

A Note on Low Order Assumptions in RSA groups

István András Seres and Péter Burcsi

Eötvös Loránd University

April 2020

Abstract

In this short note, we show that substantially weaker Low Order assumptions are sufficient to prove the soundness of Pietrzak’s protocol for proof of exponentiation in groups of unknown order. This constitutes the first step to a better understanding of the asymptotic computational complexity of breaking the soundness of the protocol. Furthermore, we prove the equivalence of the (weaker) Low Order assumption(s) and the Factoring assumption in RSA groups for a non-negligible portion of moduli. We argue that in practice our reduction applies for a considerable amount of deployed moduli. Our results have cryptographic applications, most importantly in the theory of recently proposed verifiable delay function constructions. Finally, we describe how to certify RSA moduli free of low order elements.

1 Introduction

Verifiable delay functions (VDF) are powerful cryptographic tools [BBBF18] that opened up a plethora of applications, such as non-interactive timestamping [LSS19], proof of replication [FBGB19] or randomness beacons [BGB17]. A VDF is a function whose evaluation takes $\mathcal{O}(T)$ sequential steps and cannot be sped up by parallelism. Additionally, a prover, or evaluator, can produce publicly verifiable and succinct proofs that the function evaluation was correct. A crucial requirement for a VDF that there needs to be an exponential gap between function evaluation and proof verification time, more precisely verification time should be in $\mathcal{O}(\log T)$. Naturally, we require correctness and soundness from the applied proof systems. Specifically, an honest prover should always be able to convince the verifier, while a malicious prover should only be able to produce correct proofs with negligible probability.

Recent VDF constructions [Pie18, Wes19] proposed by Pietrzak and Wesolowski instantiate VDFs in groups of unknown order, i.e. groups for which the order cannot be computed efficiently [RSA78]. The existence of verifiable delay functions in the random oracle model is ruled out [MSW], moreover, groups of unknown order are shown to be mandatory for generic group delay functions [RSS]. Both constructions [Pie18, Wes19] rely on novel, non-standard cryptographic assumptions. The soundness of these constructions can be proved by assuming the Low Order (LO) or Adaptive Root (AR) assumptions in groups of unknown order. Therefore there is an emerging need to understand better these new, non-standard cryptographic assumptions. In this note, we turn our attention to the LO assumption as it is a potentially weaker assumption than the AR assumption [BBF18].

Our contributions. In this note, we provide the following contributions.

- We observe that for the soundness of Pietrzak’s proof of exponentiation succinct argument, one can assume substantially weaker LO assumptions than as previously defined in [BBF18]. In other words, we show that potentially it is harder to break soundness of Pietrzak’s argument than as it was argued in [BBF18].
- We prove the equivalence of the LO and Factoring assumptions in RSA groups for a non-negligible portion of moduli. We argue that this result has practical consequences and that in practice one can deem the LO assumption to be equivalent to Factoring for the majority of used RSA moduli.
- We show how one could certify RSA moduli being free of low order elements using a non-interactive honest-verifier zero-knowledge proof system by Goldberg et al [GRSB19].

The rest of this note is organized as follows. In Section 2 we provide background on the recently introduced LO and AR assumptions. We show the sufficiency of weaker LO assumptions in Section 3. In Section 4 we provide our reduction from Factoring to LO assumption for non-negligible RSA moduli. We describe a method to certify RSA moduli free of low order elements in Section 5. Finally, we point out open problems in Section 6.

2 Preliminaries

2.1 Notations

Let \mathbb{G} be a group of unknown order. As usual λ denotes the security parameter. Integers denoted as p, q are always primes, while N is a semiprime, ie. $N = pq$, sometimes referred to as RSA-modulus. All logarithms have two as their base, unless stated otherwise. The $\phi(\cdot)$ denotes Euler's totient function. Let $ord_m(a)$ denote the order of element a in \mathbb{Z}_m^\times . In the following we assume, that the size of the moduli is bounded by polynomial of the security parameter $s(\lambda)$, such that $p, q, \phi(N), N \approx \mathcal{O}(2^{s(\lambda)})$. We mean by $x \stackrel{\$}{\leftarrow} S$, that x is uniformly at random sampled from set S .

2.2 Proof of exponentiation and Pietrzak's succinct argument

In a proof of exponentiation protocol in \mathbb{G} the prover wants to convince the verifier that $h = g^{(2^T)}$ holds in \mathbb{G} . That is, the protocol is an argument system for the relation

$$\mathcal{L}_{\text{EXP}} = \{(\mathbb{G}, g, h, T) : h = g^{(2^T)} \in \mathbb{G}\} \quad (1)$$

Pietrzak's proof system works as follows [Pie18].

1. The verifier checks that $g, h \in \mathbb{G}$ and outputs *reject* if not,
2. If $T = 1$ the verifier checks that $h = g^2$ in \mathbb{G} , outputs *accept* or *reject*, and stops.
3. If $T > 1$ the prover and verifier do:
 - (a) The prover computes $v \leftarrow g^{(2^{T/2})} \in \mathbb{G}$ and sends v to the verifier. The verifier checks that $v \in \mathbb{G}$ and outputs *reject* and stops, if not.
Next, the prover needs to convince the verifier that $h = v^{(2^{T/2})}$ and $v = g^{(2^{T/2})}$, which proves that $h = g^{(2^T)}$. Since the same exponent is used in both equalities, they can be verified simultaneously by checking a random linear combination, namely that

$$v^r h = (g^r v)^{(2^{T/2})}, \quad \text{where } r \stackrel{\$}{\leftarrow} \{1, \dots, 2^\lambda\}.$$

The verifier and prover do so as follows.

- (b) The verifier sends the prover a random $r \stackrel{\$}{\leftarrow} \{1, \dots, 2^\lambda\}$.
- (c) Both the prover and verifier compute $g_1 \leftarrow g^r v$ and $h_1 \leftarrow v^r h \in \mathbb{G}$.
- (d) The prover and verifier recursively engage in an interactive proof that $(\mathbb{G}, g_1, h_1, T/2) \in \mathcal{L}_{\text{EXP}}$, namely that $h_1 = g_1^{(2^{T/2})} \in \mathbb{G}$.

Figure 1: Pietrzak's succinct argument for the proof of exponentiation language \mathcal{L}_{EXP} , verbatim from [BBF18]

For simplicity, we assume that T is a power of two in which case the protocol takes $\log T$ rounds. The protocol can be adjusted to handle arbitrary T , including a T that is not a power of two [Pie18].

Naturally, the protocol can be made non-interactive by using the Fiat-Shamir heuristic. The prover generates the challenge r at every level of the recursion by hashing the quantities (\mathbb{G}, g, h, T, v) at that level, and appends v to the overall proof π . Hence, the overall proof π contains $\log T$ elements in \mathbb{G} . Proof generation has a complexity of $\frac{2T}{s\sqrt{T}}$ with s being the amount of processors. At every level of the recursion, the verifier does two small exponentiations in \mathbb{G} to compute g_i and h_i for the i th level of the recursion. Therefore, verifying the proof takes $\mathcal{O}(\log T)$ small exponentiations in \mathbb{G} .

An implementation study of Pietrzak's and Wesolowski's VDF construction showed that it is faster to verify Pietrzak's VDF than that of Wesolowski. However, Pietrzak's VDF comes with larger proofs, hence demanding slightly larger bandwidth [AVD].

2.3 RSA assumptions

Informally, the LO assumption states that it is computationally infeasible to find a low order element in a random RSA group. Let us recall the formal definition of the LO assumption [BBF18].

Definition 1. *The Low Order assumption holds for $GGen$ if there is no efficient adversary \mathcal{A} finding any element of low order:*

$$\Pr \left[\begin{array}{l} u^l = 1, u \notin \{1, -1\} : \\ \mathbb{G} \xleftarrow{\$} GGen(\lambda) \\ (u, l) \leftarrow \mathcal{A}(\mathbb{G}) \\ \text{and } l < 2^{\text{poly}(\lambda)} \end{array} \right] \leq \text{negl}(\lambda) \quad (2)$$

We remark that the LO assumption unconditionally holds in QR_N , the group of quadratic residues mod N , since it contains no elements of low order. In case of RSA groups, we model $GGen(\lambda)$ as uniformly randomly sampling primes from an appropriate domain with respect to λ .

Definition 2. *The Factoring assumption states that for random primes p, q it is difficult to factor $N = pq$.*

It is trivial to see, that if there was an adversary \mathcal{A} breaking the *Factoring* assumption, then one could easily calculate any roots mod N by applying the Chinese Remainder Theorem. Computing arbitrary roots enables an adversary finding low order elements. Note that it is also true that technically, having a factoring algorithm that works for a non-negligible portion of RSA moduli, does not imply that the LO assumption is broken. It could be the case that factoring works if and only if the modulus is the product of two safe primes. For these moduli, we do not have low order elements at all. Certainly, what one implicitly understands by saying that the LO assumption is stronger than the factoring assumption is that a reduction exists for $GGen$ that only outputs moduli with low order elements.

Even though we did not introduce Wesolowski's succinct argument for \mathcal{L}_{EXP} , for sake of self-containedness we introduce the assumption needed to prove its soundness.

Definition 3. *The Adaptive Root Assumption holds for $GGen$ if there is no efficient adversary $(\mathcal{A}_0, \mathcal{A}_1)$ that succeeds in the following task. First, \mathcal{A}_0 outputs an element $w \in \mathbb{G}$ and some state st . Then, a random prime in $\text{Primes}(\lambda)$ is chosen and $\mathcal{A}_1(w, l, st)$ outputs $w^{1/l} \in \mathbb{G}$. For all efficient $(\mathcal{A}_0, \mathcal{A}_1)$:*

$$\Pr \left[\begin{array}{l} \mathbb{G} \xleftarrow{\$} GGen(\lambda) \\ (w, st) \leftarrow \mathcal{A}_0(\mathbb{G}) \\ u^l = w \neq 1 : \quad l \xleftarrow{\$} \Pi_\lambda = \text{Primes}(\lambda) \\ u \leftarrow \mathcal{A}_1(w, l, st) \end{array} \right] \leq \text{negl}(\lambda) \quad (3)$$

We note that the number of primes in Π_λ should be exponential in λ : it is possible to precompute w using $2^{|\Pi_\lambda|}$ exponentiations. Then, an adversary with 2^M memory can store intermediate exponents and compute adaptive roots using $2^{|\Pi_\lambda| - M}$ exponentiations for each. Moreover, it was shown that the Adaptive Root Assumption holds in the generic group model [BBF19].

2.4 Number theoretic tools

We recall the following lemma without proof.

Lemma 1. *The map $x \rightarrow x^e \pmod N$ is a permutation of Z_N^* if and only if $\gcd(e, \phi(N)) = 1$.*

Furthermore, let us define the language of RSA public keys (N, e) , such that the map $x \rightarrow x^e \pmod N$ is a permutation over Z_N^* .

$$\mathcal{L}_{\text{perm}Z_N^*} = \{(N, e) \mid N, e > 0 \wedge \gcd(e, \phi(N)) = 1\} \quad (4)$$

For this particular language, Goldberg et al. devised a public-coin protocol [GRSB19] with perfect completeness, perfect honest-verifier zero-knowledge, and statistical soundness.

Later we will need to count the number of integers without factors from an interval. Therefore we introduce the following function and notation.

Definition 4. *Denote by $\Gamma(x, y, z)$ the number of all positive integers less than x which are free of prime divisors from the interval $(z, y]$.*

Theorem 1. (Weingartner [Wei01]) Let $u = \frac{\log x}{\log y}$ and $v = \frac{\log x}{\log z}$. Then uniformly for $\frac{3}{2} \leq z \leq y$, whenever $x \geq yz$, we have the following asymptotic relationship for $\Gamma(x, y, z)$:

$$\Gamma(x, y, z) = x\eta(u, v) \left(1 + \mathcal{O}\left(\frac{1}{\log z}\right)\right), \quad (5)$$

where $1 \leq u \leq v$ and $\eta(u, v) \geq \frac{u}{2v}$.

We remark that a similar result and asymptotic was obtained by Warlimont [War90], however, he solely proved his results for fixed z . We will crucially rely on the uniform convergence of the asymptotics in Equation 5.

3 Soundness of Pietrzak’s argument and weaker LO assumptions

Boneh et al. introduced the low order assumption as a sufficient and necessary assumption to prove the soundness of Pietrzak’s argument [BBF18]. In this section, we show that the original definition of the LO assumption, cf. Section 2.3, is not necessary for proving soundness. We will show that one needs to assume almost exponentially weaker assumptions as a necessary and sufficient assumption for soundness of Pietrzak’s proof of exponentiation protocol.

3.1 (Non)-necessity of the LO assumption and weaker LO assumptions

Let us assume that the LO assumption, cf. Definition 1, is broken. What is the probability that such a potent adversary could break the soundness of Pietrzak’s argument system? As it turns out, it can still be negligible.

The main observation in the soundness analysis is that whenever a malicious prover finds (u, l) , a low order element u , with order $l < 2^{\text{poly}(\lambda)}$ can potentially break the soundness of the argument system. This is because if $(\mathbb{G}, g, h, T) \in \mathcal{L}_{\text{EXP}}$, then $(\mathbb{G}, g, hu, T) \notin \mathcal{L}_{\text{EXP}}$ and will be incorrectly accepted by the verifier with probability $1/l$. Towards breaking soundness malicious prover sends $v \leftarrow g^{(2^{T/2})} u \in \mathbb{G}$. Soundness of the argument system does not hold whenever $r + 1 \equiv 2^{T/2} \pmod{l}$, since $(\mathbb{G}, g^r v, v^r(hu), T/2) \in \mathcal{L}_{\text{EXP}}$.

However, there might be an adversary \mathcal{A} breaking the LO assumption with non-negligible probability, even though their success probability in breaking the soundness of the argument system is negligible. This can occur, if \mathcal{A} is only able to find low order elements with their order in $2^{\Theta(\text{poly}(\lambda))}$. In this case the probability that \mathcal{A} breaks soundness is at most $1/2^{\Theta(\text{poly}(\lambda))}$, hence negligible.

After this discussion, one can see that the $2^{\text{poly}(\lambda)}$ upper bound for the order of the low order element needs to be decreased to get the weakest necessary assumption for proving soundness of Pietrzak’s argument. Therefore, in quest to define a sufficiently weak LO assumption with the lowest permissible bound on the order of the low order element, we introduce the following smallest subexponential LO assumption.

Definition 5. *The Subexponential Low Order assumption. For any probabilistic polynomial time adversary \mathcal{A} , and for any $0 < \epsilon$, finding any element of subexponentially low order is hard:*

$$\Pr \left[u^l = 1, u \notin \{1, -1\} : \begin{array}{l} \mathbb{G} \xleftarrow{\$} \text{GGen}(\lambda) \\ (u, l) \leftarrow \mathcal{A}(\mathbb{G}) \\ \text{and } l < 2^{\log^{1+\epsilon}(\lambda)} \end{array} \right] \leq \text{negl}(\lambda) \quad (6)$$

We note, that the $2^{\log^{1+\epsilon}(\lambda)}$ upper bound cannot be lowered to a polynomial-bound as then the assumption would not be sufficient, see next subsection. In a nutshell, for sufficiency to hold one needs to assume a superpolynomial upper bound for the order of the low order element. Furthermore, we remark that even the weaker LO assumption introduced in Definition 5 is not necessary, because there are non-negligible RSA-moduli with $\phi(N)$ having divisors between any $\text{poly}(\lambda)$ and $2^{\log^{1+\epsilon}(\lambda)}$ [BS13]. Again, there might be adversaries who are only able to find low order elements with their order being between $\text{poly}(\lambda)$ and $2^{\log^{1+\epsilon}(\lambda)}$, therefore their success probability in breaking soundness of Pietrzak’s argument would be negligible.

3.2 Sufficiency of weaker LO assumptions for soundness of Pietrzak’s argument

Let $f(\lambda)$ denote the function limiting the maximal order of the low order element in an LO assumption. We see that in Definition 1, $f(\lambda) = 2^{\text{poly}(\lambda)}$, while in Definition 5, $f(\lambda) = 2^{\log^{1+\epsilon}(\lambda)}$. We are interested in finding the minimal $f(\lambda)$ such that the LO assumption with that $f(\lambda)$ is still sufficient for proving the soundness of Pietrzak’s argument system for \mathcal{L}_{EXP} .

Theorem 2. *If the soundness of Pietrzak’s succinct argument for proof of exponentiation is broken, then so is the subexponential LO assumption, cf. Definition 5.*

Proof. Hereby we reuse the proof given by Boneh et al [BBF18] with modifications to our specific setting. Hence, we recall their proof for the sufficiency of the LO assumption for breaking the soundness of Pietrzak’s argument. Let \mathcal{A} be an adversary who breaks the soundness of Pietrzak’s argument with non-negligible probability ϵ . We use a forking argument to construct an adversary \mathcal{B} that breaks the low order assumption using \mathcal{A} .

Recall that 2^t is an upper bound on the value T output by \mathcal{A} . Let $\mathcal{A}(\mathbb{G}, r_0, \dots, r_{t-1}; R)$ denote an execution of \mathcal{A} with random tape R , where r_0, \dots, r_{t-1} are the verifier’s challenges at each level of the recursion. The adversary \mathcal{A} outputs (J, σ) . The protocol transcript is denoted by σ which is a sequence of $t+1$ tuples: $\sigma = (P_0, v_0), \dots, (P_t, v_t)$, where $P_i = (\mathbb{G}, g_i, h_i, T/2^i)$ is the input to the recursion at level i , and v_i is the prover’s message at level i . Adversary \mathcal{A} also outputs the smallest index J for which $P_J \notin \mathcal{L}_{\text{EXP}}$ but $P_{J+1} \in \mathcal{L}_{\text{EXP}}$ whenever $P_0 \in \mathcal{L}_{\text{EXP}}$ and $P_t \notin \mathcal{L}_{\text{EXP}}$. Otherwise $J = -1$ if such an index does not exist¹. Recall that $g_i \leftarrow g_{i-1}^{r_{i-1}} v_{i-1}$ and $h_i \leftarrow v_{i-1}^{r_{i-1}} h_{i-1}$ for $i = 1, \dots, t$. Here we assume $T = 2^t$, but if $T < 2^t$ then we replicate the last pair $(P_{\log T}, v_{\log T})$ to get a full transcript of $t+1$ tuples.

1. Let $K = [1, 2^{\log^{1+\epsilon}(\lambda)}]$.
2. Input generator IG samples $\mathbf{x} = (x_0, x_1, \dots, x_{t-1}) \xleftarrow{\$} [0, 2^{\lambda - 2^{\log^{1+\epsilon}(\lambda)}}]$.
3. We define algorithm \mathcal{A}' on random tape R and $k_i \in K$ for all i , as follows: $\mathcal{A}'(\mathbf{x}, k_0, k_1, \dots, k_{t-1}; R)$ invokes $\mathcal{A}(\mathbb{G}, r_0, \dots, r_{t-1}; R)$ with $r_i = x_i \cdot 2^{\log^{1+\epsilon}(\lambda)} + k_i$. Algorithm \mathcal{A}' returns the same output as \mathcal{A} . Note that whenever $k_i \xleftarrow{\$} K$ and $\mathbf{x} \xleftarrow{\$} \text{IG}$, then r_i is also sampled uniformly at random from $[1, 2^\lambda]$. Hence, the success probability of \mathcal{A}' equals that of \mathcal{A} .
4. Next, define the following probabilistic experiment $F_{\mathcal{A}'}(\mathbf{x})$: let $P_0 \notin \mathcal{L}_{\text{EXP}}$ but $P_t \in \mathcal{L}_{\text{EXP}}$ (i.e., the verifier incorrectly accepts P_0) then:
 - choose a random tape R for \mathcal{A}' .
 - Sample $k_i \xleftarrow{\$} K$ for $i \in [0, t-1]$.
 - We obtain $(I, \sigma) \leftarrow \mathcal{A}'(\mathbf{x}, k_0, k_1, \dots, k_{t-1}; R)$.
 - If $I = -1$ output *fail*.
 - If $I \geq 0$, then sample new $k'_i \xleftarrow{\$} K$ for $i = (I+1, \dots, t-1)$.
 - We obtain $(I', \sigma') \leftarrow \mathcal{A}'(\mathbf{x}, k_0, \dots, k_I, k'_{I+1}, \dots, k'_{t-1}; R)$
 - If $I = I' \wedge k_{I+1} \neq k'_{I+1}$ then return (I, σ, σ') and *success*.
 - Else return *fail*.

Let \mathcal{E} be the event that $F_{\mathcal{A}'}(\mathbf{x})$ outputs *success*. By applying the general forking lemma by Bellare and Neven [BN06] we have that \mathcal{E} happens with probability $(\epsilon^2/t) - (\epsilon/2^{\log^{1+\epsilon}(\lambda)})$. This probability is non-negligible, whenever ϵ is non-negligible.

Now we establish why event \mathcal{E} produces low order element for adversary \mathcal{B} . When \mathcal{E} happens we have $P_I \notin \mathcal{L}_{\text{EXP}}$ and $P_{I+1}, P'_{I+1} \in \mathcal{L}_{\text{EXP}}$. Therefore, if $F_{\mathcal{A}'}(\mathbf{x})$ outputs (I, σ, σ') adversary \mathcal{B} obtains the sextuple $(g, h, \hat{T}, v, r, r')$ with the following properties.

$$h \neq h^{(2^{2^T})} \quad \text{and} \quad (g^r v)^{(2^{\hat{T}})} = v^r h \quad \text{and} \quad (g^{r'} v)^{(2^{\hat{T}})} = v^{r'} h \quad (7)$$

Re-arranging terms of the two equalities on the right we get

$$(g^{(2^{\hat{T}})}/v)^r = h/v^{(2^{\hat{T}})} \quad \text{and} \quad (g^{(2^{\hat{T}})}/v)^{r'} = h/v^{(2^{\hat{T}})} \quad (8)$$

Dividing the left equality by the right we obtain

$$(g^{(2^{\hat{T}})}/v)^{r-r'} = 1 \quad (9)$$

¹Outputting $J \xleftarrow{\$} [-1, t-1]$ only reduces adversary’s success probability with a factor of $t+1$.

Let $u := g^{(2^T)}/v$. Next we establish an upper bound on the order of $u \in \mathbb{G}$. It is guaranteed that $r \neq r'$, since $k_{I+1} \neq k'_{I+1}$. We observe that the order of u can be bounded by:

$$|r - r'| = |(x_{I+1} \cdot 2^{\log^{1+\epsilon}(\lambda)} + k_{I+1}) - (x_{I+1} \cdot 2^{\log^{1+\epsilon}(\lambda)} + k'_{I+1})| = |k_{I+1} - k'_{I+1}| \leq 2^{\log^{1+\epsilon}(\lambda)} \quad (10)$$

Hence, we conclude that $(u, r - r')$ is a pair which breaks the subexponential low order assumption in \mathbb{G} . \square

We remark that one could replace $f(\lambda) = 2^{\log^{1+\epsilon}(\lambda)}$ with any superpolynomial function of λ in Theorem 2.

4 A partial reduction of Factoring to the LO assumption

In this section we prove the equivalence of the LO and the Factoring assumption for non-negligible moduli. We use the LO assumption introduced by Boneh et al [BBF18], cf. Section 2.3. However, the proof enclosed hereby would equally work well for the weaker, subexponential variants of the LO assumption. Specifically, in the following theorem we assume that the RSA-modulus N is generated in a way such that $\phi(N)$ has no prime factor in $(\mathfrak{B}, 2^{\text{poly}(\lambda)}]$ for a constant \mathfrak{B} and $\gcd(p-1, q-1) = 2$. In Section 4.1 we argue that in practice the majority of RSA moduli satisfy these requirements in a typical parameter setting.

Theorem 3. *Let \mathfrak{B} be a fixed integer. The Factoring assumption is reducible in polynomial time to the Low Order assumption for RSA-moduli when $\phi(N)$ has no prime factor between \mathfrak{B} and $2^{\text{poly}(\lambda)}$ and $\gcd(p-1, q-1) = 2$.*

Proof. Let's assume there exists an efficient adversary \mathcal{A} , who can break the LO assumption with non-negligible probability. Express differently, there exists a polynomial $q(\lambda)$, such that

$$\Pr[\mathcal{A} \text{ breaks LO}] \geq \frac{1}{q(\lambda)}. \quad (11)$$

We devise an efficient adversary \mathcal{B} who can factor non-negligible fraction of random RSA moduli by using \mathcal{A} as a subroutine. Adversary \mathcal{B} operates as follows. Upon receiving a random semiprime N it invokes \mathcal{A} on \mathbb{Z}_N^* . By our assumption, adversary \mathcal{A} with non-negligible probability outputs a (u, l) pair such that $u^l \equiv 1 \pmod N$ and $2 \leq l \leq 2^{\text{poly}(\lambda)} \wedge u \neq -1$. Note that, the order l of $u \in (\mathbb{Z}/pq\mathbb{Z})^\times \cong (\mathbb{Z}/p\mathbb{Z})^\times \times (\mathbb{Z}/q\mathbb{Z})^\times$ is the least common multiple of its (multiplicative) orders modulo p and modulo q , ie. $l = \text{lcm}(\text{ord}_p(u), \text{ord}_q(u))$.

Note that, whenever $\text{ord}_p(u) \neq \text{ord}_q(u)$, adversary \mathcal{B} could factor $N = pq$ if l was smooth enough. The reason being that, adversary \mathcal{B} raises u to the power of $\frac{l}{r}$ for all prime factors r of l , until modulo one prime factor of N , but not the other. The order of u divides $\frac{l}{r}$, which can be detected by $0 < \gcd(u^{\frac{l}{r}} - 1 \pmod N, N) < N$, hence factoring the modulus N . In our reduction adversary \mathcal{B} tries to find all prime factors of l in order to find a non-trivial factor of N as described above. This will be the technique employed by adversary \mathcal{B} . Hence, towards our goal one thing that we need to show is that $\text{ord}_p(u) \neq \text{ord}_q(u)$ with non-negligible probability.

First, let us assess the probability when $\text{ord}_p(u) = \text{ord}_q(u)$ for randomly chosen primes p, q . This probability needs to be established as in this case one cannot factor N with the aforementioned technique. We show that $\gcd(\frac{p-1}{2}, \frac{q-1}{2}) = 1$ with constant probability. Since p and q were chosen uniformly random, also $(p-1)/2$ and $(q-1)/2$ behave almost like random integers if we consider their divisibility by other primes (there is one excluded residue class for each prime). The probability of coprimality of random integers is $\frac{1}{\zeta(2)} = \frac{6}{\pi^2} \approx 0.61$. In our case, we have to adjust this value because of the excluded classes, but we still have that $\gcd(p-1, q-1) = 2$ with constant positive probability, therefore whenever $\text{ord}_p(u) = \text{ord}_q(u)$, then this quantity is either 1 or 2. We examine these two cases in more detail.

- $\text{ord}_p(u) = \text{ord}_q(u) = 1$. This is only possible if $u = 1$, which cannot be the case by the definition of the LO assumption, see Definition 1.
- $\text{ord}_p(u) = \text{ord}_q(u) = 2$. In this case $\text{ord}_N(u) = \text{lcm}(\text{ord}_p(u), \text{ord}_q(u)) = 2$. Since $u \notin \{1, -1\}$, therefore u is another non-trivial square root of 1. This also implies that one can factor N as $pq = N|u^2 - 1 = (u-1)(u+1) \wedge u \notin \{1, -1\}$. Therefore $\gcd(N, u-1) = p \vee \gcd(N, u+1) = p$.

Hence, we can conclude that for randomly chosen primes p, q with constant probability $\text{ord}_p(u) \neq \text{ord}_q(u)$.

Second, the goal of adversary \mathcal{B} is to obtain all the prime factors of l . For that end, adversary chooses a constant bound \mathfrak{B} , say 2^{10} . Hence, the adversary would like l to be a \mathfrak{B} -smooth integer in order to be able to

factor it efficiently. Whenever adversary receives order l ($1 \leq l \leq 2^{\text{poly}(\lambda)}$) of u , then adversary would like to find all of its prime factors in a brute force-manner, but still in polynomial-time in λ . Namely, adversary \mathcal{B} wants to find l 's smallest prime factor l_1 , which is smaller than \mathfrak{B} . Suppose a_1 is the largest integer such that $l_1^{a_1} | l$. Then, recursively we would like to find the smallest prime factor of $\frac{l}{l_1^{a_1}}$, denoted l_2 which is smaller than \mathfrak{B} and so on.

This approach succeeds whenever \mathcal{B} receives a pair (u, l) from \mathcal{A} , where l is \mathfrak{B} -smooth. This can be guaranteed if $\phi(N)$ has no prime factors between \mathfrak{B} and $2^{\text{poly}(\lambda)}$. Therefore, we compute now the fraction of those primes up to N , that do not have prime factors between \mathfrak{B} and $2^{\text{poly}(\lambda)}$. We need to establish the fraction of primes up to $N = \mathcal{O}(2^{s(\lambda)})$, that do not have prime factors in $(\mathfrak{B}, 2^{\text{poly}(\lambda)})$. Let us call these integers as low order smooth integers, additionally let \mathcal{P}_{los} be the probability that a randomly chosen integer is low order smooth. We obtain the following asymptotic by applying Equation 5.

$$\mathcal{P}_{\text{los}}(\lambda) = \frac{\Gamma(2^{s(\lambda)}, 2^{\text{poly}(\lambda)}, \mathfrak{B})}{2^{s(\lambda)}} \approx \frac{2^{s(\lambda)} \eta(s(\lambda)/\text{poly}(\lambda), s(\lambda)/\mathfrak{B})}{2^{s(\lambda)}} \geq \frac{s(\lambda)/\text{poly}(\lambda)}{2s(\lambda)/\mathfrak{B}} = \frac{\mathfrak{B}}{2\text{poly}(\lambda)}. \quad (12)$$

Hence, we established that \mathcal{P}_{los} is non-negligible in λ . It follows that the probability that the order of a random RSA-modulus does not have a prime factor in $(\mathfrak{B}, 2^{\text{poly}(\lambda)})$ is $\mathcal{P}_{\text{los}}^2$, i.e. non-negligible. We can establish now the success probability of adversary \mathcal{B} breaking the Factoring assumption.

$$\Pr[\mathcal{B} \text{ breaks Factoring}] \geq \frac{6}{\pi^2} q(\lambda) \mathcal{P}_{\text{los}}^2(\lambda) \quad (13)$$

Therefore, we conclude our proof that the success probability of the efficient adversary \mathcal{B} is non-negligible. \square

We remark that in the reduction we could have allowed $\phi(N)$ to have a single prime factor in $(\mathfrak{B}, 2^{\text{poly}(\lambda)})$. Once \mathcal{B} factors out all the prime factors smaller than \mathfrak{B} from the low order l , adversary \mathcal{B} can establish in probabilistic polynomial time whether the resulting integer is a prime power. If yes, then also in those cases \mathcal{B} can factor efficiently the low order l .

4.1 Practical consequences of the reduction

Hereby, we give an estimate on the portion of RSA moduli used in practice, for which the reduction presented above guarantees an equal level of security for the LO problem, as for the classical Factoring problem. Let $\epsilon = 5$ and $\lambda = 80$, consequently, the size of the moduli should be 1024 bits. OpenSSL is an open-source cryptographic library, which is widely used and the most popular on the internet [ŠNS⁺16]. OpenSSL generates primes in a way that it ensures that no prime from 3 to 17863 divides $p-1$. Let S denote the following set, $S = \{(p-1)/2 | p_i \nmid p-1, 2 \leq i \leq 2048\}$, where p_i is the i th prime. Hence, the probability of coprimality for two randomly sampled integers from S is $\lim_{n \rightarrow \infty} \prod_{i \geq 2049}^n (1 - 1/p_i^2) \approx 0.7499$. Hence we let $\mathfrak{B} = 17863$. The portion of integers having no prime factor between \mathfrak{B} and $\log^6(80)$ amounts to $\frac{\log(\mathfrak{B})}{6 \log \log(80)} = 0.8848$ that can be obtained by using Weingartner's theorem, cf. Equation 5. Overall we have that the probability that an RSA modulus randomly generated by OpenSSL provides the same security guarantees for the LO assumption as for the Factoring is $0.8848 * 0.7499 = 0.6635$. Respectively, for the majority of RSA moduli used in practice if one could find a low order element, then they would be able to factor those moduli as well.

5 Certifying RSA moduli free of low order elements

In certain use cases, e.g. in a public key infrastructure setting or for a VDF, it might be useful if users could prove that their RSA-modulus is free of low order elements. To that end, one could certify RSA moduli applying techniques developed by Goldberg et al [GRSB19].

Specifically, we assume a user wants to prove in zero-knowledge that their RSA modulus N is free of low order elements. Express differently, a user wants to show that $\phi(N)$ has no divisors smaller than a certain bound \mathfrak{B} . Let p_n denote the largest prime smaller than \mathfrak{B} . Then let $e = \prod_{i=1}^n p_i$, where p_i is odd prime, i.e. e is the n th primorial divided by two. Consequently, there cannot be elements mod N with order smaller than \mathfrak{B} if and only if $\gcd(e, N) = 1$. Put differently, N has no low order elements if and only if $x \rightarrow x^e \pmod N$ is a permutation. Applying the zero-knowledge proof system for the language $\mathcal{L}_{\text{perm}\mathbb{Z}_N^*}$ this can be proved efficiently [GRSB19].

In a typical parameter setting (1024-bit RSA moduli, 2^{-80} soundness error for the certification, $\mathfrak{B} = 2^{10}$) the low order RSA public key certification would consist of 50 elements of \mathbb{Z}_N . Hence the size of the proof amounts to 6,4KB. Generating the certification costs roughly 50 full-length RSA exponentiations modulo N , hence it is

also feasible to calculate the proof in a distributed RSA key generation scenario. Meanwhile, each verifier pays the one-time cost of verifying the certification, which is also roughly equal to 50 full-length exponentiations. We note that the number of elements consisting of the proof and number of exponentiations the verifier needs to compute depends only on the admitted soundness error of the proof system and not the size of the moduli. More precisely, both quantities are roughly $\approx \lambda/\log 3$.

6 Open Problems

It is a fascinating open problem to explore more connections between novel and standard RSA assumptions². For instance, it would be fruitful to establish the relation of the Adaptive Root assumption [BBF18] and the (Strong) RSA assumption. The Adaptive Root assumption underpins the security of Wesolowski’s VDF construction [Wes19] and several batching techniques proposed for RSA accumulators [BBF19].

Acknowledgements

We thank Benjamin Wesolowski, Dankrad Feist, Dmitry Khovratovich and Justin Drake for insightful discussions.

References

- [AVD] Vidal Attias, Luigi Vigneri, and Vassil Dimitrov. Implementation study of two verifiable delay functions.
- [BBBF18] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Annual international cryptology conference*, pages 757–788. Springer, 2018.
- [BBF18] Dan Boneh, Benedikt Bünz, and Ben Fisch. A survey of two verifiable delay functions. *IACR Cryptology ePrint Archive*, 2018:712, 2018.
- [BBF19] Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to iops and stateless blockchains. In *Annual International Cryptology Conference*, pages 561–586. Springer, 2019.
- [BGB17] Benedikt Bünz, Steven Goldfeder, and Joseph Bonneau. Proofs-of-delay and randomness beacons in ethereum. *IEEE Security and Privacy on the blockchain (IEEE S&B)*, 2017.
- [BN06] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In *Proceedings of the 13th ACM conference on Computer and communications security*, pages 390–399, 2006.
- [BS13] Eric Bach and Jonathan P Sorenson. Approximately counting semismooth integers. In *Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation*, pages 23–30, 2013.
- [FBGB19] Ben Fisch, Joseph Bonneau, Nicola Greco, and Juan Benet. Scaling proof-of-replication for filecoin mining. Technical report, Technical Report. Stanford University. Accessed May, 2019.
- [GRSB19] Sharon Goldberg, Leonid Reyzin, Omar Sagga, and Foteini Baldimtsi. Efficient noninteractive certification of rsa moduli and beyond. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 700–727. Springer, 2019.
- [LSS19] Esteban Landerreche, Marc Stevens, and Christian Schaffner. Non-interactive cryptographic timestamping based on verifiable delay functions. *IACR Cryptology ePrint Archive*, 2019:197, 2019.
- [MSW] Mohammad Mahmoody, Caleb Smith, and David J Wu. A note on the (im) possibility of verifiable delay functions in the random oracle model.
- [Pie18] Krzysztof Pietrzak. Simple verifiable delay functions. In *10th innovations in theoretical computer science conference (itsc 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

²See: <https://rsa.cash/rsa-assumptions/>

- [RSA78] Ronald L Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.
- [RSS] Lior Rotem, Gil Segev, and Ido Shahaf. Generic-group delay functions require hidden-order groups.
- [ŠNS⁺16] Petr Švenda, Matúš Nemeč, Peter Sekan, Rudolf Kvašňovský, David Formánek, David Komárek, and Vashek Matyáš. The million-key question—investigating the origins of {RSA} public keys. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, pages 893–910, 2016.
- [War90] Richard Warlimont. Sieving by large prime factors. *Monatshefte für Mathematik*, 109(3):247–256, 1990.
- [Wei01] Andreas Weingartner. Integers free of prime divisors from an interval, i. *Acta Arithmetica*, 98:117–131, 2001.
- [Wes19] Benjamin Wesolowski. Efficient verifiable delay functions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 379–407. Springer, 2019.