

On the (in)security of ROS

Fabrice Benhamouda¹, Tancrede Lepoint², Julian Loss³, Michele Orrù⁴, and Mariana Raykova⁵

¹ Algorand Foundation, New York, NY, USA fabrice.benhamouda@gmail.com

² Independent researcher, New York, NY, USA, crypto@tancre.de

³ University of Maryland, College Park, MD, USA, lossjulian@gmail.com

⁴ UC Berkeley, Berkeley, CA, USA michele.orrù@berkeley.edu

⁵ Google, New York, NY, USA marianar@google.com

Abstract. We present an algorithm solving the ROS (Random inhomogeneities in a Overdetermined Solvable system of linear equations) problem mod p in polynomial time for $\ell > \log p$ dimensions. Our algorithm can be combined with Wagner’s attack, and leads to a sub-exponential solution for any dimension ℓ with the best complexity known so far.

When concurrent executions are allowed, our algorithm leads to practical attacks against unforgeability of blind signature schemes such as Schnorr and Okamoto–Schnorr blind signatures, threshold signatures such as GJKR and the original version of FROST, multisignatures such as CoSI and the two-round version of MuSig, partially blind signatures such as Abe–Okamoto, and conditional blind signatures such as ZGP17. Schemes for e-cash (such as Brands’ signature) and anonymous credentials (such as Anonymous Credentials Light) inspired from the above are also affected.

1 Introduction

One of the most fundamental concepts in cryptanalysis is the *birthday paradox*. Roughly, it states that among $O(\sqrt{p})$ random elements from the range $\{0, \dots, p-1\}$ (where p is a prime), there exist two elements a and b such that $a = b$, with high probability. In a seminal work, Wagner [Wag02] gave a generalization of the birthday paradox to ℓ dimensions which asks to find $x_i \in L_i, i \in [\ell]$ such that $x_1 + \dots + x_\ell = 0 \pmod{p}$, where L_i are lists of random elements.

Wagner’s work also showed a simple and elegant algorithm to solve the problem in subexponential time $O((\ell + 1) \cdot 2^{\lceil \log p \rceil / (1 + \lceil \log(\ell + 1) \rceil)})$ and explained how it could be applied to perform cryptanalysis on various schemes. Among the most important applications of Wagner’s technique is a subexponential solution to the ROS (Random inhomogeneities in a Overdetermined Solvable system of linear equations) problem [Sch01, FPS20], which is defined as follows. Given a prime number p and access to a random oracle H_{ROS} with range in \mathbb{Z}_p , the ROS problem (in dimension ℓ) asks to find $(\ell + 1)$ vectors $\hat{\rho}_i \in \mathbb{Z}_p^\ell$ for $i \in [\ell + 1]$, and a vector $\mathbf{c} = (c_1, \dots, c_\ell)$ such that:

$$H_{\text{ROS}}(\hat{\rho}_i) = \langle \hat{\rho}_i, \mathbf{c} \rangle \quad \text{for all } i \in [\ell + 1].$$

This problem was originally studied by Schnorr [Sch01] in the context of blind signature schemes. Using a solver for the ROS problem, Wagner showed that the unforgeability of the Schnorr and Okamoto–Schnorr blind signature schemes can be attacked in subexponential time whenever more than $\text{polylog}(\lambda)$ signatures are issued concurrently. In this work, we revisit the ROS problem and its applications. We make the following contributions.

- We give the first polynomial time solution to the ROS problem for $\ell > \log p$ dimensions.
- We show how the above solution can be combined with Wagner’s techniques to yield an improved subexponential algorithm for dimensions lower than $\log p$. The resulting algorithm offers a smooth trade-off between the work and the dimension needed to solve the ROS problem. It outperforms the runtime of Wagner’s algorithm for a broad range of dimensions.
- Finally, we describe how to apply our new attack to an extensive list of schemes. These include: blind signatures [PS00, Sch01], threshold signatures [GJKR07, KG20a], multisignatures [STV+16, MPSW18a], partially blind signatures [AO00], conditionally blind signatures [ZGP17, GPZZ19], and anonymous

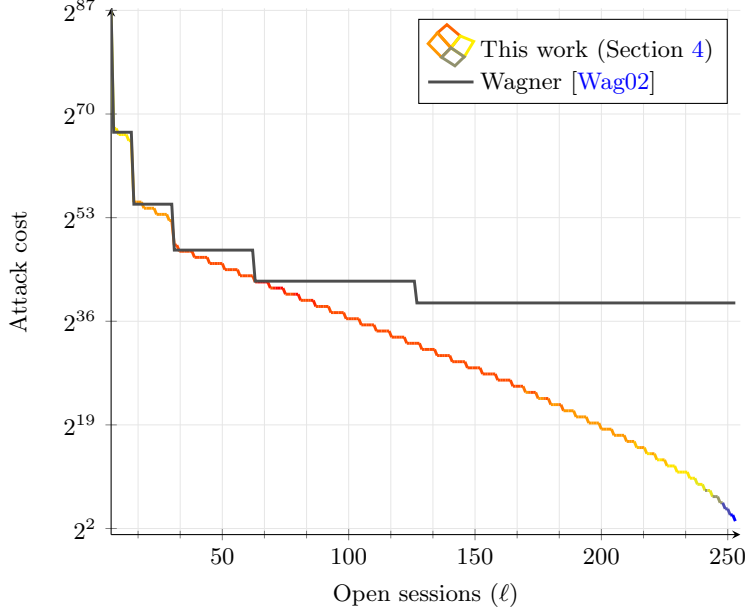


Fig. 1. Concrete cost of our combined attack compared to Wagner’s [Wag02] for $\lambda = 256$ and $\ell < 256$. The color key indicates the different values of w used to estimate the cost. For $\ell \geq 256$, the attack of Section 3 applies.

credentials [BL13, Bra94], in a concurrent setting with $\ell > \log p$ parallel executions. While our attacks do not contradict the security arguments of those schemes (which are restricted only to sequential or bounded number of executions), they prove that these schemes are unpractical for some real-world applications (cf. Section 7).

1.1 Technical overview

Let $\text{Pgen}(1^\lambda)$ be a parameter generation algorithm that given as input the security parameter λ in unary form, outputs an odd prime p of length $\lambda = \lceil \log p \rceil$. In this work, we prove the following main theorem:

Theorem 1 (ROS attack). *If $\ell \geq \lambda$, then there exists an adversary that runs in polynomial time and solves the ROS problem relative to Pgen with dimension ℓ .*

Let us first introduce some notation. Given a polynomial $\boldsymbol{\rho} = \rho_0 + \rho_1 x_1 + \dots + \rho_\ell x_\ell \in \mathbb{Z}_p[x_1, \dots, x_\ell]$ of total degree 1, we denote with $\hat{\boldsymbol{\rho}}$ the vector in \mathbb{Z}_p^ℓ having at the i -th position the coefficient of x_i . It is always possible to find ℓ (out of $\ell + 1$) “partial solutions” to the ROS problem: define the polynomials $\boldsymbol{\rho}_i(\mathbf{x}) = x_i$ in $\mathbb{Z}_p[x_1, \dots, x_\ell]$. Remark that the elements $\hat{\boldsymbol{\rho}}_i$ are the rows of the identity matrix of size ℓ . Define $c_i := H_{\text{ros}}(\hat{\boldsymbol{\rho}}_i)$, such that, for all $i \in [\ell]$, it holds that

$$\langle \hat{\boldsymbol{\rho}}_i, (c_1, \dots, c_\ell) \rangle = H_{\text{ros}}(\hat{\boldsymbol{\rho}}_i) .$$

In general, any list of ℓ polynomials of the form $\boldsymbol{\rho}_i = \rho_{i,i} x_i$ for $\rho_{i,i} \in \mathbb{Z}_p^\times$ ($i \in [\ell]$) is a valid partial solution, as long as $c_i := \rho_{i,i}^{-1} H_{\text{ros}}(\hat{\boldsymbol{\rho}}_i)$. The supposedly computationally hard problem is to find the last partial solution, that is, a non-trivial linear combination $\hat{\boldsymbol{\rho}}_{\ell+1}$ of these values c_i , that matches the hash image $H_{\text{ros}}(\hat{\boldsymbol{\rho}}_{\ell+1})$. Wagner solves the problem in the following way. Fix $\hat{\boldsymbol{\rho}}_{\ell+1} = (1, 1, \dots, 1)$, and build ℓ lists L_1, \dots, L_ℓ such that the i -th list is populated with polynomials of the form $\boldsymbol{\rho}_i = \rho_{i,i} x_i$ for random $\rho_{i,i}$ in \mathbb{Z}_p^\times . For every element in the i -th list L_i , consider its respective coefficient $c_i = \rho_{i,i}^{-1} H_{\text{ros}}(\hat{\boldsymbol{\rho}}_i)$. Build an efficient algorithm

that finds c_i 's satisfying:

$$\langle \hat{\rho}_{\ell+1}, (c_1, \dots, c_\ell) \rangle = c_1 + c_2 + \dots + c_\ell = \mathbf{H}_{\text{ros}}(\hat{\rho}_{\ell+1}).$$

Wagner shows in [Wag02] that the above problem (called the ℓ -list birthday problem) can be solved in time $O((\ell + 1) \cdot 2^{\lceil \log p \rceil / (1 + \lceil \log(\ell+1) \rceil)})$.

However, the ROS problem itself allows for much more flexibility to the attacker: for instance, the attacker can consider a subset of the c_i 's (by setting some entries of $\rho_{\ell+1}$ to zero), in which case we end up with a subset-sum problem that is, in general, NP-hard. In Section 3, we manage to express the ROS problem as a subset-sum of powers of two (modulo p), which can be solved in polynomial time.

Then, to circumvent the restriction $\ell \geq \lambda = \lceil \log p \rceil$, we prove a second theorem, under the same conjecture that the Wagner's algorithm is using (see Section 4.1 for details about the conjecture).

Conjecture 1 (Wagner [Wag02]). Let $L, w \geq 0$ be integers, let p be a odd prime and let $k = 2^w$. Then Wagner's algorithm on k lists of 2^L uniformly random elements in \mathbb{Z}_p (as defined in Fig. 4) has constant failure probability. In particular, when repeating this algorithm in case of failure, the resulting algorithm outputs a solution to the k -list problem over \mathbb{Z}_p in expected time $O(2^{w+L})$.

Theorem 2 (Generalized ROS attack). *Let $L, w \geq 0$ be integers. Under Wagner's conjecture (Conjecture 1), if $\ell \geq \max\{2^w - 1, \lceil 2^w - 1 + \lambda - (w + 1) \cdot L \rceil\}$, then there exists an adversary that runs in expected time $O(2^{w+L})$ and solves the ROS problem relative to \mathbf{P}_{gen} and dimension ℓ .*

The core idea behind the generalized ROS attack is to combine the technique from the attack from Theorem 1 with the basic subexponential attack of Wagner. In the first attack, with a bird's-eye view, the reason why we need $\ell \geq \lambda$ is to be able to write $y = \mathbf{H}_{\text{ros}}(\hat{\rho}_{\ell+1})$ in binary: each bit of the representation corresponds to a power of two in a subset sum which is trivial to solve in polynomial time. However, to make it go through, it uses one dimension (i.e., one c_i) per bit, and y has $\lambda = \lceil \log p \rceil$ bits.

In our generalized ROS attack, instead of writing y entirely in binary as above, which requires λ dimensions, we first find a sum s of 2^w values which include y , but satisfies $|s| \in [0, \frac{p}{2^{(w+1) \cdot L}} - 1] \pmod{p}$. Note that s can then be represented with $\lambda - (w + 1) \cdot L$ many bits in binary representation. This approach requires, in total, $\lceil 2^w + \lambda - (w + 1) \cdot L - 1 \rceil$ dimensions and 2^{w+L} overall work. As illustrated in Fig. 1, this improves over Wagner's attack as the dimension ℓ of the ROS problem increases. We remark that, while in our first attack we give a concrete probability of failure, our second attack is based on the conjecture that Wagner's algorithm for \mathbb{Z}_p succeeds with constant probability. While we are not aware of any formal analysis of Wagner's algorithm over \mathbb{Z}_p , we remark that it is considered a standard cryptanalytic tool. Our attack can be seen as strictly improving over its (conjectured) performance when applied to solve the ROS problem.

1.2 Impact of the attacks

Any cryptographic construction that bases its security guarantees on the hardness of the ROS problem is potentially affected by our attacks.

Blind signatures. An immediate consequence of our findings is the first polynomial-time attack against Schnorr blind signatures [Sch01] and Okamoto–Schnorr blind signatures [PS00] in the concurrent setting with $\ell > \log p$ parallel executions.¹ Structurally, our attack builds on the one shown by Schnorr [Sch01], who showed that a solver to the ROS problem can be turned into an attacker against one-more unforgeability of blind Schnorr and Okamoto–Schnorr signatures. As a concrete example, we implemented in Appendix A the attack of Section 5, illustrating how to break one-more unforgeability of blind Schnorr signatures over 256-bit elliptic curves in a few seconds (when implemented in Sage [S+20]), provided that the attacker can open 256 concurrent sessions.

¹ Okamoto–Schnorr signatures are proven secure only for ℓ parallel executions s.t. $Q^\ell/p \ll 1$, where Q is the number of queries to \mathbf{H}_{ros} . Our attack does not contradict their analysis as our attack requires $\ell > \log_2 p > \log_Q p$.

Game $\text{ROS}_{\text{Pgen,A},\ell}(\lambda)$
$p \leftarrow \text{Pgen}(1^\lambda)$
$((\hat{\rho}_i)_{i \in [\ell+1]}, \mathbf{c}) \leftarrow \mathbf{A}_{\text{ros}}^{\text{H}}(p)$
return $(\forall i \neq j \in [\ell+1] : \hat{\rho}_i \neq \hat{\rho}_j \wedge \langle \hat{\rho}_i, \mathbf{c} \rangle = \text{H}_{\text{ros}}(\hat{\rho}_i))$

Fig. 2. The $\text{ROS}_{\text{Pgen,A},\ell}(\lambda)$ game. H_{ros} is a random oracle with image in \mathbb{Z}_p .

Other affected constructions. Our attack can be adapted to an extensive list of schemes which include threshold signatures [GJKR07, KG20a], multisignatures [STV⁺16, MPSW18a], partially blind signatures [AO00], conditionally blind signatures [ZGP17, GPZZ19], blind anonymous group signatures [CFLW04], blind identity-based signcryption [YW05], and blind signature schemes from bilinear pairings [CHYC05]. We note that some of the previous works claim security only for non-concurrent executions or with a bounded number of executions; therefore, our attacks do not contradict their security claims but render these schemes unsuitable for a broad range of real-world use cases.

Scope of our attacks and countermeasures. Our attacks do not extend to the modified-ROS [FPS20] and the generalized-ROS [HKLN20] problems. The concrete hardness of both problems remains an intriguing open question.

2 Preliminaries

In this work, we assume that logarithm is always base 2, and we use the usual Landau notation. Let $\text{Pgen}(1^\lambda)$ be a parameter generation algorithm that given as input the security parameter λ in unary outputs an odd prime p of length $\lambda = \lceil \log p \rceil$. For an integer q , we let $[q]$ be the integer set $\{1, \dots, q\}$, and \mathbb{Z}_q be the ring of integers modulo q . The ROS problem for ℓ dimensions, displayed in Fig. 2, is *hard* if no adversary can solve the ROS problem in time polynomial in the security parameter λ , i.e.:

$$\text{Adv}_{\text{Pgen,A},\ell}^{\text{ros}}(\lambda) := \Pr[\text{ROS}_{\text{Pgen,A},\ell}(\lambda) = 1] = \lambda^{-\omega(1)}.$$

Alternative formulations of ROS. Fuchsbauer, Plouviez, and Seurin [FPS20, Fig. 7] present a variant of $\text{ROS}_{\text{Pgen,A},\ell}(\lambda)$ with an additional input $\text{aux} \in \{0, 1\}^*$, needed for including the message used in a Schnorr blind signature. Hauck, Kiltz, and Loss [HKL19, Fig. 3] consider an adversary returning a pair $(A, \mathbf{c}) \in \mathbb{Z}_p^{\ell+1 \times \ell+1} \times \mathbb{Z}_p^{\ell+1}$ such that $\mathbf{A}\mathbf{c} = 0$, $A_i \neq A_j \forall i \neq j \in [\ell+1]$, $\text{H}_{\text{ros}}(A_{i,1}, \dots, A_{i,\ell}) = A_{i,\ell+1}$ and $c_{\ell+1} = -1$. These formulations are all equivalent.

3 Attack

We introduce the following notation: for a polynomial $\rho = \rho_0 + \rho_1 x_1 + \rho_2 x_2 + \dots + \rho_\ell x_\ell \in \mathbb{Z}_p[x_1, \dots, x_\ell]$, we set $\hat{\rho}$ to be the vector containing at the i -th position the coefficient of x_i , that is, $\hat{\rho} = (\rho_1, \rho_2, \dots, \rho_\ell)$. Note that the constant term is not included.

Theorem 1 (ROS attack). *If $\ell \geq \lambda$, then there exists an adversary that runs in polynomial time and solves the ROS problem relative to Pgen with dimension ℓ .*

Proof. We construct an adversary for $\text{ROS}_{\text{Pgen,A},\ell}(\lambda)$, where $\ell > \log p$. The goal for the adversary \mathbf{A} is to output $(\hat{\rho}_i)_{i \in [\ell+1]}$ and $\mathbf{c} = (c_1, \dots, c_\ell)$ such that:

$$\text{H}_{\text{ros}}(\hat{\rho}_i) = \langle \hat{\rho}_i, \mathbf{c} \rangle \quad \text{for } i = 1, \dots, \ell + 1.$$

Define:

$$\boldsymbol{\rho}_i^0 := x_i, \quad \boldsymbol{\rho}_i^1 := 2x_i \quad \text{for } i = 1, \dots, \ell,$$

and let $c_i^b := 2^{-b} \mathbf{H}_{\text{ros}}(\hat{\boldsymbol{\rho}}_i^b)$, for $b = 0$ and $b = 1$. If there exists $i^* \in [\ell]$ such that $c_{i^*}^0 = c_{i^*}^1$, then \mathbf{A} stops immediately and returns the ROS solution $(\hat{\boldsymbol{\rho}}_1^0, \dots, \hat{\boldsymbol{\rho}}_\ell^0, \hat{\boldsymbol{\rho}}_{i^*}^1)$ and (c_1^0, \dots, c_ℓ^0) . Otherwise, if $c_i^0 \neq c_i^1$ for all $i \in [\ell]$, define the degree-1 polynomial:

$$\mathbf{f}_i(x_i) := \frac{x_i - c_i^0}{c_i^1 - c_i^0}.$$

We remark that, for $b \in \{0, 1\}$, $\mathbf{f}_i(c_i^b) = b$. Define $\boldsymbol{\rho}_{\ell+1} := \sum_{i=1}^{\ell} 2^{i-1} \mathbf{f}_i$, and note that $\boldsymbol{\rho}_{\ell+1}$ is a multivariate polynomial of total degree 1, i.e., $\boldsymbol{\rho}_{\ell+1} = \rho_{\ell+1,0} + \rho_{\ell+1,1}x_1 + \rho_{\ell+1,2}x_2 + \dots + \rho_{\ell+1,\ell}x_\ell$. Define $y := \mathbf{H}_{\text{ros}}(\hat{\boldsymbol{\rho}}_{\ell+1}) + \rho_{\ell+1,0} = \mathbf{H}_{\text{ros}}((\rho_{\ell+1,1}, \rho_{\ell+1,2}, \dots, \rho_{\ell+1,\ell})) + \rho_{\ell+1,0}$. Finally, write y in binary as:

$$y = \sum_{i=1}^{\ell} 2^{i-1} b_i \pmod{p}.$$

(As $2^\ell > p$, it is possible to write y this way, and this implicitly defines the b_i 's.) The adversary \mathbf{A} outputs the solution $(\hat{\boldsymbol{\rho}}_1^{b_1}, \dots, \hat{\boldsymbol{\rho}}_\ell^{b_\ell}, \hat{\boldsymbol{\rho}}_{\ell+1})$ and $\mathbf{c} := (c_1^{b_1}, \dots, c_\ell^{b_\ell})$. We have indeed that, for $i \in [\ell]$, $\langle \hat{\boldsymbol{\rho}}_i^{b_i}, \mathbf{c} \rangle = 2^{b_i} c_i^{b_i} = \mathbf{H}_{\text{ros}}(\hat{\boldsymbol{\rho}}_i^{b_i})$ and:

$$\langle \hat{\boldsymbol{\rho}}_{\ell+1}, \mathbf{c} \rangle = \boldsymbol{\rho}_{\ell+1}(\mathbf{c}) - \rho_{\ell+1,0} = \sum_{i=1}^{\ell} 2^{i-1} \mathbf{f}_i(c_i^{b_i}) - \rho_{\ell+1,0} = \sum_{i=1}^{\ell} 2^{i-1} b_i - \rho_{\ell+1,0} = \mathbf{H}_{\text{ros}}(\hat{\boldsymbol{\rho}}_{\ell+1}).$$

□

Remark 1. Fuchsbauer, Plouviez, and Seurin [FPS20, Sec. 5] propose a variant of ROS, called modified ROS (mROS). Tessaro and Zhu [TZ22] introduce a variant of ROS called weighted fractional ROS (WFROS). The attack above does not apply to mROS and WFROS.

4 Generalized attack

We present a combination of Wagner's subexponential k -list attack and the polynomial time attack from Section 3. This combined attack yields a subexponentially efficient algorithm against ROS which requires fewer dimensions than the attack in the previous section (i.e., less than $\lambda = \lceil \log p \rceil$). However, for some practical cases, the attack significantly outperforms Wagner's attack in terms of work, for the same number of dimensions. The intuition behind our attack is as follows. We set $k_1 = 2^w - 1$, $k_2 = \max(0, \lceil \lambda - (w+1) \cdot L \rceil)$, and the dimension $\ell = k_1 + k_2$, for some integer w and some real number $L > 0$.

First, we use a generalization of Wagner's algorithm to find a "small" sum $s = y_{k_2}^* + \dots + y_\ell^*$ of $k_1 + 1$ values $y_i^* := \mathbf{H}_{\text{ros}}(\hat{\boldsymbol{\rho}}_i)$, where the polynomials $\boldsymbol{\rho}_i(x_1, \dots, x_\ell)$ are chosen to make the second step of the attack work.² As we describe below, we can obtain that $|s| < 2^{k_2-1}$ using $O(2^{w+L})$ hash queries and space $O(w2^L)$. Then, we use the technique from the previous section in order to represent the sum s as a binary sum of at most k_2 terms. This solves the ROS problem. The attack runs in overall time $O(2^{w+L})$, space $O(w2^L)$, and requires $\ell = \max(2^w - 1, \lceil 2^w - 1 + \lambda - (w+1) \cdot L \rceil)$ dimensions.

We remark that the attack is a generalization of both Wagner's attack and our polynomial-time attack from Section 3. Wagner's attack corresponds to the case where $L = \lambda/(w+1)$ and $\ell = 2^w - 1$. Our polynomial-time attack corresponds to the case $w = 0$, $L = 0$, $\ell = \lambda$.

² In the actual attack, part of the second step is executed before to allow to choose these polynomials properly.




Fig. 3. Tree of lists for the k -list algorithm (\bowtie represents the join operation in the algorithm; the sets in the right handside are the sets to which the elements of the lists of a given level belong).

Examples. For a prime p of $\lambda = 256$ bits, a concrete example yields $w = 5, L = 15$, i.e., $\ell = 32 + 256 - 6 \cdot 15 - 1 = 197$ dimensions and time roughly 2^{20} and space roughly $5 \cdot 2^{15}$ (elements of \mathbb{Z}_p). On the other hand, Wagner's algorithm for 197 dimensions requires time roughly $2^{\lceil \log 197 \rceil} \cdot 2^{\frac{256}{\lceil \log 197 \rceil + 1}} = 2^7 \cdot 2^{32} = 2^{39}$ and space roughly $\lceil \log 197 \rceil \cdot 2^{\frac{256}{\lceil \log 197 \rceil + 1}} = 7 \cdot 2^{32}$.

For a 512 bit modulus, a concrete example yields $w = 6, L = 46$, i.e., $\ell = 64 + 512 - 7 \cdot 46 - 1 = 253$ dimensions and time roughly 2^{53} and space roughly $6 \cdot 2^{46}$. Wagner's algorithm for 254 dimensions requires time roughly $2^{\lceil \log 254 \rceil} \cdot 2^{\frac{512}{\lceil \log 254 \rceil + 1}} = 2^7 \cdot 2^{64} = 2^{71}$ and space roughly $\lceil \log 254 \rceil \cdot 2^{\frac{512}{\lceil \log 254 \rceil + 1}} = 7 \cdot 2^{64}$.³

4.1 Generalized k -list algorithm

In this section, we write elements \mathbb{Z}_p as signed integers in $[-\frac{p-1}{2}, \frac{p-1}{2}]$. Let w and L be two positive integers. We define the following integer intervals:

$$I_i := \left[- \left\lfloor \frac{p-1}{2^{(w-i) \cdot L + 1}} \right\rfloor, \left\lfloor \frac{p-1}{2^{(w-i) \cdot L + 1}} \right\rfloor \right].$$

Remark that $\mathbb{Z}_p = I_w$.

We now describe the k -list algorithm, which is the core of Wagner's algorithm. We generalize it to match our needs and to output elements that sum to something in I_{-1} rather than to exactly 0. (This essentially corresponds to executing Wagner's attack as usual, but stopping early.) The algorithm is defined relative to random oracles H_1, \dots, H_k (with input in \mathbb{Z}_p^* and output in \mathbb{Z}_p). It takes as input (w, L) and outputs $(\rho_1^*, \dots, \rho_k^*) \in (\mathbb{Z}_p^*)^k$ with $k = 2^w$ such that:

$$s := y_1^* + \dots + y_k^* \in I_{-1} \quad \text{where } y_i^* := H_i(\rho_i^*) .$$

The high-level idea of the algorithm is to use $2^{w+1} - 1$ lists of about 2^L values organized as a tree, as depicted in Fig. 3, and to ensure that lists \mathfrak{L}_i^w at level i contains elements from the set I_i .

³ Indeed, when considering the exact values of the constants in the asymptotics, the actual complexity of Wagner's attack is $2^{\lceil \log(\ell+1) \rceil} \cdot 2^{\frac{\lambda}{\lceil \log(\ell+1) \rceil + 1}}$.

```

Algorithm  $k$ -list $^{\text{H}_1, \dots, \text{H}_k}(w, L)$ 
// Setup
for  $i \in [k]$ ,  $\rho \in [2^L]$ :  $\mathfrak{L}_{i,j}^w := \text{H}_i(\rho)$ 
// Collisions
for  $i = w$  downto 1:
  for  $j \in [2^{i-1}]$ :
     $\mathfrak{L}_j^{i-1} = \{a + b : a \in \mathfrak{L}_{2j-1}^i, b \in \mathfrak{L}_{2j}^i, a + b \in I_{i-1}\}$ 
// Output
look for an element  $s = y_1^* + \dots + y_k^* \in \mathfrak{L}^0 \cap I_{-1}$ 
if such an element does not exist then return  $\perp$ 
return  $(\rho_1^*, \dots, \rho_k^*)$  such that  $y_i^* = \text{H}_i(\rho_i^*)$ 

```

Fig. 4. The k -list algorithm.

- **Setup/Leaves:** k -list fills the lists \mathfrak{L}_i^w in the leaves with all 2^L points of the form $\text{H}_i(\rho) \in I_w = \mathbb{Z}_p$, where $\rho \in \{1, 2, 3, \dots, 2^L\}$.
- **Collisions/Join:** The algorithm now proceeds to find collisions in levels from w to 1. At level i , process the 2^{i-1} pairs of lists $(\mathfrak{L}_1^i, \mathfrak{L}_2^i), \dots, (\mathfrak{L}_{2^{i-1}-1}^i, \mathfrak{L}_{2^{i-1}}^i)$ into 2^{i-1} lists $\mathfrak{L}_1^{i-1}, \dots, \mathfrak{L}_{2^{i-1}}^{i-1}$ as follows:

$$\mathfrak{L}_j^{i-1} := \{a + b \quad : \quad a \in \mathfrak{L}_{2j-1}^i, b \in \mathfrak{L}_{2j}^i, a + b \in I_{i-1}\} .$$

(Remember that $a, b \in \mathbb{Z}_p$ and $a + b$ is computed modulo p .) Moreover, we implicitly assume that the algorithm stores back pointers to a and b such that they can efficiently be recovered at a later point.

- **Output:** Let $\mathfrak{L}^0 = \mathfrak{L}_1^0$ denote the (only) list created at level 0. The algorithm finds an element $s \in \mathfrak{L}^0$ such that $s \in I_{-1}$. If no such element exists, it returns \perp . Otherwise, it recovers $k = 2^w$ values $\rho_1^*, \dots, \rho_k^*$ such that $y_i^* = \text{H}_i(\rho_i^*) \in \mathfrak{L}_i^w$ and $s = y_1^* + \dots + y_k^*$. It returns $(\rho_1^*, \dots, \rho_k^*)$.

We formally write the algorithm k -list in Fig. 4.

Correctness. Our algorithm’s correctness follows directly from the correctness of Wagner’s original algorithm. More precisely, our algorithm performs identical steps as Wagner’s, but stops upon finding a sum of values with a suitably small absolute value, i.e., one that falls into I_{-1} . On the other hand, Wagner’s algorithm keeps continuing with more levels until it finds values who sum to 0.

We remark that we are not aware of a formal analysis of Wagner’s algorithm for values in \mathbb{Z}_p . The work of Minder and Sinclair [MS09] analyses the case of finding a weighted sum of *vectors* of \mathbb{Z}_p values that sum to zero in each component, but uses a different technique from the one presented in Wagner’s paper (and used here). Our attack can be seen as working under the assumption that Wagner’s algorithm works correctly, i.e., has constant failure probability. We can repeat the attack until it succeeds, which makes the resulting algorithm expected polynomial time. Formally analyzing the failure probability of Wagner’s algorithm over \mathbb{Z}_p remains an important open problem.

Complexity. Overall, the algorithm runs in time $O(2^{w+L})$ and is conjectured to succeed with constant probability. As described [Wag02], this running time is made possible using an optimized join operation such as Hash Join or Merge Join. The algorithm uses space $O(2^{w+L})$, but by evaluating the collisions/joins in postfix order (in the tree), this can be reduced to $O(w2^L)$.

Concretely, we state the following conjecture:

Conjecture 1 (Wagner [Wag02]). Let $L, w \geq 0$ be integers, let p be an odd prime and let $k = 2^w$. Then Wagner’s algorithm on k lists of 2^L uniformly random elements in \mathbb{Z}_p (as defined in Fig. 4) has constant failure probability. In particular, when repeating this algorithm in case of failure, the resulting algorithm outputs a solution to the k -list problem over \mathbb{Z}_p in expected time $O(2^{w+L})$.

4.2 Combined attack

Theorem 2 (Generalized ROS attack). *Let $L, w \geq 0$ be integers. Under Wagner's conjecture (Conjecture 1), if $\ell \geq \max\{2^w - 1, \lceil 2^w - 1 + \lambda - (w + 1) \cdot L \rceil\}$, then there exists an adversary that runs in expected time $O(2^{w+L})$ and solves the ROS problem relative to \mathbf{Pgen} and dimension ℓ .*

Proof. Recall that $k_1 = 2^w - 1$ and $k_2 = \max(0, \lceil \lambda - (w + 1) \cdot L \rceil)$. Set $\ell = k_1 + k_2$. For all $i \in [k_2]$, define:

$$\rho_i^0 := x_i \quad \rho_i^1 := 2x_i,$$

and define $c_i^0 := H_{\text{ros}}(\hat{\rho}_i^0)$ and $c_i^1 := 2^{-1}H_{\text{ros}}(\hat{\rho}_i^1)$. If there exists $i^* \in [k_2]$ such that $c_{i^*}^0 = c_{i^*}^1$, then the adversary already found a non-trivial ROS solution. Define, for $i \in [k_2 + 1, \ell]$, $\rho_i := x_i$ and $c_i := H_{\text{ros}}(\hat{\rho}_i)$. Output the ROS solution $(\hat{\rho}_1^0, \dots, \hat{\rho}_{k_2}^0, \hat{\rho}_{k_2+1}, \dots, \hat{\rho}_\ell, \hat{\rho}_{i^*}^1)$, and $(c_1^0, \dots, c_{k_2}^0, c_{k_2+1}, \dots, c_\ell)$. Otherwise, let:

$$\mathbf{f}_i := \frac{x_i - c_i^0}{c_i^1 - c_i^0}$$

for all $i \in [k_2]$. We remark that $\mathbf{f}_i(c_i^b) = b$ (for $b = 0, 1$). Define:

$$\bar{\rho}_{\ell+1}(x_1, \dots, x_\ell) := \sum_{i=1}^{k_2} 2^{i-1} \mathbf{f}_i + \left\lfloor \frac{p-1}{2^{(w+1) \cdot L+1}} \right\rfloor - \sum_{i=k_2+1}^{\ell} x_i .$$

Run $(\rho_{k_2+1}^*, \dots, \rho_{\ell+1}^*) := k\text{-list}^{\text{H}_{k_2+1}, \dots, \text{H}_{\ell+1}}(w, L)$, where $k = \ell - k_2 + 1 = k_1 + 1 = 2^w$ and the oracles are defined as:

$$\text{H}_i(\alpha) := \begin{cases} \text{if } i \in [k_2 + 1, \ell]: & \text{let } \mathbf{p} = \alpha x_i \quad \text{return } \alpha^{-1} \text{H}_{\text{ros}}(\hat{\mathbf{p}}) \\ \text{if } i = \ell + 1: & \text{let } \mathbf{p} = \alpha \bar{\rho}_{\ell+1} \quad \text{return } \alpha^{-1} \text{H}_{\text{ros}}(\hat{\mathbf{p}}) + \bar{\rho}_{\ell+1,0} \end{cases}$$

Define (similarly to the above):

$$\begin{aligned} \rho_i^* &:= \rho_i^* x_i, & y_i^* &:= \text{H}_i(\rho_i^*) = (\rho_i^*)^{-1} \text{H}_{\text{ros}}(\hat{\rho}_i^*) , & \text{for } i \in [k_2 + 1, \ell]; \\ \rho_{\ell+1}^* &:= \rho_{\ell+1}^* \bar{\rho}_{\ell+1}, & y_{\ell+1}^* &:= \text{H}_{\ell+1}(\rho_{\ell+1}^*) = (\rho_{\ell+1}^*)^{-1} \text{H}_{\text{ros}}(\hat{\rho}_{\ell+1}^*) + \bar{\rho}_{\ell+1,0} . \end{aligned} \quad (1)$$

Set:

$$s := \sum_{i=k_2+1}^{\ell+1} y_i^* \in I_{-1} = \left[- \left\lfloor \frac{p-1}{2^{(w+1) \cdot L+1}} \right\rfloor, \left\lfloor \frac{p-1}{2^{(w+1) \cdot L+1}} \right\rfloor \right] . \quad (2)$$

Write $s + \lfloor (p-1)/2^{(w+1) \cdot L+1} \rfloor$ in binary as:

$$s + \left\lfloor \frac{p-1}{2^{(w+1) \cdot L+1}} \right\rfloor = \sum_{i=1}^{k_2} 2^{i-1} b_i \in \left[\left\lfloor \frac{p-1}{2^{(w+1) \cdot L}} \right\rfloor \right] , \quad (3)$$

which is possible since $p < 2^\lambda$, $k_2 \geq \lceil \lambda - (w + 1) \cdot L \rceil$, hence $(p-1)/2^{(w+1) \cdot L} < 2^{k_2}$. Define:

$$\hat{\rho}_i := \begin{cases} \hat{\rho}_i^{b_i} & \text{for } i \in [1, k_2] , \\ \hat{\rho}_i^* & \text{for } i \in [k_2 + 1, \ell + 1] . \end{cases}$$

and:

$$c_i := \begin{cases} c_i^{b_i} & \text{for } i \in [1, k_2] , \\ y_i^* & \text{for } i \in [k_2 + 1, \ell] . \end{cases}$$

A outputs $(\hat{\rho}_1, \dots, \hat{\rho}_{\ell+1})$ and $\mathbf{c} = (c_1, \dots, c_\ell)$.

We have indeed that for $i \in [\ell]$:

$$\langle \hat{\rho}_i, \mathbf{c} \rangle = \begin{cases} \rho_i^{b_i}(\mathbf{c}) - \rho_{i,0}^{b_i} = \rho_i^{b_i}(\mathbf{c}) = 2^{-b_i} c_i = \mathbf{H}_{\text{ros}}(\hat{\rho}_i^{b_i}) & \text{for } i \in [1, k_2] , \\ \rho_i^*(\mathbf{c}) - \rho_{i,0}^* = \rho_i^*(\mathbf{c}) = \rho_i^* y_i^* = \mathbf{H}_{\text{ros}}(\hat{\rho}_i^*) & \text{for } i \in [k_2 + 1, \ell] . \end{cases}$$

since $\rho_{i,0}^{b_i}$ and $\rho_{i,0}^*$ are the constant coefficients of the polynomials $\rho_i^{b_i} = 2^{b_i} x_i$ and $\rho_i^* = \rho_i^* x_i$ (respectively) and hence are zero. For the case $\ell + 1$:

$$\begin{aligned} \langle \hat{\rho}_{\ell+1}, \mathbf{c} \rangle &= \rho_{\ell+1}^*(\mathbf{c}) - \rho_{\ell+1,0}^* = \rho_{\ell+1}^* \cdot \left(\sum_{i=1}^{k_2} 2^{i-1} \mathbf{f}_i(\mathbf{c}) - \left\lfloor \frac{p-1}{2^{(w+1) \cdot L+1}} \right\rfloor - \sum_{i=k_2+1}^{\ell} c_i - \bar{\rho}_{\ell+1,0} \right) \\ &= \rho_{\ell+1}^* \cdot \left(\sum_{i=1}^{k_2} 2^{i-1} b_i - \left\lfloor \frac{p-1}{2^{(w+1) \cdot L+1}} \right\rfloor - \sum_{i=k_2+1}^{\ell} y_i^* - \bar{\rho}_{\ell+1,0} \right) \\ &= \rho_{\ell+1}^* \cdot \left(s - \sum_{i=k_2+1}^{\ell} y_i^* - \bar{\rho}_{\ell+1,0} \right) = \rho_{\ell+1}^* \cdot (y_{\ell+1}^* - \bar{\rho}_{\ell+1,0}) = \mathbf{H}_{\text{ros}}(\hat{\rho}_{\ell+1}) . \end{aligned}$$

where the second equality comes from Equation (1) (and hence the constant coefficient $\rho_{\ell+1,0}^* = \rho_{\ell+1}^* \bar{\rho}_{\ell+1,0}$), the fourth equality comes from Equation (3), the fifth equality comes from Equation (2), the last equality comes from Equation (1). The attack requires $k_1 + k_2 = \max\{2^w - 1, \lceil 2^w - 1 + \lambda - (w+1) \cdot L \rceil\}$ dimensions, runs in time $O(2^{w+L})$, and in space $O(w2^L)$. \square

5 Affected blind signatures

For simplicity and clarity of exposition, we explain how to instantiate the attack presented in Section 3 only. Our attack can be easily adapted for the one presented in Section 4.

Throughout the remaining of this manuscript, we assume the existence of a group generator algorithm $\text{GrGen}(1^\lambda)$ that, given as input the security parameter in unary form outputs the description $\Gamma = (\mathbb{G}, p, G)$ of a group \mathbb{G} of prime order p generated by G . Similarly to Section 2, we assume that the prime p is of length λ . We use additive notation for the group law.

5.1 Schnorr blind signatures

In Schnorr blind signatures [Sch01, FPS20], a signing key x is a scalar sampled uniformly at random from \mathbb{Z}_p , and its respective verification key is $X = xG$ in the group \mathbb{G} . A signature for a message $m \in \{0, 1\}^*$ consists of a pair $(R, s) \in \mathbb{G} \times \mathbb{Z}_p$ such that $sG - cX = R$, where $c := \mathbf{H}(R, m)$. Fig. 5 depicts the protocol. A formal description of the protocol can be found in [FPS20, Fig. 6], using the same notation employed here.

We construct a probabilistic polynomial-time adversary \mathbf{A} that is able to produce (with overwhelming probability) $\ell + 1$ signatures after opening $\ell \geq \lceil \log p \rceil = \lambda$ parallel sessions. \mathbf{A} selects an arbitrary message $m_\ell \in \{0, 1\}^*$ for which a signature will be forged. It opens ℓ sessions, obtaining the first message from the server and receiving $\bar{\mathbf{R}} = (\bar{R}_1, \dots, \bar{R}_\ell) \in \mathbb{G}^\ell$. For $i \in [\ell + 1]$ and $b \in \{0, 1\}$, sample uniformly at random two blinding factors $(\alpha_{i,b}, \beta_{i,b}) \leftarrow_s \mathbb{Z}_p^2$, and define $R_{i,b} := \bar{R}_i + \alpha_{i,b}G + \beta_{i,b}X$. Let m_i be an arbitrary message and let $c_i^b := \mathbf{H}(R_{i,b}, m_i)$ for $i \in [\ell]$ and $b \in \{0, 1\}$. Assume $c_i^0 \neq c_i^1$. Define the polynomial $\rho \in \mathbb{Z}_p[x_1, \dots, x_\ell]$:

$$\rho(x_1, \dots, x_\ell) := \sum_{i=1}^{\ell} 2^{i-1} \cdot \frac{x_i - c_i^0}{c_i^1 - c_i^0} = \sum_{i=1}^{\ell} \rho_i x_i + \rho_0 . \quad (4)$$

We recall that, from Theorem 1, the above polynomial is such that, for any $(b_1, \dots, b_\ell) \in \{0, 1\}^\ell$, $\rho(c_1^{b_1}, \dots, c_\ell^{b_\ell}) = \sum_i 2^{i-1} b_i$. Let $R_\ell := \sum_{i=1}^{\ell} \rho_i \bar{R}_i$ (nota bene: the constant term ρ_0 is not included), and select an arbitrary




Fig. 5. The signing protocol of Schnorr blind signatures [Sch01].

message $m_{\ell+1}$. Define $c_{\ell+1} := H(R_{\ell+1}, m_{\ell+1})$ and consider the binary decomposition of $c_{\ell+1} + \rho_0$:

$$c_{\ell+1} + \rho_0 = \sum_{i=1}^{\ell} 2^{i-1} b_i. \quad (5)$$

Let $\bar{c} = (c_1^{b_1} + \beta_{1,b_1}, \dots, c_\ell^{b_\ell} + \beta_{\ell,b_\ell})$. Complete the ℓ opened sessions with \bar{c} : reply to the i -th open session with \bar{c}_i , for $i \in [\ell]$. The adversary thus obtains responses $\bar{s} := (\bar{s}_1, \dots, \bar{s}_\ell) \in \mathbb{Z}_p^\ell$, and defines $s_{\ell+1} := \sum_{i=1}^{\ell} \rho_i \bar{s}_i$. Finally, the adversary proceeds unblinding the ℓ honest signatures by computing: $\mathbf{s} := (\bar{s}_1 + \alpha_{1,b_1}, \dots, \bar{s}_\ell + \alpha_{\ell,b_\ell})$. The adversary outputs the $\ell + 1$ forgeries $(m_i, (R_i, s_i))_{i \in [\ell+1]}$, that is:

$$(R_i, s_i) = \begin{cases} (\bar{R}_i + \alpha_{i,b_i} G + \beta_{i,b_i} X, \bar{s}_i + \alpha_{i,b_i}) & \text{for } i = 1, \dots, \ell, \\ (\sum_{i=1}^{\ell} \rho_i R_i, \sum_{i=1}^{\ell} \rho_i s_i) & \text{for } i = \ell + 1. \end{cases}$$

By perfect correctness, we have that the first ℓ signatures are valid. In fact, for $i \in [\ell]$:

$$R_i = \bar{R}_i + \alpha_{i,b_i} G + \beta_{i,b_i} X = \bar{s}_i G - \bar{c}_i X + \alpha_{i,b_i} G + \beta_{i,b_i} X = s_i G - c_i^{b_i} X,$$

and $c_i^{b_i} = H(R_i, m_i)$. For the case $(m_{\ell+1}, (R_{\ell+1}, s_{\ell+1}))$:

$$R_{\ell+1} = \sum_{i=1}^{\ell} \rho_i \bar{R}_i = \sum_{i=1}^{\ell} \rho_i (\bar{s}_i G - \bar{c}_i^{b_i} X) = s_{\ell+1} G - c_{\ell+1} X,$$

where $c_{\ell+1} = H(R_{\ell+1}, m_{\ell+1})$ and the last equality come from:

$$\sum_{i=1}^{\ell} \rho_i \bar{c}_i^{b_i} = \rho(c_1^{b_1}, \dots, c_\ell^{b_\ell}) - \rho_0 = \sum_{i=1}^{\ell} 2^{i-1} b_i - \rho_0 = c_{\ell+1}$$

where the first equality come from Equation (4) last equality comes from Equation (5). We provide an implementation of the attack in Appendix A using Sage [S+20].

5.2 Okamoto–Schnorr blind signatures

An Okamoto–Schnorr blind signature [PS00] for a message m consists of a tuple $(R, s, t) \in \mathbb{G} \times \mathbb{Z}_p^2$ such that $sG + tH - cX = R$, where $c := H(R, m)$, and (G, H) are two generators of \mathbb{G} whose discrete log relation is unknown. The attack of the previous section directly extends to Okamoto–Schnorr signatures: A operates exactly as before until Equation (4). Then, the forgery is constructed as:

$$(R_\ell := \sum_{i=1}^{\ell} \rho_i R_i, \quad s_\ell := \sum_{i=1}^{\ell} \rho_i s_i, \quad t_\ell := \sum_{i=1}^{\ell} \rho_i t_i).$$

We stress again that this does not contradict the security analysis of Stern and Pointcheval [PS00], whose security was reduced to $\text{DLOG}_{\text{GrGen}, A}(\lambda)$ for a $\text{polylog}(\lambda)$ number of queries.

6 Other constructions affected

In this section, we overview how the attacks presented in Sections 3 and 4 apply to a number of other cryptographic primitives. To simplify exposition, we focus on adapting the attack of Section 3. We note that, in some cases (e.g., multi-signatures), we break the security claims of the papers, while for other primitives (e.g., threshold signatures), our attack illustrates the tightness of the security theorems, which assume either non-concurrent setting, or up to a logarithmic number of concurrent executions.

For multi-signatures and threshold signatures, we show potential fixes that thwart the attack. Intuitively, the idea of these countermeasures is to prevent the adversary from adaptively selecting its commitment rG in the protocol from the honest parties' commitments. The simplest solution is to add one initial round of communication where each party sends a commitment to their commitment (e.g., $H(rG)$ with H modeled as a random oracle), and then reveal the original commitment rG in a second round. However, for blind signatures and their variants, such a simple fix does not work, as the adversary's forgery is made on a commitment that the adversary does not need to reveal until the end of the security game (and hence cannot be committed to).

6.1 Multi-signatures

A multi-signature scheme allows a group of signers S_1, \dots, S_n , each having their own key pair $(\text{pk}_j, \text{sk}_j)$, to collaboratively sign a message m . The resulting signature can be verified given the message and the set of public keys of all signers.

6.1.1 CoSi

CoSi is a multi-signature scheme introduced by Syta et al. [STV⁺16] which features a two-round signing protocol. The signers are organized in a tree structure, where S_1 is the root of the tree. A signature for a message $m \in \{0, 1\}^*$ consists of a pair $(c, s) \in \mathbb{Z}_p^2$ such that $c = H(sG - c \cdot \text{pk}, m)$, where $\text{pk} = \sum_{j=1}^n \text{pk}_j \in \mathbb{G}$ is the aggregated verification key. A formal description of the protocol can be found in [DEF⁺19, Sec. 2.5]; we use the same notation, except that we employ additive notation xG instead of multiplicative notation g^x .

Attack. We present an attack for a two-node tree where the attacker controls the root S_1 . The attack can easily be extended to other settings, similarly to [DEF⁺19, Sec. 4.2]. Our attack allows the signer S_1 to forge one signature, for an arbitrary message $m_\ell \in \{0, 1\}^*$, after performing $\ell > \log p$ interactions with the honest signer S_2 . Recall that $\text{pk} = \text{pk}_1 + \text{pk}_2$ where $\text{pk}_i = \text{sk}_i G$. The signing protocol proceeds as follows. First, S_1 obtains a commitment $t_2 = r_2 G$ from S_2 , and computes $\bar{t} = t_1 = r_1 G + t_2$ for a random r_1 . Then, S_1 computes the challenge $c = H(\bar{t}, m)$, and sends (\bar{t}, c) to S_2 . Next, S_2 returns $s_2 := r_2 + c \cdot \text{sk}_2$. Finally, S_1 computes $s := s_2 + r_1 + c \cdot \text{sk}_1$ and outputs the signature (c, s) for the message m .

The attack proceeds as follows. S_1 opens ℓ parallel sessions with ℓ arbitrary distinct messages $m_1, \dots, m_\ell \in \{0, 1\}^*$. For each session, S_1 gets the commitments $t_i = r_i G$ from S_2 at the end of the first round

of signing. Now, it samples two random values $r_{i,0}, r_{i,1}$ for each $i \in [\ell]$, and defines $\bar{t}_i^0 = r_{i,0}G + t_i$ and $\bar{t}_i^1 = r_{i,1}G + t_i$, and computes $c_i^b = \text{H}(\bar{t}_i^b, m_i)$. (As usual, if $c_i^0 = c_i^1$, S_1 samples again $r_{i,0}$ and $r_{i,1}$ until $c_i^0 \neq c_i^1$.) S_1 then defines the polynomial $\rho := \sum_{i=1}^{\ell} 2^{i-1} x_i / (c_i^1 - c_i^0)$, computes $t_{\ell+1} := \rho(t_1, \dots, t_\ell)$ and $c_{\ell+1} := \text{H}(t_{\ell+1}, m_{\ell+1})$. S_1 computes $d_{\ell+1} = c_{\ell+1} - \rho(c_1^0, \dots, c_\ell^0)$ and writes this value in binary as $d_{\ell+1} = \sum_{i=1}^{\ell} 2^{i-1} b_i$. It then closes the ℓ sessions by using $\bar{t}_i = \bar{t}_i^{b_i}$ and $c_i = c_i^{b_i}$. At the last step of the signing sessions, S_1 obtains values $s_i = r_i + c_i \cdot \text{sk}_2$ from S_2 , and closes the sessions honestly using r_{i,b_i} . Finally, S_1 concludes its forgery by defining $s_{\ell+1} := \rho(\mathbf{s}) + c_{\ell+1} \cdot \text{sk}_1$: the pair $(c_{\ell+1}, s_{\ell+1})$ is a valid signature for $m_{\ell+1}$. In fact:

$$\begin{aligned}
s_{\ell+1}G - c_{\ell+1} \cdot \text{pk} &= (\rho(\mathbf{s}) + c_{\ell+1} \cdot \text{sk}_1)G - c_{\ell+1} \cdot \text{pk} \\
&= \sum_{i=1}^{\ell} \frac{2^{i-1} s_i}{c_i^1 - c_i^0} G - c_{\ell+1} \cdot \text{pk}_2 \\
&= \sum_{i=1}^{\ell} \frac{2^{i-1} (r_i + c_i^{b_i} \cdot \text{sk}_2)}{c_i^1 - c_i^0} G - c_{\ell+1} \cdot \text{pk}_2 \\
&= \sum_{i=1}^{\ell} \frac{2^{i-1} r_i}{c_i^1 - c_i^0} G + \left(\sum_{i=1}^{\ell} \frac{2^{i-1} c_i^{b_i}}{c_i^1 - c_i^0} - c_{\ell+1} \right) \cdot \text{pk}_2 \\
&= \sum_{i=1}^{\ell} \frac{2^{i-1} t_i}{c_i^1 - c_i^0} + \left(\sum_{i=1}^{\ell} 2^{i-1} b_i + \sum_{i=1}^{\ell} \frac{2^{i-1} c_i^0}{c_i^1 - c_i^0} - c_{\ell+1} \right) \cdot \text{pk}_2 \\
&= \sum_{i=1}^{\ell} \frac{2^{i-1} t_i}{c_i^1 - c_i^0} + \underbrace{\left(\sum_{i=1}^{\ell} 2^{i-1} b_i + \rho(c_1^0, \dots, c_\ell^0) - c_{\ell+1} \right)}_{=d_{\ell+1} - d_{\ell+1} = 0} \cdot \text{pk}_2 \\
&= \rho(t_1, \dots, t_\ell) = t_{\ell+1},
\end{aligned}$$

and $c_{\ell+1} = \text{H}(t_{\ell+1}, m_{\ell+1})$ by definition.

6.1.2 Two-Round MuSig

As in [DEF⁺19], the above technique (with some minor modifications) can be applied to the two-round MuSig as initially proposed by Maxwell et al. [MPSW18a], as the main difference between CoSi and two-round MuSig is in how the public key is aggregated in order to avoid rogue-key attacks.

6.1.3 Three-Round MuSig and fix

Our attack does not apply to the updated MuSig that uses a 3-round signing algorithm [MPSW18b]. The fix in the 3-round MuSig follows the blueprint of the simple countermeasure we described at the beginning of the section. What makes the above attack possible is that the adversary sees the commitment $t_2 = r_2G$ of the honest party S_2 before choosing its own commitment r_1G (which defines $\bar{t} = r_1G + t_2$ and $c = \text{H}(\bar{t}, m)$). To prevent the attack, in 3-round MuSig, all parties S_j first commit to their commitment r_jG , and then reveal r_jG in a second round. This additional first round of commitment of commitment ensures that S_1 cannot choose its commitment r_1G depending on S_2 's commitment r_2G , which thwarts the attack.

6.2 Threshold signatures

A (t, n) -threshold signature scheme assumes that the secret signing key is split among n parties P_1, \dots, P_n in a way that allows any subset of at least t out of the n parties to produce a valid signature. As long as the adversary corrupts less than the threshold number of parties, it is not possible to forge signatures or learn any information about the signing key.

6.2.1 GJKR07

Gennaro, Jarecki, Krawczyk, Rabin proposed a threshold signature scheme based on Pedersen’s distributed key generation (DKG) protocol in [GJKR07, Section 5.2]. At a very high level, Pedersen’s DKG protocol allows to generate a random group element $X = \chi G$ so that its discrete logarithm χ is shared both additively and according to the Feldman secret sharing [Fel87] scheme, between a set of “qualified” parties. For the attack we present below, all parties P_1, \dots, P_n (included the ones that are controlled by the adversary) will remain qualified.⁴ We denote by χ_j the additive share of party P_j . We have $\chi = \sum_{j=1}^n \chi_j$. Importantly for the attack, the adversary controlling for example P_1 , can see all the group elements $\chi_2 G, \dots, \chi_n G$ and then can choose its value χ_1 . This is due to the way the Feldman secret sharing is performed.

In the threshold signature scheme of Gennaro et al. [GJKR07], the parties execute a distributed key generation procedure to produce a verification key $\text{pk} := \text{sk} \cdot G \in \mathbb{G}$, where the secret key sk is additively shared between the parties: each party P_j has an additive share sk_j , so that $\text{sk} = \sum_{j=1}^n \text{sk}_j$. A signature (R, s) for a message $m \in \{0, 1\}^*$ is generated as follows. The participants run once again the distributed key generation protocol to produce a commitment $t = rG \in \mathbb{G}$, where r is additively shared between the parties: each party P_j has a share r_j , so that $r = \sum_{j=1}^n r_j$. Then, each party computes a share of the response:

$$s_j = r_j + c \cdot \text{sk}_j, \quad \text{where } c := H(t, m). \quad (6)$$

Let $s := \sum_{j=1}^n s_j$. Then (c, s) is a valid signature on m . In fact:

$$sG = \sum_{j=1}^n r_j G + c \cdot \sum_{j=1}^n \text{sk}_j \cdot G = t + c \cdot \text{pk}, \quad (7)$$

where $c = H(t, m)$.

Concurrent setting insecurity. Gennaro et al. [GJKR07] proved the security of the scheme in a stand-alone *sequential* setting, where no two instances of the protocol can be run in parallel. We remark that if an adversary is allowed to start $\ell \geq \lceil \log p \rceil$ sessions in parallel, the attack against CoSi in Section 6.1.1 can be directly adapted to attack this threshold signature scheme for $n = 2$. The attack of both schemes use the fact that the adversary P_1 (or signer S_1 in CoSi) can see the commitment $t_2 = r_2 G$ of the honest party P_2 (or honest signed S_2) and only then chooses r_1 that defines the commitment $t = r_1 G + t_2$. The generalization to any $n \geq 2$ is straightforward.

Scope of the attack and potential fix. Our attack is an attack against the proposed threshold signature scheme when instantiated with Pedersen’s DKG (and when considered in a concurrent setting, but not an attack against Pedersen’s DKG itself (i.e., JF-DKG from [GJKR07, Fig. 1]). Actually, Gennaro et al. already showed that the Pedersen DKG is not a secure DKG: an adversary can bias the output distribution. They proved the security of their threshold signature scheme in a stand-alone sequential setting directly, without relying on the security of the Pedersen’s DKG.

As for multi-signatures (in Section 6.1.3), adding one initial round of commitment to the commitments $r_j G$ would immediately thwart the attack. Our attack also does not apply when Pedersen’s DKG is replaced by the new DKG protocol from [GJKR07, Fig. 2]. Indeed, the intuition is as follows: this new DKG protocol actually replaces the (Feldman) commitments $r_j G$ by Pedersen commitments⁵ $r_j G + r'_j H$, with H a second generator. Doing so hides the original commitments $r_j G$ and acts similarly as adding one initial round of commitment to the commitments $r_j G$ (without actually requiring one additional round of communication).

⁴ We do not use the fact that only a threshold $t + 1$ of the parties are required to sign in our attack. We assume that all the parties come to sign, to simplify the description of the attack.

⁵ Pedersen commitments are unrelated to Pedersen’s DKG, apart from the fact that both were invented by Pedersen.

6.2.2 Original version of FROST

Komlo and Goldberg FROST [KG20a] proposed an extension of the above threshold signature scheme that was similarly affected by the above concurrent attack. On 19 July 2020, they updated the signing algorithm [KG20b] in a way that is no more susceptible to the above issue: each party now shares (D_j, E_j) and the commitment is computed as $R = \sum_j D_j + h_j E_j$, where $h_j := \mathbb{H}((D_j, E_j, j)_{j \in [t]})$. We direct the reader to [KG20b, Fig. 3] for a more detailed illustration of the problem and the fix.

6.3 Partially blind signatures

Partially blind signatures [AO00] are an extension of blind signature schemes that allow the signer to include some public metadata (e.g., expiration date, collateral conditions, server name, etc.) in the resulting signature. The original construction [AO00], as well as schemes inspired from it, such as Anonymous Credentials Light [BL13] and restrictive partially-blind signatures from bilinear pairings [CZMS06] could be vulnerable to polynomial attacks in specific scenarios where the security of the protocol hinges on the ROS assumption.

6.3.1 Abe–Okamoto

Abe and Okamoto [AO00, Fig. 1] propose a partially blind signature scheme inspired from Schnorr blind signatures. Given a verification key $X := xG$ and some public information `info` that is hashed into the group $Z := \mathbb{H}(\text{info})$, a partially blind signature for the message $m \in \{0, 1\}^*$ is a tuple $(r, c, s, d) \in \mathbb{Z}_p$ where $c + d = \mathbb{H}(rG + cX, sG + dZ, Z, m)$.

Attack. The security of the above partially blind signature is proved up to a poly-logarithmic number of parallel open sessions in the security parameter [AO00]. We show that the security claim is tight by showing that there exists a poly-time attacker against one-more unforgeability in the setting where the adversary can have $\ell = O(\lambda)$ open sessions using the same metadata `info`. The attack follows essentially the same strategy of Section 5.1. First, the attacker opens ℓ parallel sessions and obtains the commitments $(A_i, B_i) \in \mathbb{G}^2$ for $i \in [\ell]$. It then constructs the polynomial ρ_ℓ as per Equation (4). The forged signature for an arbitrary message m^* is computed using the challenge:

$$e_{\ell+1} := \mathbb{H}(\rho(\mathbf{A}) + \rho_0 X, \rho(\mathbf{B}) + \rho_0 Z, Z, m^*) - \rho_0$$

and closing the ℓ sessions as in Section 5.1, i.e., by using the challenges $e_i^{b_i}$ where b_i is the i -th bit of the canonical representation of $e_{\ell+1}$. Given the signatures $(r_i, c_i^{b_i}, s_i, d_i)$ for $i \in [\ell]$, the attacker can finally create its forgery $(\rho(\mathbf{r}), \rho(\mathbf{c}), \rho(\mathbf{s}), \rho(\mathbf{d}))$. The forgery is indeed correct because:

$$\begin{aligned} \rho(\mathbf{c}) + \rho(\mathbf{d}) &= \sum_i \rho_i (c_i^{b_i} + d_i) + 2\rho_0 \\ &= \rho(e_1^{b_1}, \dots, e_\ell^{b_\ell}) + \rho_0 \\ &= \mathbb{H}(\rho(\mathbf{r})G + \rho(\mathbf{c})X, \rho(\mathbf{s})G + \rho(\mathbf{d})Z, Z, m^*). \end{aligned}$$

6.3.2 Anonymous credentials light

Inspired from Abe’s blind signature [Abe01], Baldimtsi and Lysyanskaya [BL13] developed anonymous credentials light (ACL). The security proof of their scheme is under standard assumptions in the sequential settings. The public parameters are a so-called *real public key* $Y = xG$ and a *tag public key* $Z = wG$ (using the paper’s notation). During the signing protocol, the signer produces two shares Z_1, Z_2 of Z such that $Z_1 + Z_2 = Z$, and proves either knowledge of Y (referred to as y -side), or of Z_1, Z_2 (so-called z -side). The discrete log of Z_1, Z_2 is never known by the signer, and the z -branch is inherited by Abe’s blind signature and is necessary for the proof of security.

The essential difference between ACL and Abe’s blind signature is the computation of Z_1 : while in Abe’s scheme it is computed invoking the random oracle over a random string (so that neither the user

nor the signer know its discrete logarithm), in ACL it is computed starting from the user's commitment $C = \sum_{i=1}^n l_i H_i + rH$ (where l_1, \dots, l_n is the list of attributes) and the user might know a discrete-log relation across multiple sessions. This difference is fatal in the concurrent settings.

For simplicity, we assume that no blinding is performed during issuance: μ is removed from the signature and $\gamma := 1$. We stress that γ serves the sole purpose of re-randomizing the tag key for unlinkability, and hence it is not relevant when studying unforgeability. Under the above premises, a valid ACL signature on a message m , for a commitment \tilde{C} on the set of attributes l_1, \dots, l_n is a tuple $(Z, \tilde{C}, r, r'_1, r'_2, c, c') \in \mathbb{G}^2 \times \mathbb{Z}_p^6$ that satisfies:

$$c + c' = \text{H}(Z, \tilde{C}, \underbrace{rG + cY}_A, \underbrace{r'_1 G + c' \tilde{C}}_{A_1}, \underbrace{r'_2 G + c'(Z - \tilde{C})}_{A_2}, m)$$

(Note: $\tilde{C} = C + rnd \cdot G$ is a re-randomization of C ; A is the commitment relative to the y -side; A_1, A_2 are the commitments relative to the z -side.)

Attack. The attacker A opens ℓ parallel sessions, all with the same commitment C , and will provide a one-more forgery for an arbitrary message m^* on the same commitment C .

After opening the ℓ concurrent sessions, the attacker proves in zero-knowledge (as per protocol issuance) that the attributes required are valid, following the *registration phase* as prescribed in the protocol. Let d_1, \dots, d_ℓ denote the randomization key used by the server to re-randomize the commitment C (displayed in [BL13, Fig. 1] as rnd) and sent to the user at the end of the registration phase. Upon receiving $A_i \in \mathbb{G}$ (the commitment of the y -side) and $A'_{1,i}, A'_{2,i}$ (the commitment of the z -side), for $i \in [\ell]$, the attacker computes the polynomial ρ defined in Section 3 (using the commitments and the message of the previous sessions), and computes the commitment forgeries:

$$\begin{aligned} A_{\ell+1} &:= \rho(A_0, \dots, A_\ell) + \rho_0 Y \\ A_{1,\ell+1} &:= \rho(A'_{1,0}, \dots, A'_{1,\ell-1}) + \rho_0 C \\ A_{2,\ell+1} &:= \rho(A'_{2,0}, \dots, A'_{2,\ell-1}) + \rho_0 (Z - C) \end{aligned}$$

A sends the challenges according to the bits of $\text{H}(Z, C, A_{\ell+1}, A_{1,\ell+1}, A_{2,\ell+1}, m^*) - \rho_0$, similarly to Section 5, and receives the responses $(c_i, r_i, c'_i, r'_{1,i}, r'_{2,i}) \in \mathbb{Z}_p^5$, for $i \in [\ell]$. The adversary A computes the forged responses for the y -side:

$$\begin{aligned} c_{\ell+1} &:= \rho(\mathbf{c}) = \sum_{i=1}^{\ell} \rho_i c_i + \rho_0 \\ c'_{\ell+1} &:= \rho(\mathbf{c}') = \sum_{i=1}^{\ell} \rho_i c'_i + \rho_0 \\ r_{\ell+1} &:= \rho(\mathbf{r}) = \sum_{i=1}^{\ell} \rho_i r_i + \rho_0 \\ r'_{1,\ell+1} &:= \rho(\mathbf{r}'_1 + \mathbf{c}' \circ \mathbf{d}) = \sum_{i=1}^{\ell} \rho_i (r'_{1,i} + c'_i d_i) + \rho_0 \\ r'_{2,\ell+1} &:= \rho(\mathbf{r}'_2 - \mathbf{c}' \circ \mathbf{d}) = \sum_{i=1}^{\ell} \rho_i (r'_{2,i} - c'_i d_i) + \rho_0 \end{aligned}$$

In fact, it holds that:

$$r_{\ell+1} G + c_{\ell+1} Y = \sum_{i=1}^{\ell} \rho_i (r_i G + c_i Y) + \rho_0 (Y + G) = A_{\ell+1}$$

$$\begin{aligned}
r'_{1,\ell+1}G + c'_{\ell+1}C &= \sum_{i=1}^{\ell} \rho_i (r'_{1,i}G + c'_i(C + d_iG)) + \rho_0(C + G) = \sum_{i=1}^{\ell} \rho_i A'_{1,i} + \rho_0(C + G) = A_{1,\ell+1} \\
r'_{2,\ell+1}G + c'_{\ell+1}(Z - C) &= \sum_{i=1}^{\ell} \rho_i (r'_{2,i}G + c'_i(Z - C - d_iG)) + \rho_0(Z - C) = A_{2,\ell+1} \\
c_{\ell+1} + c'_{\ell+1} &= \boldsymbol{\rho}(c_0 + c'_0, \dots, c_{\ell} + c'_{\ell}) + \rho_0 = \text{H}(Z, C, A_{\ell+1}, A_{1,\ell+1}, A_{2,\ell+1}, m^*)
\end{aligned}$$

6.4 Brands' signature scheme and U-Prove

Brands [Bra94] designed credential-sharing system where, if the user spends twice the same credential, then anyone can recover their key. The signatures inspired a various anonymous credentials systems such as Microsoft's *U-Prove*⁶ and *credlib*.⁷ In Brands' blind signature scheme, a coin is a pair $(A, B) \in \mathbb{G}$ and a Schnorr blind signature $(X', R, R', s) \in \mathbb{G}^3 \times \mathbb{Z}_p$. Roughly speaking, withdrawal and spending of a Brands coin consists a Schnorr-type protocol where the user does an interactive Schnorr blind signature. The system's security hinges on the security of the security of blind Schnorr signatures (for which we illustrated an attack in Section 5), and hence it presents the same pitfalls we illustrated in the concurrent setting. Unfortunately, this attack extends also to the constructions inspired from it.

Attack. In this paragraph, we focus on the U-Prove cryptographic specification [PZ11, Fig. 8], as an example for our attack. We illustrate how to produce $\ell + 1$ different U-Prove tokens, after ℓ issuance sessions for the same attribute information $l_1, \dots, l_n \in \mathbb{Z}_p$. We stress that the attack is limited to the same attribute information, that is, the same commitment $M = G_0 + \sum_{i=1}^n x_i G_i$ will be used throughout the ℓ sessions. Informally, a valid U-Prove token transcript for some prover information PI is of the form $(H, Z, A, B, c, r) \in \mathbb{G}^4 \times \mathbb{Z}_p^2$, such that:

$$\begin{bmatrix} A \\ B \end{bmatrix} = s \begin{bmatrix} G \\ H \end{bmatrix} - c \begin{bmatrix} G_0 \\ Z \end{bmatrix},$$

where $c = \text{H}(\text{PI}, M, Z, A, B)$; G is the group generator, $G_0 = y_0 G$ is the public key from the server, $H = \alpha M$ is a commitment to list of attributes $l_1, \dots, l_n \in \mathbb{Z}_p$ (blinded by α), and $Z = y_0 H$.

We assume that the adversary A has access to ℓ parallel signing sessions, all for the same set of attributes.⁸ A opens ℓ parallel sessions, obtaining the commitments $A_i, B_i \in \mathbb{G}$ for $i \in [\ell]$. Let M denote the commitment to the attributes, and $Z = y_0 M$ the signature on them as provided by the server during the commitment phase. A samples different α_i for $i \in [\ell]$, and two different prover information values $\text{PI}_i^b \in \{0, 1\}^*$, for $b \in \{0, 1\}$ and $i \in [\ell]$ and computes the blinded group elements $H_i := \alpha_i M$, $Z_i := \alpha_i Z$ and stores internally the challenges $c_i^b := \text{H}(\text{PI}_i^b, M, Z, A_i, B_i)$, without sending them. It computes:

$$\begin{aligned}
A_{\ell+1} &= \sum_i \rho_i A_i + \rho_{\ell}(G - G_0) \\
B_{\ell+1} &= \sum_i \rho_i \alpha_i^{-1} B_i + \rho_{\ell}(M - Z)
\end{aligned}$$

using $\boldsymbol{\rho}$ as defined in Equation (4). We let $H_{\ell} = M, Z_{\ell} = y_0 M$ for simplicity. Z_{ℓ} can be obtained either opening a new session (without completing it), or selecting the blinding factors α_i (for $i \in [\ell + 1]$) such that $\sum_i \alpha_i = 1$. Let PI^* be some arbitrary prover information. A computes $c_{\ell+1} := (b_1, \dots, b_{\ell}) := \text{H}(\text{PI}^*, M, Z, A_{\ell+1}, B_{\ell+1})$ and closes the i -th session (for $i \in [\ell]$) using the challenge c_i^b previously computed. Upon receiving the responses $s_i \in \mathbb{Z}_p$ for $i \in [\ell]$, let $s_{\ell+1} := \boldsymbol{\rho}(s_1, \dots, s_{\ell})$. The $\ell + 1$ forged token transcript is $(A_{\ell+1}, B_{\ell+1}, H_{\ell+1}, Z_{\ell+1}, c_{\ell+1}, r_{\ell+1})$.

⁶ <https://www.microsoft.com/en-us/research/project/u-prove/>

⁷ <http://www.cypherspace.org/credlib/>

⁸ From the specification: *Multiple U-Prove tokens generated using identical common inputs MAY be issued in parallel [and the computation of M, Z] can be shared among all parallel protocol executions.*

6.5 Conditional blind signatures

Conditional blind signatures (CBS), introduced by Grontas et al. [ZGP17], allow a user to request a blind signature on messages of their choice, and the server has a secret boolean input which determines if it will issue a valid signature or not. CBS only allow a *designated* verifier to check the validity of the signature; the user will not be able to distinguish between valid and invalid signatures. Conditional blind signatures have application in e-voting schemes [GPZZ19].

Zacharakis et al. [ZGP17] propose an instantiation of CBS as an extension of Okamoto–Schnorr blind signatures, where the (designated) verifier holds a secret verification key $k \in \mathbb{Z}_p$ and publishes $K = kG$ as public information. During the execution of Okamoto–Schnorr, one of the two responses (s, t) will be computed in \mathbb{G} rather than \mathbb{Z}_p , using K as a generator. Only the designated verifier, who knows the discrete log of K can now check the verification equation.

The attack from Section 5.2 directly applies to their scheme, and leads to a poly-time adversary that with λ queries to the signing oracle for the same bit $b = 1$ can produce one-more forgery with overwhelming probability. This attack does not invalidate the security claims of [ZGP17], which are argued only for a poly-logarithmic number of parallel open sessions.

6.6 Other schemes

The following papers rely on the hardness of the ROS problem for their security proofs, and henceforth may not provide the expected security guarantees: blind anonymous group signatures [CFLW04]; blind identity-based signcryption [YW05]; blind signature schemes from bilinear pairings [CHYC05].

7 Conclusions

Our work provides a polynomial attack against $\text{ROS}_\ell(\lambda)$ when $\ell > \log p$, and a sub-exponential attack for $\ell \leq \log p$. This impacts the one-more unforgeability property of Schnorr and Okamoto–Schnorr blind signatures, plus a number of cryptographic schemes derived from them. Our attacks run in polynomial time only in the concurrent setting, and only for $\ell > \log p$ parallel signing sessions.

In practice, the cost of the attack and the number of sessions required are very small: for today’s security parameters, the attack can be already mounted with $\ell = 9$ parallel open sessions. As already pointed out by [FPS20], even just $\ell = 16$ open sessions could lead to a forgery in time roughly 2^{55} , for a 256-bit prime p . For $\ell = 128$, our attack of Section 4 leads to a forgery in time roughly 2^{32} . For $\ell = 256$, our attack of Section 3 produces a forgery in a matter of seconds on commodity hardware. Although 256 parallel signing sessions might seem at first unrealistic, modern large-scale web servers must handle more than 10 million concurrent sessions.⁹ Given our attack, the main takeaway of our work is that blind Schnorr signatures are unsuitable for wide-scale deployments.

The easiest countermeasure to our attack could be to allow only for sequential signing sessions, as Schnorr blind signatures are unforgeable in the algebraic group model for polynomially many sessions [KLRX20]. Another countermeasure to our attack could be to employ (much) larger security parameters, require the signer to enforce strong ratio limits, and perform frequent key rotations, accepting the tradeoffs given by our attacks. Finally, Fuchsbauer et al. [FPS20] recently introduced a variant of blind Schnorr signatures (the *clause* version) which is unaffected by our attack. We caution that it relies on the conjectured hardness of the so-called *modified ROS problem*, a new assumption which has not been subject to any significant cryptanalysis.

To conclude, other blind signature schemes are to this day considered secure and can be considered as alternatives: blind RSA [Cha82], blind BLS [Bol03], Abe’s blind signature scheme [Abe01, KLRX20], and Tessaro–Zhu signatures [TZ22].

⁹ For further information, read the C10K problem (’99) and the C10M problem (’11).

References

- Abe01. Masayuki Abe. A secure three-move blind signature scheme for polynomially many signatures. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 136–151. Springer, Heidelberg, May 2001.
- AO00. Masayuki Abe and Tatsuaki Okamoto. Provably secure partially blind signatures. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 271–286. Springer, Heidelberg, August 2000.
- BL13. Foteini Baldimtsi and Anna Lysyanskaya. Anonymous credentials light. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 1087–1098. ACM Press, November 2013.
- Bol03. Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 31–46. Springer, Heidelberg, January 2003.
- Bra94. Stefan Brands. Untraceable off-line cash in wallets with observers (extended abstract). In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 302–318. Springer, Heidelberg, August 1994.
- CFLW04. Tony K. Chan, Karyin Fung, Joseph K. Liu, and Victor K. Wei. Blind spontaneous anonymous group signatures for ad hoc groups. In *ESAS*, volume 3313 of *Lecture Notes in Computer Science*, pages 82–94. Springer, 2004.
- Cha82. David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 199–203. Plenum Press, New York, USA, 1982.
- CHYC05. Sherman S. M. Chow, Lucas Chi Kwong Hui, Siu-Ming Yiu, and K. P. Chow. Two improved partially blind signature schemes from bilinear pairings. In Colin Boyd and Juan Manuel González Nieto, editors, *ACISP 05*, volume 3574 of *LNCS*, pages 316–328. Springer, Heidelberg, July 2005.
- CZMS06. Xiaofeng Chen, Fangguo Zhang, Yi Mu, and Willy Susilo. Efficient provably secure restrictive partially blind signatures from bilinear pairings. In Giovanni Di Crescenzo and Avi Rubin, editors, *FC 2006*, volume 4107 of *LNCS*, pages 251–265. Springer, Heidelberg, February / March 2006.
- DEF⁺19. Manu Drijvers, Kasra Edalatnejad, Bryan Ford, Eike Kiltz, Julian Loss, Gregory Neven, and Igor Stepanovs. On the security of two-round multi-signatures. In *2019 IEEE Symposium on Security and Privacy*, pages 1084–1101. IEEE Computer Society Press, May 2019.
- Fel87. Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th FOCS*, pages 427–437. IEEE Computer Society Press, October 1987.
- FPS20. Georg Fuchsbauer, Antoine Plouviez, and Yannick Seurin. Blind schnorr signatures and signed ElGamal encryption in the algebraic group model. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 63–95. Springer, Heidelberg, May 2020.
- GJKR07. Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. *Journal of Cryptology*, 20(1):51–83, January 2007.
- GPZZ19. Panagiotis Grontas, Aris Pagourtzis, Alexandros Zacharakis, and Bingsheng Zhang. Towards everlasting privacy and efficient coercion resistance in remote electronic voting. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala, editors, *FC 2018 Workshops*, volume 10958 of *LNCS*, pages 210–231. Springer, Heidelberg, March 2019.
- HKL19. Eduard Hauck, Eike Kiltz, and Julian Loss. A modular treatment of blind signatures from identification schemes. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 345–375. Springer, Heidelberg, May 2019.
- HKLN20. Eduard Hauck, Eike Kiltz, Julian Loss, and Ngoc Khanh Nguyen. Lattice-based blind signatures, revisited. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 500–529. Springer, Heidelberg, August 2020.
- KG20a. Chelsea Komlo and Ian Goldberg. FROST: Flexible round-optimized Schnorr threshold signatures, 2020. <https://crysp.uwaterloo.ca/software/frost/frost-extabs.pdf>; version from "January 7, 2020"; accessed 2020-10-04.
- KG20b. Chelsea Komlo and Ian Goldberg. FROST: Flexible round-optimized Schnorr threshold signatures. Cryptology ePrint Archive, Report 2020/852, 2020. <https://eprint.iacr.org/2020/852>.
- KLRX20. Julia Kaster, Julian Loss, Michael Rosenberg, and Jiayu Xu. On pairing-free blind signature schemes in the algebraic group model. Cryptology ePrint Archive, Report 2020/1071, 2020.
- MPSW18a. Gregory Maxwell, Andrew Poelstra, Yannick Seurin, and Pieter Wuille. Simple Schnorr multi-signature with applications to Bitcoin. Cryptology ePrint Archive, Report 2018/068, Revision 20180118:124757, 2018. <https://eprint.iacr.org/2018/068/20180118:124757>.
- MPSW18b. Gregory Maxwell, Andrew Poelstra, Yannick Seurin, and Pieter Wuille. Simple Schnorr multi-signature with applications to Bitcoin. Cryptology ePrint Archive, Report 2018/068, Revision 20180520:191909, 2018. <https://eprint.iacr.org/2018/068/20180520:191909>.

- MS09. Lorenz Minder and Alistair Sinclair. The extended k-tree algorithm. In Claire Mathieu, editor, *20th SODA*, pages 586–595. ACM-SIAM, January 2009.
- PS00. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
- PZ11. Christian Paquin and Greg Zaverucha. U-prove cryptographic specification v1. 1. *Technical Report, Microsoft Corporation*, 2011.
- S⁺20. W. A. Stein et al. *Sage Mathematics Software (Version 9.1)*. The Sage Development Team, 2020. <http://www.sagemath.org>.
- Sch01. Claus-Peter Schnorr. Security of blind discrete log signatures against interactive attacks. In Sihan Qing, Tatsuaki Okamoto, and Jianying Zhou, editors, *ICICS 01*, volume 2229 of *LNCS*, pages 1–12. Springer, Heidelberg, November 2001.
- STV⁺16. Ewa Syta, Iulia Tamas, Dylan Visher, David Isaac Wolinsky, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ismail Khoffi, and Bryan Ford. Keeping authorities “honest or bust” with decentralized witness cosigning. In *2016 IEEE Symposium on Security and Privacy*, pages 526–545. IEEE Computer Society Press, May 2016.
- TZ22. Stefano Tessaro and Chenzhi Zhu. Short pairing-free blind signatures with exponential security. Cryptology ePrint Archive, Report 2022/047, 2022. <https://ia.cr/2022/047>.
- Wag02. David Wagner. A generalized birthday problem. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 288–303. Springer, Heidelberg, August 2002.
- YW05. Tsz Hon Yuen and Victor K. Wei. Fast and proven secure blind identity-based signcryption from pairings. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 305–322. Springer, Heidelberg, February 2005.
- ZGP17. Alexandros Zacharakis, Panagiotis Grontas, and Aris Pagourtzis. Conditional blind signatures. Cryptology ePrint Archive, Report 2017/682, 2017. <http://eprint.iacr.org/2017/682>.

A Code listing for Schnorr's blind signature forgery

```
1 # public parameters: secp256k1
2 Zq = GF(0xfffffffffffffffffffffffffffffffffffffffffffffffffffffffc2f)
3 E = EllipticCurve(Zq, [0, 7])
4 G = E.lift_x(0x79be667ef9dcbbac55a06295ce870b07029bfcd2dce28d959f2815b16f81798)
5 p = G.order()
6 Zp = GF(p)
7
8 def random_oracle(hinput, _table=dict()):
9     if hinput not in _table:
10         _table[hinput] = Zp.random_element()
11     return _table[hinput]
12
13 def verify(message, K, e, s):
14     assert random_oracle((K, message)) == e, "random oracle fails"
15     assert G * int(s) + X * int(e) == K, "verification equation fails"
16     return True
17
18 def inner_product(coefficients, values):
19     return sum(y*int(x) for x, y in zip(coefficients, values))
20
21 # server: generate public key
22 x = Zp.random_element()
23 X = G * int(x)
24
25 # adversary: open 'ell' sessions
26 ell = 256
27
28 # server: generate commitments
29 k = [Zp.random_element() for i in range(ell)]
30 K = [G * int(k_i) for k_i in k]
31
32 # adversary: generate challenges
33 e = [[random_oracle((K_i, b)) for b in range(2)] for K_i in K]
34 P = ([-sum([Zp(2)^i * e[i][0]/(e[i][1] - e[i][0]) for i in range(ell)])] +
35      [Zp(2)^i / (e[i][1] - e[i][0]) for i in range(ell)])
36
37 forged_K = inner_product(P, [G+X] + K)
38 forged_message = "message"
39 forged_e = random_oracle((forged_K, forged_message))
40 bits = [int(b) for b in bin(forged_e)[2:].rjust(256, '0')][::-1]
41 chosen_e = [e[i][b] for (i, b) in enumerate(bits)]
42
43 # server: generate the responses
44 s = [k[i] - chosen_e[i]*x for i in range(ell)]
45
46 # attacker: generate the forged response
47 forged_s = inner_product(P, [1] + s)
48
49 ## check all previous signatures were valid
50 print(all(
51     # 1 signatures generated honestly
52     [verify(m_i, K_i, e_i, s_i) for (m_i, K_i, e_i, s_i) in zip(bits, K, chosen_e, s)] +
53     # final signature
54     [verify(forged_message, forged_K, forged_e, forged_s)]
55 ))
```