \geq \frac{n^{2\delta}}{d} \left(\frac{\varphi(n)}{n} - 4n^{-\delta} \right)^2.$$

Using the fact that we have an RSA modulus, we use the lower bound $1 - \frac{3}{\sqrt{n}}$ for $\frac{\varphi(n)}{n}$ and, combining with the above, obtain the final bound

$$\frac{B^2}{d} \geq \frac{n^{2\delta}}{d} \left(1 - \frac{14}{\sqrt{n}} \right).$$

Finally, piecing bounds together, we get:

$$A - \frac{B^2}{d} \leq \frac{n^{2\delta}}{d} (22 + 8d)n^{2/3-\delta},$$

which provides a bound for $(a_1 - b_1)^2 = (a_1 - B/d)^2$. Observing that we only have to deal with $d \geq 2$, we easily get that $|a_1 - B/d|$ is at most

$$\sqrt{19} n^\delta n^{1/3-\delta/2}.$$

Replacing B/d by the constant

$$K = \frac{n^\delta \varphi(n)}{d n},$$

yields a minute difference $\leq 4/d$, which we handle by slightly raising the $\sqrt{19}$ constant. Thus, a_1 can be written $K(1 + \lambda(I))$, with

$$|\lambda(I)| \leq (\sqrt{19} + \gamma) d \frac{n}{\varphi(n)} n^{1/3-\delta/2},$$

We finally handle the term $\frac{n}{\varphi(n)}$ by raising the constant again. This gives the requested bound

$$|\lambda(I)| \leq 5dn^{1/3-\delta/2}.$$

3.3 A proof based on characters

We now show that a better bound for $\lambda(I)$, can be obtained as a consequence of the Polya-Vinogradov inequality of theorem 2.

Theorem 4. *Let n be an RSA modulus. Let I be an interval of length n^δ , with $1/2 < \delta < 1$. Let G be any subgroup of \mathbb{Z}_n^* and let d be the number of elements of the quotient group \mathbb{Z}_n^*/G . Then, for some constant M , the number of elements of $I \cap G$ is $K(1 + \lambda(I))$, where*

$$K = \frac{n^\delta}{d} \frac{\varphi(n)}{n},$$

and $|\lambda(I)|$ is bounded by $Mn^{1/2-\delta} \ln n$. Furthermore, M has the explicit bound $M \leq 5d$.

Proof. We consider the dual \hat{H} of the quotient group $H = \mathbb{Z}_n^*/G$. For any character χ over H , we can extend χ to G , by composing with the canonical homomorphism from G onto H . We still denote by χ , the resulting character. Since there are d characters altogether, we get, using the relations in theorem 1, that the number of elements of $I \cap G$ is equal to the sum

$$\frac{1}{d} \sum_{x \in I} \sum_{\chi \in \hat{G}} \chi(x),$$

Changing the order of the sums, we see that this number consists of two terms:

1. one comes from the principal character and equals: $\frac{|I \cap \mathbb{Z}_n^*|}{d}$,
2. the others come from the non trivial characters, and, by the Polya-Vinogradov inequality, each is bounded by $\frac{4}{d} n^{1/2} \ln n$.

By lemma 5, the first contribution differs from

$$K = \frac{n^\delta}{d} \frac{\varphi(n)}{n}$$

by at most $\frac{4}{d}$. Summing up with the second contribution, we obtain the bound:

$$\frac{4}{d} + \frac{4(d-1)}{d} n^{1/2} \ln n \leq 4n^{1/2} \ln n.$$

Altogether, we obtain that the number of elements of $I \cap G$ is $K(1 + \lambda(I))$, with

$$\lambda(I) \leq 4d \frac{n}{\varphi(n)} n^{1/2-\delta} \ln n,$$

Using the fact that n is an RSA modulus, we estimate $\varphi(n)$, by $n(1 - 4/\sqrt{n})$, and bound the multiplicative constant by a term

$$\simeq 4d \left(1 + \frac{4}{\sqrt{n}}\right).$$

This is bounded by 5. The result follows.

It should be noted that an even better bound has been obtained by Burgess [4]. The bound covers the case $1/4 < \delta < 1$, and reads:

$$|\lambda(I)| \leq Mdn^{\frac{1}{4r} - \frac{\delta}{r+1}} \ln n,$$

for any positive r . However, the constant M is not explicit, and therefore the improvement is not well suited for our purposes.

4 The security proof of ESIGN

In this section, we review the proof of security for ESIGN in view of the previous results. For the reader's convenience, we first provide a short description of the scheme and of the underlying mathematical problem AERP. We follow [15].

4.1 Description

The key generation algorithm of ESIGN chooses two large primes p, q of equal size k and computes the modulus $n = p^2q$. The sizes of p, q are set in such a way that the binary length $|n|$ of n equals $3k$. Additionally, an exponent $e > 4$ is chosen, possibly a small power of 2.

Signature generation uses a hash function \mathcal{H} , outputting strings of length $k - 1$, and is performed as follows:

1. Pick at random r in \mathbb{Z}_{pq}^* .
2. Convert $(0\|\mathcal{H}(m)\|0^{2k})$ into an integer y and compute $z = (y - r^e) \bmod n$.
3. Compute

$$w_0 = \left\lceil \frac{z}{pq} \right\rceil$$

$$w_1 = w_0.pq - z$$

4. If $w_1 \geq 2^{2k-1}$, return to step 1.
5. Set $u = w_0.(er^{e-1})^{-1} \bmod p$ and $s = r + upq$.
6. Output s as the signature of m .

Signature verification converts integer $s^e \bmod n$ into a bit string S of length $3k$ and checks that $[S]^k = 0\|\mathcal{H}(m)$, where $[S]^k$ denotes the k leading bits of S .

The key idea in ESIGN is that the arithmetical progression $r^e \bmod n + tpq$ consists of e -th powers of integers easily computed from r . The signature generation algorithm simply adjusts t so as to fall into a prescribed interval, with lower end-point y . The test at step 4 actually sets the length of this prescribed interval to 2^{2k-1} .

The following lemma will prove useful in the sequel.

Lemma 6. *For a fixed message m , the e -th power $s^e \bmod n$ of the output s of the signature generation algorithm is uniformly distributed over the set of e -th powers of elements of \mathbb{Z}_n^* lying in the interval $[y, y + 2^{2k-1})$.*

Proof. Denote by $S(y)$ the intersection of the set of e -th powers in \mathbb{Z}_n^* and the interval $[y, y + 2^{2k-1})$. Observe that $s = r + tpq$ uniquely defines $r = s \bmod pq$ from s . This shows that any element in $S(y)$ comes from a single r . To see that all elements in $S(y)$ are uniformly hit, pick $w \in S(y)$, consider any r in \mathbb{Z}_{pq}^* such that $r^e = w \bmod pq$, and apply the signature generation algorithm with r , disregarding the check at step 4. This produces a value of s such that $s^e = r^e = w \bmod pq$. Thus, w and $s^e \bmod n$ lie in the arithmetical progression $s^e + tpq$. Since this arithmetical progression has a single element in the interval $[y, y + 2^{2k-1})$, we get that $s^e \bmod n = w$. The check at step 4 turns out correct and the signature generation algorithm duly hits w as many times as the number of e -th roots of an e -th power.

4.2 The approximate e -th root problem

As noted in the previous section, RSA moduli of the form p^2q offer a very efficient way to solve the following problem, having knowledge of the factorization of n : given n and y in \mathbb{Z}_n^* , find x such that $x^e \bmod n$ lies in the interval $[y, y + 2^{2k-1})$, where the bit-size of n is $3k$ and $[y, y + 2^{2k-1})$ denotes $\{u \mid y \leq u < y + 2^{2k-1}\}$.

It is conjectured that the above problem, called the approximate e -th root problem (AERP) in [15], is hard to solve. More precisely, denote by $\text{Succ}^{\text{aerp}}(\tau, k)$ the probability for any adversary \mathcal{A} to find an element whose e -th power lies in the prescribed interval, within time τ , in symbols:

$$\Pr[(n, e) \leftarrow \mathcal{K}(1^k), y \leftarrow \mathbb{Z}_n^*, x \leftarrow \mathcal{A}(n, e, y) : (x^e \bmod n) \in [y, y + 2^{2k-1})],$$

then, for large enough moduli, this probability is extremely small. Variants of the above can be considered, where the length of the interval is replaced by 2^{2k} or 2^{2k+1} .

4.3 Security proof

We now complete the security proof of ESIGN, in order to cover the case where e is not prime to $\varphi(n)$. We use the random oracle model and prove the following security result, where $T_{\text{exp}}(k)$ denotes the computing time of modular exponentiation modulo a $3k$ -bit integer.

Theorem 5. *Let \mathcal{A} be a SO-CMA-adversary against the ESIGN signature scheme that produces an existential forgery, with success probability ε , within time τ , making q_H queries to the hash function and q_s distinct requests to the signing oracle respectively. Then, AERP can be solved with probability ε' , and within time τ' , where*

$$\begin{aligned} \varepsilon' &\geq \frac{\varepsilon - 2^{-k+1}}{q_H} - (q_H + q_s) \times (3/4)^k - \frac{ke^2(q_H + q_s)}{2^{k-6}} \\ \tau' &\leq \tau + k(q_s + q_H) \cdot T_{\text{exp}}(k). \end{aligned}$$

Our method of proof is inspired by Shoup [19]. It differs from [15] but extends the proof given in [20]. The security estimates are similar and show the same multiplicative loss q_H : contrary to schemes based on self-reducible problems, it does not seem that this can be avoided. Recall that earlier proofs used the assumption that e is prime to $\varphi(n)$, which we avoid. This brings additional terms in the security estimates, which account for the simulation of the random oracle. Also note that our security model is the single occurrence chosen message attack SO-CMA from [20], where the attacker is only allowed to query each message once. As already noted, it is easy to modify the scheme to withstand CMA attackers and our proof can be modified accordingly.

As usual, the proof considers a sequence of **Game**₁, **Game**₂, etc of modified attack games starting from the actual game **Game**₀. Each of the games operates on the same underlying probability space, only the rules defining how the view is computed differ from game to game.

Proof. (of Theorem 5). We consider an adversary \mathcal{A} outputting an existential forgery (m, s) , with probability ε , within time τ . We denote by q_H and q_s respectively the number of queries from the random oracle \mathcal{H} and from the signing oracle. As explained, we start by playing the game coming from the actual adversary, and modify it step by step, until we reach a final game, whose success probability has an upper-bound obviously related to solving AERP on a random instance (n, e, v) .

Game₀: The key generation algorithm $\mathcal{K}(1^k)$ is run and produces a pair of keys $(\mathbf{pk}, \mathbf{sk})$. The adversary \mathcal{A} is fed with \mathbf{pk} and, querying the random oracle \mathcal{H} and the signing oracle $\Sigma_{\mathbf{sk}}$, it outputs a pair (m, s) . We denote by S_0 the event that $V_{\mathbf{pk}}(m, s) = 1$. We use a similar notation S_i in any **Game** _{i} below. By definition, we have

$$\Pr[S_0] = \varepsilon.$$

Game₁: In this game, we discard executions, which end up outputting a valid message/signature pair (m, s) , such that m has not been queried from \mathcal{H} . This means restricting to the event **AskH** that m has been queried from \mathcal{H} . Unwinding the ESIGN format, we write: $s^e = 0 \parallel w \parallel \star \bmod n$. If **AskH** does not hold, $\mathcal{H}(m)$ is undefined, and the probability that $\mathcal{H}(m) = w$ holds is $1/2^{k-1}$: $\Pr[S_0 \mid \neg \mathbf{AskH}] \leq 2^{-k+1}$. Thus,

$$\Pr[S_1] = \Pr[S_0 \wedge \mathbf{AskH}] \geq \Pr[S_0] - 2^{-k+1}.$$

Game₂: In this game, we choose at random an index κ between 1 and q_H . We let m_κ be the κ -th message queried to \mathcal{H} by the adversary. We then discard executions which output a valid message/signature pair (m, s) , such that $m \neq m_\kappa$. Since the additional random value κ is chosen independently of the execution of **Game**₁,

$$\Pr[S_2] = \Pr[S_1]/q_H.$$

Game₃: In this game, we immediately abort if a signing query involves message m_κ . By the definition of existential forgery, this only eliminates executions outside S_2 . Thus:

$$\Pr[S_3] = \Pr[S_2].$$

Game₄: We now simulate the random oracle \mathcal{H} , by maintaining an appropriate list, which we denote by H-List. For any fresh query m , other than the κ -th query, we pick at random $u \in \mathbb{Z}_n$ and compute $z = u^e \bmod n$, until the most significant bit of z is 0. We next parse z as $0 \| w \| \star$, where w is of length $k - 1$ and check whether $z - w \cdot 2^{2k}$ is less than 2^{2k-1} . If this is true, we store (m, u, w) in H-List and returns w as the answer to the oracle call. Otherwise we restart the simulation of the current query. From theorem 4, we see that the game differs from the previous due to a slightly biased simulated distribution. This distribution is obtained by setting $z = w2^{2k}$, counting the number of e -th powers of elements of \mathbb{Z}_n^* lying in the interval $[z, z + 2^{2k-1})$, and multiplying by a suitable constant for normalisation. Recall that, an element x of $[z, z + 2^{2k-1})$ is an e -th power modulo n if and only if $x \bmod pq$ is an e -th power modulo pq . This is basically a restatement of the key idea of ESIGN. Thus, setting $z' = z \bmod pq$, we have to count the number $\nu(z)$ of elements of the interval $[z', z' + 2^{2k-1})$, which belong to the subgroup G of e -th powers in \mathbb{Z}_{pq}^* . By theorem 4, the result is $K(1 + \lambda(z))$, where $|\lambda(z)|$ is bounded by $M(pq)^{1/2}2^{-2k+1} \ln pq$, and where K, M are appropriate constants. This yields

$$|\lambda(z)| \leq M2^{-k+1/2} \ln pq$$

Upperbounding $\ln pq$ by $3/2 \log pq$, we get:

$$|\lambda(z)| \leq 3Mk2^{-k+1/2}$$

Now, it is easily seen that any probability distribution obtained by normalizing a function $\nu(z) = K(1 + \lambda(z))$, where $\lambda(z)$ is bounded by λ , differs from the uniform distribution by at most $\frac{2\lambda}{1-\lambda} \simeq 2\lambda$. Taking into account the bound $M \leq 5d$, where d is the number of elements of the quotient of \mathbb{Z}_{pq}^* by the subgroup of e -th powers, and bounding d by e^2 , we conclude that the statistical distance of the simulated distribution to the uniform distribution is bounded by twice the bound on λ , which is $30e^2k2^{-k+1/2} \leq 64e^2k2^{-k}$. Summing up for all oracle calls, we get:

$$|\Pr[S_4] - \Pr[S_3]| \leq \frac{ke^2(q_H + q_s)}{2^{k-6}}.$$

Game₅: Here, we modify the previous simulation stopping and aborting the game when the \mathcal{H} query cannot be simulated after k trials. This game differs from the previous one when w remains undefined after k attempts.

$$\Pr[S_5] \geq \Pr[S_4] - (q_H + q_s) \times (3/4)^k.$$

Game₆: We complete the simulation by replacing $\mathcal{H}(m_\kappa)$ by v , where v is an additional random string, which serves as an input to the AERP problem. The distribution of \mathcal{H} -outputs is unchanged:

$$\Pr[S_6] = \Pr[S_5].$$

Game₇: We finally simulate the signing oracle: for any m , whose signature is queried, we know that $m = m_\kappa$ does not hold, since corresponding executions have been aborted in **Game₃**. Thus **H-List** includes a triple (m, u, w) , such that $u^e \bmod n$ has its k leading bits of the form $0 \parallel \mathcal{H}(m)$. Accordingly, u provides a valid signature of m . Furthermore, referring to lemma 6, we see that the signing oracle outputs a value s , such that $s^e \bmod n$ is uniformly distributed over all elements of \mathbb{Z}_n^* whose $k+1$ leading bits match up with $0 \parallel \mathcal{H}(M) \parallel 0$. Keeping in mind that $\mathcal{H}(m)$ is chosen at random, we conclude that s and u follow an identical distribution. We now argue that the simulation is perfect. The key fact is that, due to the **SO-CMA** setting, all inputs m submitted to the \mathcal{H} oracle by the signing oracle during execution are distinct. This implies that the values of s returned at each invocation of the signing oracle are independent. Since the values of u are also independent, the overall distribution of simulated signatures obtained at **Game₇** is identical to the distribution of actual signatures from **Game₆**. Therefore,

$$\Pr[S_7] = \Pr[S_6].$$

Summing up the above inequalities, we obtain

$$\begin{aligned} \Pr[S_7] &\geq \Pr[S_4] - (q_H + q_s) \times \left(\frac{3}{4}\right)^k \geq \Pr[S_3] - (q_H + q_s) \times \left(\frac{3}{4}\right)^k - \frac{ke^2(q_H + q_s)}{2^{k-6}} \\ &\geq \frac{\Pr[S_1]}{q_H} - (q_H + q_s) \times \left(\frac{3}{4}\right)^k - \frac{ke^2(q_H + q_s)}{2^{k-6}} \\ &\geq \frac{\varepsilon - 2^{-k+1}}{q_H} - (q_H + q_s) \times \left(\frac{3}{4}\right)^k - \frac{ke^2(q_H + q_s)}{2^{k-6}} \end{aligned}$$

When **Game₇** terminates outputting a valid message/signature pair (m, s) , we unwind the **ESIGN** format and get $s^e = (0 \parallel v \parallel \star) \bmod n$, with $v = \mathcal{H}(m)$. If S_7 holds, we know that $m = m_\kappa$ and $\mathcal{H}(m) = v$. This leads to an element whose e -th power lies in the interval $[v2^{2k}, v2^{2k} + 2^{2k})$, thus solving an instance of **AERP**. We finally have: $\Pr[S_7] \leq \text{Succ}^{\text{aerp}}(\tau', k)$, where τ' denotes the running time of **Game₇**. This is the requested bound. Observe that τ' is the sum of the time for the original attack, plus the time required for simulations, which amounts to at most $k(q_s + q_H)$ modular exponentiations.

Remark. The security proof that appears in [15] replaces the k multiplicative factor in the running time by 4. This is intuitively related to the fact that, on average, it takes at most 4 steps to perform the simulation of each call to \mathcal{H} in **Game₄**. It is actually possible to improve our time estimate

$$\tau' \leq \tau + k(q_s + q_H) \cdot T_{exp}(k),$$

to

$$\tau' \leq \tau + 4(q_s + q_H) \cdot T_{exp}(k),$$

This uses a method due to Jonsson [12]. It modifies the strategy for the simulation of \mathcal{H} in Game_5 : instead of limiting the number of trials allowed, at each execution, to find a value of z in the correct range, it sets a counter that bounds the overall number of retries, during the entire algorithm.

References

1. M. Bellare and P. Rogaway. Random Oracles Are Practical: a Paradigm for Designing Efficient Protocols. In *Proc. of the 1st CCS*, pages 62–73, ACM Press, New York, 1993.
2. M. Bellare and P. Rogaway. The Exact Security of Digital Signatures – How to Sign with RSA and Rabin. In *Eurocrypt '96*, LNCS 1070, pages 399–416, Springer-Verlag, Berlin, 1996.
3. E. Brickell and J. M. DeLaurentis. An Attack on a Signature Scheme proposed by Okamoto and Shiraishi. In *Crypto '85*, LNCS 218, pages 28–32, Springer-Verlag, Berlin, 1986.
4. D.A. Burgess. On character sums and primitive roots, *Proc. London Math. Soc.*, 12 (1962), 179–192.
5. H. Davenport. *Multiplicative Number theory*, Graduate Texts in Mathematics, Vol 74, Springer Verlag, (1980).
6. W.J. Ellison and M. Mendes France. *Les nombres premiers*, Hermann, Paris (1975).
7. M. Girault, P. Toffin and B. Vallée. Computation of Approximate L-th Roots Modulo n and Application to Cryptography. In *Crypto '88*, LNCS 403, pages 100–118, Springer-Verlag, Berlin, 1989.
8. S. Goldwasser, S. Micali, and R. Rivest. A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. *SIAM Journal of Computing*, 17(2):281–308, April 1988.
9. L. Granboulan. How to repair ESIGN, NESSIE internal document, may 2002. See <http://www.cryptonessie.org>, Docuemnyt NES/DOC/ENS/WP5/019.
10. IEEE Standard 1363–2000. Standard Specifications for Public Key Cryptography. IEEE. Available from <http://grouper.ieee.org/groups/1363>, August 2000.
11. IEEE P1363a Draft Version 9. Standard Specifications for Public Key Cryptography: Additional Techniques.
12. J. Jonsson. Security Proofs for RSA–PSS and Its Variants. Cryptology ePrint Archive 2001/053. June 2001. Available from <http://eprint.iacr.org/>.
13. A. K. Lenstra, H. W. Lenstra and L. Lovász, Factoring polynomials with rational coefficients, *Mathematische Ann.*, 261, (1982), 513–534.
14. T. Okamoto. A Fast Signature Scheme Based on Congruential Polynomial Operations. *IEEE Transactions on Information Theory*, IT–36 (1), pages 47–53, 1990.
15. T. Okamoto, E. Fujisaki and H. Morita. TSH-ESIGN: Efficient Digital Signature Scheme Using Trisection Size Hash, Submission to P1363a, 1998.
16. T. Okamoto and A. Shiraishi. A Fast Signature Scheme Based on Quadratic Inequalities. Proc. of the ACM Symp. Security and Privacy, ACM Press, pages 123–132, 1985.
17. G. Pólya. Über die Verteilung des quadratischen Reste und Nichtreste, *Göttinger Nachrichten* (1918), 21–26.

18. R. Rivest, A. Shamir, and L. Adleman. A Method for Obtaining Digital Signatures and Public Key Cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.
19. V. Shoup. OAEP Reconsidered. In *Crypto '2001*, LNCS 2139, pages 239–259. Springer-Verlag, Berlin, 2001. Also appeared in the Cryptology ePrint Archive 2000/060. November 2000. Available from <http://eprint.iacr.org/>.
20. J. Stern, D. Pointcheval, J. Malone-Lee, and N. Smart. Flaws in Applying Proof Methodologies to Signature Schemes. In *Crypto '02*, LNCS 2442, pages 93–110.
21. B. Vallée, M. Girault, and P. Toffin. How to break Okamoto's Cryptosystem by Reducing Lattice Bases. In *Eurocrypt '88*, LNCS 330, pages 281–292, Springer-Verlag, Berlin, 1988.
22. B. Vallée, M. Girault and P. Toffin. How to Guess ℓ th Roots Modulo n by Reducing Lattice Bases. In *AAECC-6*, LNCS 357, pages 427–442, Springer-Verlag, Berlin, 1988.
23. I.M. Vinogradov. Sur la distributions des résidus et des non-résidus des puissances, *J. Phys.-Math. Soc. Perm.* 1 (1918), 94-96.