



On Lattice-Based Interactive Protocols: An Approach with Less or No Aborts

Nabil Alkeilani Alkadri^{1(✉)}, Rachid El Bansarkhani²,
and Johannes Buchmann¹

¹ Technische Universität Darmstadt, Darmstadt, Germany
nabil.alkadri@tu-darmstadt.de, buchmann@cdc.informatik.tu-darmstadt.de
² QuantiCor Security GmbH, Darmstadt, Germany
rachid.elbansarkhani@quanticor-security.de

Abstract. A canonical identification (CID) scheme is a 3-move protocol consisting of a commitment, challenge, and response. It constitutes the core design of many cryptographic constructions such as zero-knowledge proof systems and various types of signature schemes. Unlike number-theoretic constructions, CID in the lattice setting usually forces provers to abort and repeat the whole authentication process once the distribution of the computed response does not follow a target distribution independent from the secret key. This concept has been realized by means of rejection sampling, which makes sure that the secrets involved in a protocol are concealed after a certain number of repetitions. This however has a negative impact on the efficiency of interactive protocols because it leads to a number of communication rounds that is multiplicative in the number of aborting participants (or rejection sampling procedures). In this work we show how the CID scheme underlying many lattice-based protocols can be designed with smaller number of aborts or even without aborts. Our new technique exploits (unbalanced) binary hash trees and thus significantly reduces the communication complexity. We show how to apply this new method within interactive zero-knowledge proofs. We also present BLAZE⁺: a further application of our technique to the recently proposed lattice-based blind signature scheme BLAZE (FC'20). We show that BLAZE⁺ has an improved performance and communication complexity compared to BLAZE while preserving the size of keys and signatures.

Keywords: Lattice-based cryptography · Aborts · Hash trees

1 Introduction

A canonical identification (CID) scheme allows a prover \mathcal{P} to prove to a verifier \mathcal{V} the possession of a secret key s in the following way: \mathcal{P} sends a commitment to \mathcal{V} , who then sends a challenge c back to \mathcal{P} . Upon receiving c , \mathcal{P} answers with a response z . This response allows \mathcal{V} to verify \mathcal{P} 's authenticity while not leaking any information about the secret key. In number-theoretic constructions like

Schnorr’s CID scheme [34], the response z already hides s , since it is computed by adding a secret masking term y to the term sc , i.e., $z = y + sc$. The term y is chosen uniformly at random from a large distribution and is also used to compute the commitment. This approach has been generalized in [22] to include aborting provers for the lattice setting. In a lattice-based CID scheme y is required to be chosen from a narrow distribution (typically, Gaussian or uniform) and the so called rejection sampling procedure [29] is used to hide the distribution of sc . If the sum z is not accepted, a new masking term is sampled. This procedure is repeated until the sum becomes independently distributed from the secret term sc . Lattice-based CID is a fundamental building block of many cryptographic constructions including zero-knowledge protocols (e.g., [6, 9]) as well as signature schemes (e.g., [5, 12, 23]) and even those with advanced functionalities such as ring signatures (e.g., [7, 35]), blind signatures [4, 33], and multisignatures [14].

While aborting does not affect the efficiency of constructions with one rejection sampling process like ordinary signatures, it has a significant negative impact on the performance and communication complexity of lattice-based interactive protocols with multiple rejection sampling procedures. For instance, the multisignature scheme proposed in [14] entails a repetition rate that grows exponentially in the number of users participating in the signing protocol. Though it is efficient for a small set of users, one would need to restart the protocol very often when instantiated with a large set because each user has to carry out rejection sampling. Another example is the blind signature scheme BLAZE [4] and its predecessor introduced in [33]. In both constructions not only signers have to carry out rejection sampling and repeat the signing process M_S times until the secret key is concealed, but for maintaining blindness even users have to apply rejection sampling M_U times and request a protocol restart in case of failure. This imposes a multiplicative repetition rate $M_S \cdot M_U$ and an additional communication step due to the possibility of failures causing protocol restarts. In this case, a proof of failure is sent to the signer, i.e., a proof that allows the signer to verify the occurrence of a failure. Although BLAZE has been shown to be practical [4], this additional step increases the time and communication complexity required to generate valid signatures and forces the use of statistically hiding and computationally binding commitments to retain security.

Therefore, masking secrets in lattice-based interactive protocols with multiple rejection sampling procedures such that aborting occurs as little as possible while maintaining efficiency and security remained a very important research question. This would improve the running time and decrease the total amount of communication required to successfully complete the protocol.

Contributions. In this work we show how to reduce the number of repetitions in lattice-based protocols by means of a tool that we call *trees of commitments*. A tree of commitments is an (unbalanced) binary hash tree of height $h \geq 1$, whose leaves are the hash values of $\ell > 1$ commitments computed from masking terms sampled during an instance of a CID-based protocol. The number ℓ is chosen such that rejection sampling succeeds for at least one masking term at a

Table 1. Comparing BLAZE⁺ (this work) with BLAZE [4] at approximately 128 bits of security. The parameter δ_{abort} denotes the aborting probability by the user, and ℓ denotes the related number of masking terms. Performance is given in cycles and milliseconds (in parentheses), sizes and communication complexity in kilobytes. The corresponding parameters can be found in Table 3. Benchmarking the parameters were carried out on an Intel Core i7-6500U, operating at 2.3 GHz and 8 GB of RAM.

Scheme	δ_{abort}	ℓ	Complexity	BS.KGen	BS.Sign	BS.Verify	sk	pk	Signature
BLAZE ⁺	2^{-128}	71	177.8	222, 151 (0.11)	112, 540, 972 (56.49)	348, 724 (0.18)	0.75	3.9	6.7
BLAZE ⁺	2^{-40}	32	189.1	222, 151 (0.11)	56, 193, 762 (28.21)	348, 724 (0.18)	0.75	3.9	6.7
BLAZE ⁺	2^{-10}	8	189.2	222, 151 (0.11)	24, 443, 555 (12.27)	348, 724 (0.18)	0.75	3.9	6.6
BLAZE	0.38	1	351.6	204, 671 (0.10)	35, 547, 397 (17.85)	276, 210 (0.14)	0.8	3.9	6.6

given probability bound. This allows to aggregate ℓ commitments in one tree and send only the root of the tree as a new commitment rather than ℓ commitments. The new response now further includes the authentication path of the leaf with index k ($0 \leq k < \ell$), where at step k rejection sampling accepts for the first time after $k - 1$ trials. Note that by choosing ℓ large enough we can remove aborts completely. Interestingly, only trees with small heights are required to reduce aborts to very small probabilities, e.g., $h = 3$ for a probability of at most 2^{-10} .

We demonstrate the effectiveness of using our method in interactive zero-knowledge proofs and blind signature schemes. More concretely, we show how to reduce the communication complexity of interactive zero-knowledge protocols by using trees of commitments in a lattice-based zero-knowledge proof of knowledge. Furthermore, we utilize trees of commitments in the blind signature scheme BLAZE [4]. We call the new scheme BLAZE⁺. In the new scheme a user constructs a tree of commitments using ℓ masking terms such that blindness is ensured at a given probability bound. More precisely, given a security level of λ bits we fix an aborting probability δ_{abort} and compute ℓ such that signatures are blind with probability of at least $1 - \delta_{\text{abort}}$. For approximately 128 bits of security, our results (summarized in Table 1) show that while preserving the size of keys and signatures, the communication complexity is significantly decreased and the signing speed is improved for $\delta_{\text{abort}} = 2^{-10}$. Note that choosing $\delta_{\text{abort}} = 2^{-128}$ implies blindness with overwhelming probability. In this case (i.e., when $\delta_{\text{abort}} = 2^{-\lambda}$) we can safely remove the last step of the protocol, hence proof of failures and the use of commitment schemes. Thus, we obtain a 3-move version of the protocol similar to the basic structure of CID. We present this version in Sect. 4 and the 4-move scheme in the full version of this paper [3], where aborts at the user side occur with probability of choice. We leave applying trees of commitments to multisignatures [14] as a future work.

Finally, we note that the impossibility results of 3-move blind signature schemes due to [16] do not apply to our 3-move version of BLAZE⁺. These results show that finding black-box reductions from successful forgers to some

non-interactive cryptographic assumption is infeasible in the standard model (i.e., without random oracles) for statistically blind schemes with 3 (or less) moves such that one can verify that an honest user was able to obtain a valid signature from the interaction with the (malicious) signer. In our 3-move protocol, there is no way to check that the user has obtained a valid signature, since he does not reveal the secret information that are involved in generating the signature and are required to check its validity. Furthermore, BLAZE⁺ is proven secure in the random oracle model rather than the standard model.

Techniques. We show how to reduce the number of repetitions or even remove aborts in CID-based protocols, completely. To this end, we give a brief description of the CID scheme that underlies many lattice-based constructions and was originally introduced in [22]. Let \mathbf{A} be a public matrix selected uniformly at random from $\mathbb{Z}_q^{n \times m}$. The prover \mathcal{P} would like to prove to a verifier \mathcal{V} the possession of a secret matrix $\mathbf{S} \in \mathbb{Z}^{m \times n}$ with small entries such that $\mathbf{B} = \mathbf{AS} \pmod{q}$. We let χ denote some distribution over \mathbb{Z} . Typically, χ is either the discrete Gaussian distribution over \mathbb{Z} or the uniform distribution over a small subset of \mathbb{Z} . The challenge space is defined by $\mathcal{C} = \{\mathbf{c} = (c_1, \dots, c_n) \in \mathbb{Z}^n : c_i \in \{-1, 0, 1\}, \sum_1^n |c_i| = \kappa\}$. We let RejSamp denote an algorithm that carries out rejection sampling. The commitment is a vector $\mathbf{v} = \mathbf{Ay} \pmod{q}$, where \mathbf{y} is a masking vector chosen from χ^m . For a challenge $\mathbf{c} \in \mathcal{C}$ the response is given by $\mathbf{z} = \mathbf{y} + \mathbf{Sc}$. The verifier accepts if and only if $\mathbf{v} = \mathbf{Az} - \mathbf{Bc} \pmod{q}$ and $\|\mathbf{z}\|_p \leq B$, where B is a predefined bound and $p \in \{2, \infty\}$ depending on the distribution χ . Aborting occurs if $\text{RejSamp}(\mathbf{z})$ does not accept. The protocol is always repeated by sampling a fresh \mathbf{y} until RejSamp accepts such that \mathbf{z} is statistically independent from \mathbf{Sc} .

Consider a lattice-based interactive protocol with $N \geq 1$ rejection sampling procedures, where each of them is repeated $x \geq 1$ times on average. The main motivation of this work is the observation that the total average number of repetitions M in such a protocol is multiplicative in N , i.e., $M = x^N$. Thus, the main question is: Can we improve it?

One can use a large enough distribution χ such that RejSamp accepts after a fixed number of repetitions M , e.g., $M \leq 2$. This is already established in previous works as a trade-off between performance and sizes (see, e.g., [5, 12, 23]), but it does not solve the problem for all interactive protocols as explained above.

Our first attempt is the following. Rather than sampling one masking term \mathbf{y} and repeating this process until RejSamp accepts, \mathcal{P} generates $\ell > 1$ masking vectors \mathbf{y}_j at once and computes the commitment $(\mathbf{v}_0, \dots, \mathbf{v}_{\ell-1})$, where $\mathbf{v}_j = \mathbf{Ay}_j \pmod{q}$ and $j = 0, \dots, \ell - 1$. The response is then \mathbf{z}_k , where k ($0 \leq k < \ell$) is the first index for which RejSamp accepts. This reduces aborts, but the amount of exchanged data grows in ℓ . In particular, any type of lattice-based signature following this approach becomes very large. While this can be decreased by using some cryptographic hash function F and sending $F(\mathbf{v}_j)$ instead of \mathbf{v}_j , this is still not satisfactory. An approach with some similarities has been taken in [30] in a different context for zero-knowledge proofs, where all the hash values of

commitments of potential masking terms are sent. We note that no tree structure for commitments has been applied in [30] and furthermore the challenge size increases linearly in the number of masking terms, which is not the case in our attempt. The protocol is then repeated multiple times to achieve negligible soundness error. Thus, such an approach is still inefficient.

Our final solution to this issue is to use a *tree of commitments*: an (unbalanced) binary hash tree of height $h = \lceil \log(\ell) \rceil$, whose leaves are $F(\mathbf{v}_j)$. The commitment is simply the root of the tree \mathbf{root} , and the response is the pair $(\mathbf{z}_k, \mathbf{auth})$, where \mathbf{auth} is the authentication path of the leaf with index k . Verification is carried out by checking that $\|\mathbf{z}_k\|_p \leq B$ and \mathbf{root} is equal to the root of the tree associated to the leaf $F(\mathbf{A}\mathbf{z}_k - \mathbf{B}\mathbf{c} \pmod{q})$ and its given authentication path \mathbf{auth} . Using a tree of commitments obviously reduces the communication complexity. It can also improve the performance of interactive protocols with multiple rejection sampling procedures as we demonstrate in this work. We note that the number of masking terms can be chosen such that the aborting probability is bounded by some given bound. In Sect. 3.3 we show how to optimize this number. We note that our technique may be used in [30] to improve efficiency.

Finally, we briefly explain two further optimizations that can be exploited when using trees of commitments. The first one is to generate trees with randomized hashing similar to the standard of the hash-based signature scheme XMSS [18]. This allows to save space and further reduce the communication complexity, since randomized hashing requires the hash function F to be only second preimage resistant rather than collision resistant. This means the output of F is required to be $\geq \lambda$ rather than $\geq 2\lambda$ bits assuming λ bits of classical security. The second optimization allows to reuse already generated, but not consumed, masking terms in subsequent executions of the protocol. This further improves the performance of the protocol, since complete subtrees of the tree can be reused. This reduces the number of new masking terms to be sampled in addition to the number of multiplications and hash computations.

Related Work. In the context of analyzing the hardness of computational lattice problems, previous works such as [10, 11, 17] point to techniques called “noise swallowing” or “super-polynomial noise flooding”, which use Gaussian masking terms entailing a super-polynomial Gaussian parameter in order to swallow a polynomially large secret term. However, the negative impact on the efficiency is tremendous as the parameters become also super-polynomial. By generating many masking terms at once and capturing them in a tree of commitments, the secret and masking terms remain polynomially bounded while the number of repetitions is reduced. As mentioned above, the approach of sending hashed commitments has been used in [30] for zero-knowledge proofs of small secrets, but without the use of tree structures for commitments and the other efficiency improvements. However, sending commitments in a tree structure has been suggested, e.g., in [20] to reduce the communication complexity of proof systems, but not repetitions of lattice-based protocols. Our work exploits hash trees in the context of lattice-based interactive protocols with aborting participants.

Outline. In Sect. 2 we review the relevant background. In Sect. 3 we define trees of commitments and show how they can be utilized in lattice-based canonical identification schemes and hence, in interactive zero-knowledge protocols. In Sect. 4 we demonstrate the practical relevance of our new technique by introducing a new blind signature scheme that we call BLAZE⁺.

2 Preliminaries

We let $\mathbb{N}, \mathbb{Z}, \mathbb{R}$ denote the set of natural numbers, integers, and real numbers, respectively. We denote column vectors with bold lower-case letters and matrices with bold upper-case letters. The identity matrix of dimension n is denoted by \mathbf{I}_n . For any positive integer q we write \mathbb{Z}_q to denote the set of integers in the range $[-\frac{q}{2}, \frac{q}{2}] \cap \mathbb{Z}$. The Euclidean norm (ℓ_2 -norm) of a vector \mathbf{v} with entries v_i is defined as $\|\mathbf{v}\| = (\sum_i |v_i|^2)^{1/2}$, and its ℓ_∞ -norm as $\|\mathbf{v}\|_\infty = \max_i |v_i|$. We define the ring $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ and its quotient $R_q = R/qR$, where n is a power of 2. We assume that R is an integral domain. A ring element $a_0 + a_1x + \dots + a_{n-1}x^{n-1} \in R_q$ is denoted by \hat{a} and it corresponds to a vector $\mathbf{a} \in \mathbb{Z}_q^n$ via coefficient embedding, hence $\|\hat{a}\| = \|\mathbf{a}\|$ and $\|\hat{a}\|_\infty = \|\mathbf{a}\|_\infty$. We write $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_k) \in R_q^k$ to denote a vector of ring elements and $\hat{\mathbf{A}}$ for a matrix with entries from R_q . The norms of $\hat{\mathbf{a}}$ are defined by $\|\hat{\mathbf{a}}\| = (\sum_{i=1}^k \|\hat{a}_i\|^2)^{1/2}$ and $\|\hat{\mathbf{a}}\|_\infty = \max_i \|\hat{a}_i\|_\infty$. We let \mathbb{T}_κ^n denote the set of all $(n-1)$ -degree polynomials with coefficients from $\{-1, 0, 1\}$ and Hamming weight κ . All logarithms in this work are to base 2, i.e., $\log(\cdot) = \log_2(\cdot)$. We always denote the security parameter by $\lambda \in \mathbb{N}$. A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is called *negligible* if there exists an $n_0 \in \mathbb{N}$ such that for all $n > n_0$, it holds $f(n) < \frac{1}{p(n)}$ for any polynomial p . With $\text{negl}(\lambda)$ we denote a negligible function in λ . A probability is called *overwhelming* if it is at least $1 - \text{negl}(\lambda)$. The *statistical distance* between two distributions X, Y over a countable domain D is defined by $\Delta(X, Y) = \frac{1}{2} \sum_{n \in D} |X(n) - Y(n)|$. The distributions X, Y are called *statistically close* if $\Delta(X, Y) = \text{negl}(\lambda)$. We write $x \leftarrow D$ to denote that x is sampled according to a distribution D . We let $x \leftarrow_{\S} S$ denote choosing x uniformly random from a finite set S .

2.1 Cryptographic Primitives

A canonical identification (CID) scheme is a 3-move interactive protocol of the following form: A prover \mathcal{P} initiates the protocol by sending a commitment message y to a verifier \mathcal{V} . Upon receiving y , \mathcal{V} sends a uniform random challenge c to \mathcal{P} . Afterwards, a response z is sent from \mathcal{P} back to \mathcal{V} , which then allows \mathcal{V} to make a deterministic decision about \mathcal{P} 's authenticity. The tuple (y, c, z) represents a protocol transcript. A formal definition follows.

Definition 1 (Canonical Identification Scheme). *A canonical identification scheme with commitment space \mathcal{Y} , challenge space \mathcal{C} , and response space \mathcal{Z} is defined as a tuple of the following polynomial-time algorithms:*

- $\text{KG}(1^\ell)$ is a key generation algorithm that outputs a pair of keys (pk, sk) from some key space \mathcal{K} , where pk is a public key and sk is a secret key.
- $P = (P_1(\text{sk}), P_2(\text{sk}, y, c, \text{st}))$ is a prover algorithm consisting of two algorithms: P_1 takes as input a secret key sk and returns a commitment $y \in \mathcal{Y}$ and a state st , whereas P_2 on input sk, y , a challenge $c \in \mathcal{C}$, and st , outputs a response $z \in \mathcal{Z} \cap \{\perp\}$, where the symbol $\perp \notin \mathcal{Z}$ indicates failure.
- $V(\text{pk}, y, c, z)$ is a verification algorithm that takes as input a public key pk and a transcript (y, c, z) , and outputs 1 if it is valid and 0 otherwise.

The standard security notion of CID schemes is impersonation under the active or passive attack model. In the active attack model, any adversary \mathcal{A} interacting with \mathcal{P} must not be able to extract any useful information. Passive attacks correspond to eavesdropping, i.e., \mathcal{A} is in possession of transcripts generated by interactions between the real prover and verifier. According to [1], impersonation under passive attacks is stronger than the active attack model.

Definition 2 (Blind Signature Scheme). A blind signature scheme BS is a tuple of polynomial-time algorithms $\text{BS} = (\text{BS.KGen}, \text{BS.Sign}, \text{BS.Verify})$ such that:

- $\text{BS.KGen}(1^\lambda)$ is a key generation algorithm that outputs a pair of keys (pk, sk) , where pk is a public key and sk is a secret key.
- $\text{BS.Sign}(\text{sk}, \text{pk}, \mu)$ is an interactive protocol between a signer \mathcal{S} and a user \mathcal{U} . The input of \mathcal{S} is a secret key sk , whereas the input of \mathcal{U} is a public key pk and a message $\mu \in \mathcal{M}$, where \mathcal{M} is the message space. The output of \mathcal{S} is a view \mathcal{V} (interpreted as a random variable) and the output of \mathcal{U} is a signature σ , i.e., $(\mathcal{V}, \sigma) \leftarrow \langle \mathcal{S}(\text{sk}), \mathcal{U}(\text{pk}, \mu) \rangle$. We write $\sigma = \perp$ to denote failure.
- $\text{BS.Verify}(\text{pk}, \mu, \sigma)$ is a verification algorithm that outputs 1 if the signature σ is valid and 0 otherwise.

Security of blind signatures is captured by two security notions: blindness and one-more unforgeability [19, 31]. The former prevents a malicious signer to learn information about user’s messages (see [4] for a formal definition). The latter ensures that each completed execution of BS.Sign yields at most one signature.

Definition 3 (One-More Unforgeability). Let \mathcal{H} be a family of random oracles. A blind signature scheme BS is called $(t, q_{\text{Sign}}, q_{\text{H}}, \varepsilon)$ -one-more unforgeable in the random oracle model if for any adversarial user \mathcal{U}^* running in time at most t and making at most q_{Sign} signing and q_{H} hash queries, the game $\text{Forge}_{\text{BS}, \mathcal{U}^*}(\lambda)$ depicted in Fig. 1 outputs 1 with probability $\Pr[\text{Forge}_{\text{BS}, \mathcal{U}^*}(\lambda) = 1] \leq \varepsilon$. The

Game $\text{Forge}_{\text{BS}, \mathcal{U}^*}(\lambda)$

- 1: $(\text{pk}, \text{sk}) \leftarrow \text{BS.KGen}(1^\lambda)$
- 2: $\text{H} \leftarrow \mathcal{H}(1^\lambda)$
- 3: $((\mu_1, \sigma_1), \dots, (\mu_l, \sigma_l)) \leftarrow \mathcal{U}^{*\text{H}(\cdot), \langle \mathcal{S}(\text{sk}), \cdot \rangle^\infty}(\text{pk})$
- 4: $k :=$ number of successful signing invocations
- 5: **if** $(\mu_i \neq \mu_j \text{ for all } 1 \leq i < j \leq l \wedge \text{BS.Verify}(\text{pk}, \mu_i, \sigma_i) = 1, \forall i \in [l] \wedge k + 1 = l)$ **then**
- 6: **return** 1
- 7: **return** 0

Fig. 1. The security game of one-more unforgeability of blind signatures.

scheme is strongly $(t, q_{\text{Sign}}, q_H, \varepsilon)$ -one-more unforgeable if the condition $\mu_i \neq \mu_j$ in the game changes to $(\mu_i, \sigma_i) \neq (\mu_j, \sigma_j)$ for all $1 \leq i < j \leq l$.

2.2 Lattices and Gaussians

Let $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\} \in \mathbb{R}^{m \times k}$ be a set of linearly independent vectors for $k \leq m$. The m -dimensional lattice \mathcal{L} of rank k generated by \mathbf{B} is given by $\mathcal{L}(\mathbf{B}) = \{\mathbf{B}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}^k\} \subset \mathbb{R}^m$. The discrete Gaussian distribution $D_{\mathcal{L}, \sigma, \mathbf{c}}$ over a lattice \mathcal{L} with standard deviation $\sigma > 0$ and center $\mathbf{c} \in \mathbb{R}^n$ is defined as follows: For every $\mathbf{x} \in \mathcal{L}$ the probability of \mathbf{x} is $D_{\mathcal{L}, \sigma, \mathbf{c}}(\mathbf{x}) = \rho_{\sigma, \mathbf{c}}(\mathbf{x}) / \rho_{\sigma, \mathbf{c}}(\mathcal{L})$, where $\rho_{\sigma, \mathbf{c}}(\mathbf{x}) = \exp(-\frac{\|\mathbf{x} - \mathbf{c}\|^2}{2\sigma^2})$ and $\rho_{\sigma, \mathbf{c}}(\mathcal{L}) = \sum_{\mathbf{x} \in \mathcal{L}} \rho_{\sigma, \mathbf{c}}(\mathbf{x})$. The subscript \mathbf{c} is taken to be $\mathbf{0}$ when omitted. We recall a lemma that gives a tail bound on discrete Gaussians and a rejection sampling lemma.

Lemma 1 ([23, Lemma 4.4]). *For any $t, \eta > 0$ we have*

1. $\Pr_{x \leftarrow D_{\mathbb{Z}, \sigma}}[|x| > t \cdot \sigma] \leq 2 \exp(-t^2/2)$.
2. $\Pr_{\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \sigma}}[\|\mathbf{x}\| > \eta \sigma \sqrt{m}] \leq \eta^m \exp(\frac{m}{2}(1 - \eta^2))$.

Lemma 2 ([23, Theorem 4.6, Lemma 4.7]). *Let $V \subseteq \mathbb{Z}^m$ with elements having norms bounded by T , $\sigma = \omega(T\sqrt{\log m})$, and $h : V \rightarrow \mathbb{R}$ be a probability distribution. Then there exists a constant $M = O(1)$ such that $\forall \mathbf{v} \in V$: $\Pr[D_{\mathbb{Z}^m, \sigma}(\mathbf{z}) \leq M \cdot D_{\mathbb{Z}^m, \sigma, \mathbf{v}}(\mathbf{z}); \mathbf{z} \leftarrow D_{\mathbb{Z}^m, \sigma}] \geq 1 - \varepsilon$, where $\varepsilon = 2^{-\omega(\log m)}$. Furthermore, the following two algorithms are within statistical distance $\delta = \varepsilon/M$.*

1. $\mathbf{v} \leftarrow h, \mathbf{z} \leftarrow D_{\mathbb{Z}^m, \sigma, \mathbf{v}}$, output (\mathbf{z}, \mathbf{v}) with probability $\frac{D_{\mathbb{Z}^m, \sigma}(\mathbf{z})}{M \cdot D_{\mathbb{Z}^m, \sigma, \mathbf{v}}(\mathbf{z})}$.
2. $\mathbf{v} \leftarrow h, \mathbf{z} \leftarrow D_{\mathbb{Z}^m, \sigma}$, output (\mathbf{z}, \mathbf{v}) with probability $1/M$.

Moreover, the probability that the first algorithm outputs something is at least $(1 - \varepsilon)/M$. If $\sigma = \alpha T$ for any positive α , then $M = \exp(\frac{12}{\alpha} + \frac{1}{2\alpha^2})$ with $\varepsilon = 2^{-100}$.

We let $\text{RejSamp}(x)$ denote an algorithm that carries out rejection sampling on input x . The algorithm outputs 1 if it accepts and 0 otherwise. Next, we define the lattice problems related to this work.

Definition 4 (Module Short Integer Solution (MSIS) Problem). *Let n, q, k_1, k_2 be positive integers and β a positive real. Given a uniform random matrix $\hat{\mathbf{A}} \in R_q^{k_1 \times k_2}$, the Hermite Normal Form of MSIS asks to find a non-zero vector $\hat{\mathbf{x}} \in R^{k_1 + k_2}$ such that $[\mathbf{I}_{k_1} \ \hat{\mathbf{A}}] \cdot \hat{\mathbf{x}} = \mathbf{0} \pmod{q}$, where $\|\hat{\mathbf{x}}\| \leq \beta$.*

Definition 5 (Module Learning With Errors (MLWE) Problem). *Let n, q, k_1, k_2 be positive integers and $\hat{\mathbf{A}}$ be a matrix chosen uniformly at random from $R_q^{k_1 \times k_2}$. Given $(\hat{\mathbf{A}}, \hat{\mathbf{b}})$, the decision MLWE problem asks to distinguish (with non-negligible advantage) whether $\hat{\mathbf{b}}$ were chosen from the uniform distribution over $R_q^{k_1}$ or from the following distribution: Given $\hat{\mathbf{s}} \leftarrow \chi^{k_2}$ and $\hat{\mathbf{e}} \leftarrow \chi^{k_1}$, output the vector $\hat{\mathbf{A}}\hat{\mathbf{s}} + \hat{\mathbf{e}} \pmod{q}$, where χ is an error distribution (typically, either $D_{\mathbb{Z}^n, \sigma}$ or the uniform distribution over a small subset of R_q).*

The MLWE problem [21] generalizes LWE [32] and RLWE [24]. More precisely, by setting $k_1 = 1$ in the definition above we obtain the ring version RLWE, while setting $k_1 > 1$ and $R_q = \mathbb{Z}_q$ yields a definition of the LWE problem. The same applies for MSIS [21] and its special versions SIS [2] and RSIS [27].

3 How to Reduce Aborts in Lattice-Based Protocols

In this section we show how aborting in lattice-based protocols can be reduced or even be removed at all. As stated in Sect. 1, when the number of rejection sampling procedures N in an interactive CID-based protocol grows, the total number of repetitions becomes multiplicative in N , e.g., [4, 14, 33], and a large amount of communication is required to successfully complete the protocol. Consider the CID protocol sketched in Sect. 1. If rejection sampling fails, a new masking term is sampled, hence a new commitment has to be computed and sent in order to receive a new challenge c . Suppose that c does not change for certain number of masking terms and related commitments, which are sent in an aggregated form while any successfully computed response can be verified and related to the corresponding commitment. In this case repetition does not have to occur often or even not at all. We realize this concept by means of *tree of commitments*: a method by which different commitments belong to one challenge in an aggregated form and only the valid response and its related commitment will be revealed. Masking terms that are rejected or not consumed during rejection sampling remain hidden and will never be revealed.

3.1 Trees of Commitments

In this section we describe trees of commitments. We first define relevant functions and algorithms. For a positive integer $\omega \geq 2\lambda$, we let $F : \{0, 1\}^* \rightarrow \{0, 1\}^\omega$ be a collision resistant hash function. We define the algorithms related to binary hash trees in a way that fits to our purposes.

HashTree: An algorithm that computes an (unbalanced) binary hash tree of height $h \geq 1$. Its input consists of $\ell \leq 2^h$ commitments $v_0, \dots, v_{\ell-1}$ whose hash values are the leaves of the tree, i.e., $(\text{root}, \text{tree}) \leftarrow \text{HashTree}(v_0, \dots, v_{\ell-1})$, where root is the root of the tree and tree is a sequence of all other nodes.

BuildAuth: An algorithm that on input an index k , a sequence of nodes tree , and a height h outputs the corresponding authentication path auth including the index k , i.e., $\text{auth} \leftarrow \text{BuildAuth}(k, \text{tree}, h)$.

RootCalc: An algorithm that computes the root of a hash tree given a commitment v and its authentication path auth , i.e., $\text{root} \leftarrow \text{RootCalc}(v, \text{auth})$.

In the following we define trees of commitments. The leaves are the hash values of commitments v_j , i.e., $v_0[j] = F(v_j)$ for $0 \leq j < \ell$. The inner nodes of height i are denoted by $v_i[j]$, where $0 < i \leq h$, $0 \leq j < 2^{h-i}$. They are typically computed as $v_i[j] = F(v_{i-1}[2j] \parallel v_{i-1}[2j+1])$. The root is the only node of height h , i.e., $v_h[0] = \text{root}$. A formal definition follows.

Definition 6 (Tree of Commitments). Let v_j be commitments of $\ell > 1$ secrets y_j , where $0 \leq j < \ell$. A tree of commitments is an (unbalanced) binary hash tree of height $h = \lceil \log(\ell) \rceil$, whose leaves are the hash values of v_j , i.e., $F(v_j)$. The root constitutes an aggregated commitment root, and \mathbf{auth} is the authentication path of the commitment v_k generated using the secret y_k , where $0 \leq k < \ell$.

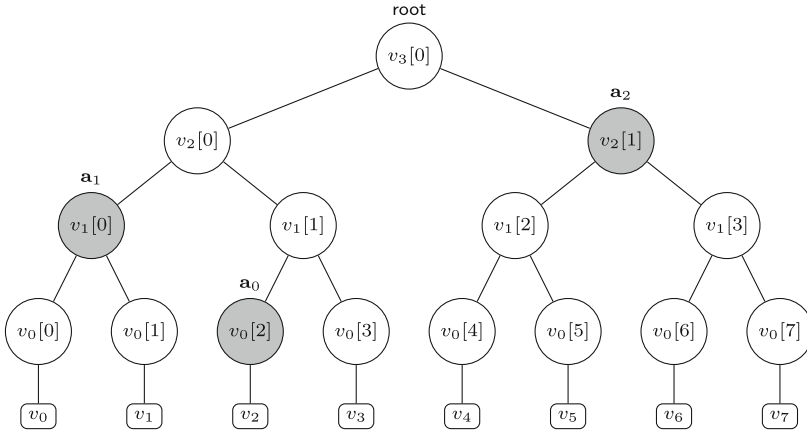


Fig. 2. A tree of commitments of height $h = 3$ and $\ell = 8$ commitments. Assume that the first time RejSamp accepts at step $k = 3$ ($0 \leq k < \ell$), then the gray colored nodes represent the authentication path required to compute the root starting from v_3 .

Next we define trees of commitments for lattice-based canonical identification (CID) schemes. Fig. 2 illustrates such a tree of height $h = 3$.

Definition 7 (Tree of Commitments for CID). Let CID be a lattice-based canonical identification scheme. Let v_j be commitments of CID generated using $\ell > 1$ masking terms y_j ($0 \leq j < \ell$). A tree of commitments for CID is an (unbalanced) binary hash tree of height $h = \lceil \log(\ell) \rceil$, whose leaves are the hash values of v_j , i.e., $F(v_j)$, and its root constitutes an aggregated commitment root to ℓ masking terms for up to ℓ repetitions within CID for the same challenge c . A response is composed of (z_k, \mathbf{auth}) , where $z_k = y_k + sc$ and y_k is the first masking term for which rejection sampling succeeds (i.e., $\text{RejSamp}(z_k) = 1$ for $0 \leq k < \ell$), and \mathbf{auth} is the authentication path of the commitment v_k generated by use of the masking term y_k .

3.2 Canonical Identification Using Trees of Commitments

Figure 3 describes a variant of the CID protocol briefly explained in Sect. 1. The variant shown here is based on MLWE and MSIS and utilizes trees of commitments. Using the Fiat-Shamir transform [15] we obtain a digital signature scheme. In the full version of this paper [3] we give a formal description of

this signature scheme and prove its correctness and security. By the equivalence results of [1], we deduce the soundness property of the CID protocol described in this section as well as its security against impersonation under passive attacks. More concretely, the main goal of providing the signature scheme and its security proof in the full version is to show how trees of commitments can also be used in lattice-based Fiat-Shamir signatures, and to establish the security of the CID protocol shown in this section based on the results of [1].




Fig. 3. Canonical identification based on MLWE and MSIS using trees of commitments.

We can choose ℓ such that at least one of the masking pairs $(\hat{\mathbf{y}}_1^{(k)}, \hat{\mathbf{y}}_2^{(k)})$ (see Fig. 3) hides $\hat{\mathbf{S}}\hat{c}$ with probability of at least $1 - \delta_{\text{abort}}$ for a given bound δ_{abort} , where $\hat{\mathbf{S}} = (\hat{\mathbf{s}}_1, \hat{\mathbf{s}}_2)$. This can be established as follows. Since the entries of the masking pairs are chosen from $D_{\mathbb{Z}^n, \sigma}$, the probability of successfully outputting $(\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2)$ with only one masking pair is $\approx 1/M$, where M is the expected number of repetitions (see Lemma 2). Consequently, one of the ℓ masking pairs conceals the secret key with probability $1 - (1 - 1/M)^\ell$. Hence, by choosing ℓ satisfying $(1 - 1/M)^\ell \leq \delta_{\text{abort}}$, the protocol aborts with probability at most δ_{abort} . For instance, to obtain a probability negligible in λ we have to select ℓ such that $(1 - 1/M)^\ell \leq 2^{-\lambda}$, which allows to completely eliminate aborts.

Let us consider an illustrative example. Suppose that we set $\delta_{\text{abort}} = 2^{-10}$ and use masking pairs with entries sampled from $D_{\mathbb{Z}^n, \sigma}$, where $\sigma = \alpha \|\hat{\mathbf{S}}\hat{c}\|$ and $M = \exp(\frac{12}{\alpha} + \frac{1}{2\alpha^2})$ (Lemma 2). Then, by setting $\alpha = 23$ we need only

$\ell = 8$ masking pairs in order to hide $\hat{\mathbf{S}}\hat{c}$ with probability at least 0.999. This means we need a tree of commitments of height $h = 3$, which is a very small tree. Regarding communication complexity, both the commitment and response consist of only 4 hash values and a pair of Gaussian vectors with $\sigma = 23 \cdot \|\hat{\mathbf{S}}\hat{c}\|$, i.e., $(\text{root}, \hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2, \text{auth} = (\mathbf{a}_0, \mathbf{a}_1, \mathbf{a}_2))$. The choice of $\alpha = 23$ increases σ in this example and hence the size of $(\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2)$ by at most 1.1 bits per integer entry in comparison to $\alpha = 11$, which is a typical choice (see, e.g., [13]) that induces a repetition rate of $M \approx 3$ and a communication complexity consisting of 3 vectors from $R_q^{k_1}$ and 3 Gaussian vectors with $\sigma = 11 \cdot \|\hat{\mathbf{S}}\hat{c}\|$. Note that in order to hide $\hat{\mathbf{S}}\hat{c}$ with probability $1 - 2^{-10}$ using a single masking pair we need to set $\alpha > 2^{13.6}$, which increases the size of the response to at least 10.1 bits per integer entry when compared with $\alpha = 11$. Hence, a larger modulus q is required and the communication complexity increases to a vector from $R_q^{k_1}$ and a Gaussian vector with a very large σ , i.e., $\sigma > 2^{13.6} \cdot \|\hat{\mathbf{S}}\hat{c}\|$.

Furthermore, we can improve the performance of protocols employing trees of commitments as follows. For subsequent executions of the protocol we can reuse the masking pairs that were sampled in previous executions but were not consumed during rejection sampling. For instance, consider the tree in Fig. 2, where the first time `RejSamp` accepts at step $k = 3$. For the next protocol run we can simply reuse the whole subtree with root $\mathbf{a}_2 = v_2[1]$ such that we only need to compute a new subtree of height $h - 1$ and combine both subtrees to obtain a new tree of height h . This decreases the number of new masking terms to be sampled and reduces the number of hash computations and multiplications modulo q . We can also lower the security requirement of the hash function F following the standard of the hash-based signature scheme XMSS [18] and using randomized tree hashing. This allows to generate trees of commitments, where F is required to be only second preimage resistant rather than collision resistant. This reduces the size of the authentication path to one half of its original size.

3.3 Optimizing the Number of Masking Terms

The previous section shows how to reduce the overall repetition rate of lattice-based protocols with multiple rejection sampling procedures. In this section we show how to minimize the height of the tree of commitments when using Gaussian distributed masking terms. This improves the performance of interactive protocols significantly. A similar approach can be taken for masking terms sampled from other distributions such as the uniform distribution.

Lemma 3. *Let $\epsilon = 2^{-\omega(\log n)}$ and M be the repetition rate of sampling masking terms from $D_{\mathbb{Z}^n, \sigma}$ such that rejection sampling succeeds. Let δ_{abort} be the desired aborting probability. Then, the number of masking terms ℓ required to conceal a secret-related term with norm bounded by T and with probability at most $1 - \delta_{\text{abort}}$ is minimized by solving the following optimization problem:*

$$\min(\ell) \text{ conditioned on } \left(1 - \frac{1 - \epsilon}{M}\right)^\ell \leq \delta_{\text{abort}}.$$

Table 2. Values for the required number of Gaussian masking terms ℓ and the bit length of the standard deviation $\sigma = \alpha T = \alpha \cdot 500$ at given aborting probabilities.

Aborting probability δ_{abort}	2^{-128}	2^{-100}	2^{-80}	2^{-40}	2^{-40}	2^{-10}	2^{-10}
Number of masking terms ℓ	64	63	62	31	16	16	8
Height of the binary hash tree h	6	6	6	5	4	4	3
Parameter α	42	30	23	23	62	12	23
Bit length of the standard deviation σ	15	14	14	14	15	13	14

Proof. Given a fixed δ_{abort} we can write ℓ as a function in M using the inequality given above. In particular, if $\sigma = \alpha T$ for $\alpha > 0$, then $M = \exp(\frac{12}{\alpha} + \frac{1}{2\alpha^2})$, $\epsilon = 2^{-100}$, and the probability of aborting using only one masking term is given by $1 - (1 - \epsilon)/M$ (see Lemma 2). Hence, ℓ can also be considered as a function in α , i.e., $\ell(\alpha) = \log(\delta_{\text{abort}}) / \log\left(1 - \frac{1 - 2^{-100}}{\exp(\frac{12}{\alpha} + \frac{1}{2\alpha^2})}\right)$. Note that increasing α directly reduces ℓ . Therefore, this problem translates to finding a local minimum of the function $\ell(\alpha)$ within a given range of α , which can be solved using Lagrange optimization. \square

The above lemma shows that reducing the number of masking terms ℓ for a fixed aborting probability δ_{abort} increases σ , hence the size of the responses (or signatures). In Table 2 we exhibit examples for various values of ℓ and $\sigma = \alpha T$ given δ_{abort} and $T = 500$.

4 Applications

As mentioned in Sect. 1, there are various advanced lattice-based constructions that are based on canonical identification (CID) and thus may benefit from using trees of commitments as described in Sect. 3. Our approach can also be applied to interactive zero-knowledge proof systems in a straightforward way. For instance, the scheme depicted in Fig. 3 can be seen as a zero-knowledge proof of knowledge of RLWE secrets with reduced communication complexity.

As a further practical application, we exploit trees of commitments within the blind signature scheme BLAZE [4] resulting in major efficiency gains. The signing protocol of BLAZE consists of 4 moves between a signer \mathcal{S} and a user \mathcal{U} . It can be aborted due to 2 rejection sampling procedures; the first one is carried out by \mathcal{S} in order to hide the secret key and the second one by \mathcal{U} to ensure blindness. In case the latter fails, \mathcal{U} must send \mathcal{S} a proof of failure in order to restart the signing protocol. This is why the last move is needed in the protocol as opposed to the standard 3-move structure of the CID scheme underlying BLAZE. Due to the possibility of failures the user must also use a statistically hiding and computationally binding commitment scheme in order to hide the message from the signer.

In the following we redesign BLAZE such that signatures can be generated in 3 moves. We call the new scheme BLAZE⁺. In particular, we are able to

completely remove the rejection sampling procedure carried out by \mathcal{U} . This is accomplished by generating enough masking terms at once such that blindness is achieved with overwhelming probability. This allows to safely eliminate the last move in the protocol and hence the need for proof of failures. Consequently, statistically hiding and computationally binding commitments concealing the message from \mathcal{S} are not required anymore. We also describe a 4-move version of BLAZE⁺ in the full version of this paper [3]. In that version aborts at the user side occur with probability of choice. We note that a similar approach may be applied at the signer side.

In addition to the functions defined in Sect. 3 we need some additional tools. Let $H : \{0, 1\}^* \rightarrow \mathbb{T}_\kappa^n$ be a public hash function modeled as a random oracle. Further, let E be a public function that expands given strings to any desired length. Sampling from $D_{\mathbb{Z},s}^n$ using randomness ρ is denoted by $D_{\mathbb{Z},s}^n(\rho)$. We let $\hat{\mathbb{T}} = \{\pm x^i : i \in [n]\} \subset R_q$. Let **Compress** and **Decompress** be functions for (de)compressing Gaussian elements (see [4] for description). Next we describe the new blind signature scheme BLAZE⁺. The respective algorithms are formalized in Fig. 4.

Key Generation. As in BLAZE, the algorithm BS.KGen generates an instance of RSIS (Fig. 4). It's secret vector \hat{s} is sampled from $D_{\mathbb{Z}^n, \sigma}^{m+1}$. However, BLAZE⁺ employs an additional condition on \hat{s} , which can also be used in BLAZE. More concretely, the ℓ_2 -norm of \hat{s} is bounded by $\gamma\sigma\sqrt{(m+1)n}$. This condition represents a trade-off between the speed of generating keys and the size of signatures, since the standard deviation s^* of masking terms used by the signer is a multiple of $\|\hat{s}\|$. Therefore, a smaller γ decreases s^* , but reduces the success probability of passing the given condition (see Lemma 1). Note that \hat{s} can also be sampled from the uniform distribution over a subset of R^{m+1} , in which the coefficients of each polynomial from R are in the set $\{-d, \dots, 0, \dots, d\}$, where $d \in \mathbb{Z}_{>0}$.

Signing. The signing algorithm is similar to that of BLAZE [4]. The difference is that in BLAZE⁺ the user \mathcal{U} generates $\ell > 1$ masking vectors $\hat{e}^{(0)}, \dots, \hat{e}^{(\ell-1)}$ chosen from $D_{\mathbb{Z}^n, s}^{m+1}$. These vectors are then used to compute $\hat{t}^{(k)} = \hat{a} \cdot \hat{e}^{(k)} + \hat{y} \pmod{q}$, which are needed to generate a tree of commitments of height $h = \lceil \log(\ell) \rceil$ via the algorithm **HashTree**. We note that generating $\hat{e}^{(k)}$ and $\hat{a} \cdot \hat{e}^{(k)} \pmod{q}$, for $k = 0, \dots, \ell - 1$, can be precomputed by \mathcal{U} before starting the protocol with \mathcal{S} . The sum $\hat{t}^{(k)}$ containing \hat{y} and the construction of the tree cannot be carried out in advance, since \hat{y} is computed from the commitment sent by \mathcal{S} (see Fig. 4). We also note that $\hat{e}^{(k)}$ can be reused when \mathcal{S} restarts the protocol, since $\hat{e}^{(k)}$ are not revealed and \hat{c}^* is always fresh. After receiving the vector \hat{z}^* , \mathcal{U} computes \hat{z} and the authentication path **auth** associated to the first index $k < \ell$ for which the vector $\hat{e}^{(k)}$ ensures blindness. Note that ℓ is chosen such that this happens with probability at least $1 - 2^{-\lambda}$, i.e., \mathcal{U} outputs a valid signature with overwhelming probability. Also note that for each signature a new root is generated.

Verification. Verifying a signature is straightforward as described in Fig. 4.

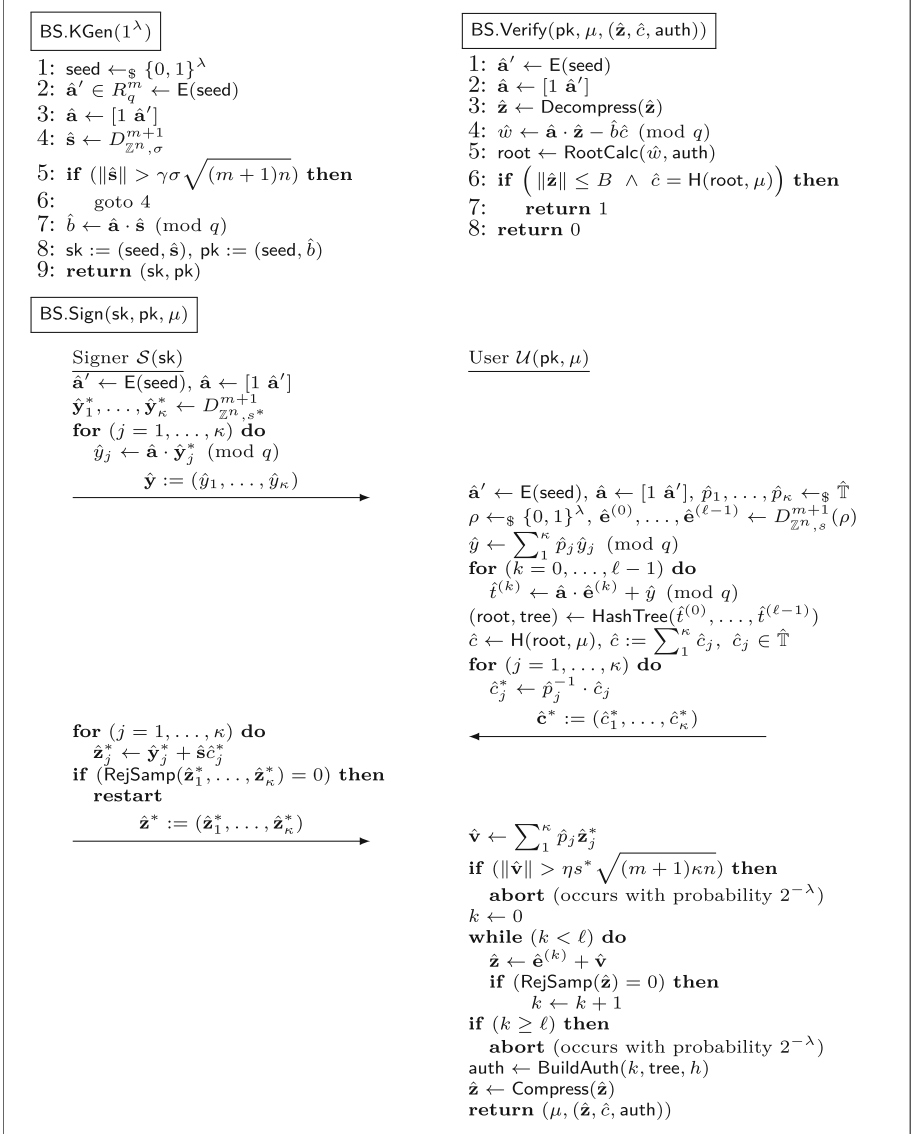


Fig. 4. A formal description of the new blind signature scheme BLAZE⁺.

In the following we give the main security statements of this section comprising completeness, blindness, and strong one-more unforgeability of BLAZE⁺. The proofs of both correctness and blindness directly follow from [4] and are hence omitted. In particular, proving blindness requires to show that the exchanged messages during protocol execution together with the user's output does not leak any information about the message being signed. In comparison to BLAZE, the

authentication path \mathbf{auth} , which is a part of the signature in BLAZE^+ is the only additional information exchanged between the signer and the user. Obviously, \mathbf{auth} does not give any information about the signed message.

Theorem 1 (Completeness). *Let $\alpha^*, \alpha, \gamma, \eta > 0$, $s^* = \alpha^* \gamma \sigma \sqrt{(m+1)\kappa n}$, $s = \eta \alpha s^* \sqrt{(m+1)\kappa n}$, and $B = \eta s \sqrt{(m+1)n}$. Further, let $(1 - \frac{1-2^{-100}}{U})^\ell \leq 2^{-\lambda}$, where $U = \exp(\frac{12}{\alpha} + \frac{1}{2\alpha^2})$. After at most M repetitions, any blind signature produced by BLAZE^+ is validated with probability of at least $1 - 2^{-\lambda}$, where $M = \exp(\frac{12}{\alpha^*} + \frac{1}{2\alpha^{*2}})$.*

Theorem 2 (Blindness). *The scheme BLAZE^+ is statistically blind. The statistical distance between two executions of its signing protocol is given by $2^{-100}/M$.*

Next, we prove the strong one-more unforgeability of BLAZE^+ . In the proof we assume that in the unforgeability game (see Definition 3) any forgery output by the adversary \mathcal{A} is considered valid if and only if its associated root was not queried to the signing oracle by \mathcal{A} , i.e., a forgery must contain a new root that is distinct from the roots queried to the signing oracle. We note that our 3-move protocol achieves completeness with overwhelming probability, i.e., each completed interaction yields a valid blind signature, where users do not have to request a protocol restart. Therefore, the one-more unforgeability proof does not need to consider aborts at the user side as opposed to BLAZE .

Theorem 3 (Unforgeability). *The scheme BLAZE^+ is strongly one-more unforgeable in the random oracle model (ROM) if F is a collision resistant hash function and RSIS is hard. More precisely, suppose that F is collision resistant and it is hard to find a vector $\hat{\mathbf{x}} \neq \mathbf{0}$ satisfying $[1 \ \hat{\mathbf{a}}'] \cdot \hat{\mathbf{x}} = 0 \pmod{q}$ and $\|\hat{\mathbf{x}}\| \leq \beta$ for $\beta = 2(B + \eta \sigma \sqrt{(m+1)\kappa n})$, then BLAZE^+ is strongly one-more unforgeable in the ROM.*

Proof. We assume that there exists a forger \mathcal{A} that wins the one-more unforgeability game given in Definition 3 with probability $\varepsilon_{\mathcal{A}}$. We construct a reduction algorithm \mathcal{D} that finds collisions in the hash function F or computes a vector $\hat{\mathbf{x}} \neq \mathbf{0}$ as described in the theorem statement with probability $\varepsilon_{\mathcal{D}}$ as given below.

Setup. The input of \mathcal{D} is a uniformly random vector $\hat{\mathbf{a}}' \in R_q^m$ and a hash function F . It also has access to an oracle \mathcal{O}_F for F . The reduction \mathcal{D} samples a vector $\hat{\mathbf{s}}$ from $D_{\mathbb{Z}^n, \sigma}^{m+1}$ and randomly selects answers for random oracle queries $\{\hat{h}_1, \dots, \hat{h}_{q_H}\}$. Then, it runs the forger \mathcal{A} with public key $(\hat{\mathbf{a}}, \hat{b})$, where $\hat{\mathbf{a}} = [1 \ \hat{\mathbf{a}}']$ and $\hat{b} = \hat{\mathbf{a}} \cdot \hat{\mathbf{s}} \pmod{q}$.

Random Oracle Query. The reduction \mathcal{D} maintains a list L_H , which includes pairs of random oracle queries and their answers. If H was previously queried on some input, then \mathcal{D} looks up its entry in L_H and returns its answer $\hat{c} \in \mathbb{T}_{\kappa}^n$. Otherwise, it returns the first unused \hat{c} and updates the list.

Hash Query. Hash queries to F sent by \mathcal{A} are forwarded to the oracle \mathcal{O}_F . The reduction \mathcal{D} also maintains a list L_F , which includes pairs of hash queries to F and their answers as well as the structure of the trees.

Blind Signature Query. Upon receiving signature queries from the forger \mathcal{A} as a user, \mathcal{D} interacts as a signer with \mathcal{A} according to the signing protocol (see Fig. 4).

Output. After $k \leq q_{\text{Sign}}$ successful executions of the signing protocol, \mathcal{A} outputs $k + 1$ distinct messages and their valid signatures $(\mu_1, \text{sig}_1), \dots, (\mu_{k+1}, \text{sig}_{k+1})$. Then, one of the following two cases applies.

Case 1. \mathcal{D} finds two signatures of messages $\mu, \mu' \in \{\mu_1, \dots, \mu_{k+1}\}$ with the same random oracle answer \hat{c} . In this case BS.Verify yields $H(\text{root}, \mu) = H(\text{root}', \mu')$. If $\mu \neq \mu'$ or $\text{root} \neq \text{root}'$, then a second preimage of \hat{c} has been found by \mathcal{A} . If $\mu = \mu'$ and $\text{root} = \text{root}'$, then both signatures were generated using the same hash tree. This does not follow the unforgeability game because the output of \mathcal{A} does not include a valid forgery.

Case 2. If all signatures output by \mathcal{A} have distinct random oracle answers, then \mathcal{D} guesses an index $i \in [k+1]$ such that $\hat{c}_i = \hat{h}_j$ for some $j \in [q_H]$. Then, it records the pair $(\mu_i, (\hat{\mathbf{z}}_i, \hat{c}_i, \text{auth}_i))$ and invokes \mathcal{A} again with the same random tape and the random oracle queries $\{\hat{h}_1, \dots, \hat{h}_{j-1}, \hat{h}'_j, \dots, \hat{h}'_{q_H}\}$, where $\{\hat{h}'_j, \dots, \hat{h}'_{q_H}\}$ are fresh random elements. After the second invocation, the output of \mathcal{A} includes a pair $(\mu'_i, (\hat{\mathbf{z}}'_i, \hat{c}'_i, \text{auth}'_i))$. By the General Forking Lemma [8], \mathcal{A} outputs a forgery containing \hat{c}'_i with probability $\varepsilon_{\text{fork}}$ (see below), where $\hat{c}_i \neq \hat{c}'_i$ and $\text{root} = \text{root}'$. Let $\hat{w} = \hat{\mathbf{a}} \cdot \hat{\mathbf{z}}_i - \hat{b}\hat{c}_i \pmod{q}$ and $\hat{w}' = \hat{\mathbf{a}} \cdot \hat{\mathbf{z}}'_i - \hat{b}\hat{c}'_i \pmod{q}$. Then, one of the following holds:

1. $\hat{\mathbf{z}}_i \neq \hat{\mathbf{z}}'_i$ and $\text{auth}_i = \text{auth}'_i$. If $\hat{w} = \hat{w}'$, then $\hat{\mathbf{a}}(\hat{\mathbf{z}}_i - \hat{\mathbf{z}}'_i) - \hat{b}(\hat{c}_i - \hat{c}'_i) = 0 \pmod{q}$. Therefore, by setting $\hat{b} = \hat{\mathbf{a}} \cdot \hat{\mathbf{s}} \pmod{q}$ we obtain $\hat{\mathbf{a}} \cdot \hat{\mathbf{x}} = 0 \pmod{q}$, where $\hat{\mathbf{x}} = \hat{\mathbf{z}}_i - \hat{\mathbf{z}}'_i - \hat{\mathbf{s}}(\hat{c}_i - \hat{c}'_i)$. Since both signatures are valid, we have $\|\hat{\mathbf{z}}_i\| \leq B$ and $\|\hat{\mathbf{z}}'_i\| \leq B$. Moreover we have $\|\hat{\mathbf{s}}(\hat{c}_i - \hat{c}'_i)\| \leq 2\eta\sigma\sqrt{(m+1)\kappa n}$. Hence, $\|\hat{\mathbf{x}}\| \leq 2(B + \eta\sigma\sqrt{(m+1)\kappa n})$. This constitutes a solution to RSIS with norm bound β . If $\hat{w} \neq \hat{w}'$, then a collision in F has been found in the leaves of the hash tree.
2. $\hat{\mathbf{z}}_i \neq \hat{\mathbf{z}}'_i$ and $\text{auth}_i \neq \text{auth}'_i$. If $\hat{w} = \hat{w}'$, then we have a solution to RSIS (as in 1.). If $\hat{w} \neq \hat{w}'$, then we consider two cases: If \hat{w}, \hat{w}' belong to different hash trees, then a collision has occurred in F similar to [26], i.e., there exists an index $j \in \{0, \dots, h-1\}$ such that $\mathbf{a}_j \neq \mathbf{a}'_j$, where $\mathbf{a}_j \in \text{auth}_i$, $\mathbf{a}'_j \in \text{auth}'_i$ and $\text{RootCalc}(\hat{w}, \text{auth}_i) = \text{root} = \text{RootCalc}(\hat{w}', \text{auth}'_i)$. If \hat{w}, \hat{w}' belong to the same hash tree, then \mathcal{D} keeps the pair $(\mu_i, (\hat{\mathbf{z}}_i, \hat{c}_i, \text{auth}_i))$ and invokes \mathcal{A} at most ℓ times with the same random tape and the random oracle queries $\{\hat{h}_1, \dots, \hat{h}_{j-1}, \hat{h}_j^{(t)}, \dots, \hat{h}_{q_H}^{(t)}\}$ (where $t \in \{0, \dots, \ell-1\}$) until we obtain two forgeries such that the associated leaves have the same index in the tree and the same hash value. If $\hat{w} = \hat{w}'$, then we have a solution to RSIS (as in 1.). Otherwise, we have a collision in F .

3. $\hat{\mathbf{z}}_i = \hat{\mathbf{z}}'_i$ and $\mathbf{auth}_i = \mathbf{auth}'_i$. If $\hat{w} = \hat{w}'$, then $\hat{\mathbf{a}} \cdot \hat{\mathbf{s}}(\hat{c}_i - \hat{c}'_i) = \hat{b}(\hat{c}_i - \hat{c}'_i) = 0 \pmod{q}$ and $\|\hat{\mathbf{s}}(\hat{c}_i - \hat{c}'_i)\| \leq \beta$. Since $\hat{c}_i \neq \hat{c}'_i$ and R is an integral domain, then $\hat{\mathbf{s}}(\hat{c}_i - \hat{c}'_i) \in R^{m+1} \setminus \{\mathbf{0}\}$. This constitutes a solution to RSIS. Note that if \hat{b} is invertible in R_q , then we obtain $\hat{c}_i - \hat{c}'_i = 0 \pmod{q}$. This contradicts $\hat{c}_i \neq \hat{c}'_i$. Moreover, if $(\hat{c}_i - \hat{c}'_i)$ is invertible in R_q , then $\hat{b} = 0 \pmod{q}$, which is not the case. If $\hat{w} \neq \hat{w}'$, then a collision has occurred in F (in the leaves of the hash tree).
4. $\hat{\mathbf{z}}_i = \hat{\mathbf{z}}'_i$ and $\mathbf{auth}_i \neq \mathbf{auth}'_i$. If $\hat{w} = \hat{w}'$, then we have a solution to RSIS (as in **3.**). If $\hat{w} \neq \hat{w}'$, then one of the cases considered in **2.** (for $\hat{w} \neq \hat{w}'$) applies.

The reduction \mathcal{D} retries at most $q_{\mathbf{H}}^{(k+1)}$ times with different random tape and random oracle queries.

Analysis. First, we note that the environment of \mathcal{A} is perfectly simulated by \mathcal{D} and signatures are generated with the same probability as in the real execution of the signing protocol. Next, one of the $k+1$ pairs output by \mathcal{A} is by assumption not generated during the execution of the signing protocol. The probability of correctly guessing the index i corresponding to this pair is $1/(k+1)$, where there are $q_{\mathbf{H}}^{k+1}$ index pairs (i, j) such that $\hat{c}_i = \hat{h}_j$. Therefore, one of the $q_{\mathbf{H}}^{k+1}$ reruns of \mathcal{A} yields the correct index pair (i, j) . The probability that \hat{c}_i was a random oracle query made by \mathcal{A} is at least $1 - 1/|\mathbb{T}_{\kappa}^n|$. Thus, the probability that $\hat{c}_i = \hat{h}_j$ is $\varepsilon_{\mathcal{A}} - 1/|\mathbb{T}_{\kappa}^n|$. By the General Forking Lemma with at most $\ell = O(1)$ rewindings and distinct $\hat{h}_j^{(t)}, \dots, \hat{h}_{q_{\mathbf{H}}}^{(t)}$, we have $\varepsilon_{\text{fork}} \geq (\varepsilon_{\mathcal{A}} - \frac{1}{|\mathbb{T}_{\kappa}^n|}) \cdot ((\frac{\varepsilon_{\mathcal{A}} - 1/|\mathbb{T}_{\kappa}^n|}{q_{\text{Sign}} + q_{\mathbf{H}}})^{\ell} - \tau)$, where $\tau = 1 - \prod_{t=1}^{\ell} \frac{|\mathbb{T}_{\kappa}^n| - t}{|\mathbb{T}_{\kappa}^n|} \leq 1 - \left(\frac{|\mathbb{T}_{\kappa}^n| - \ell}{|\mathbb{T}_{\kappa}^n|}\right)^{\ell} \leq \frac{\ell^2}{|\mathbb{T}_{\kappa}^n|}$. Since $\hat{\mathbf{s}}$ is not uniquely defining \hat{b} when $(m+1)\log(2d) > \log(q)$ (see, e.g., [28]) for a sufficiently large d that is related to the size of the coefficients of $\hat{\mathbf{s}}$, \mathcal{A} does not know which $\hat{\mathbf{s}}$ is being used to construct $\hat{\mathbf{x}}$. Hence, $\hat{\mathbf{x}} \neq \mathbf{0}$ with probability at least $1/2$ (see, e.g., [25, 33]). This can be easily shown, e.g., when the coefficients of $\hat{\mathbf{s}}$ are uniformly distributed over $\{-d, \dots, 0, \dots, d\}$. The success probability of \mathcal{D} is given by $\varepsilon_{\mathcal{D}} \geq \frac{\varepsilon_{\text{fork}}}{2(k+1)}$, which is non-negligible if $\varepsilon_{\mathcal{A}}$ is non-negligible. \square

Parameters. Table 3 shows our proposed parameters for BLAZE⁺, which are selected for approximately 128 bits of security. The table also reviews the parameters of BLAZE proposed in [4] for the same security level. Table 1 gives the related communication complexity, performance, and sizes of keys and signatures. For the sake of comparison with BLAZE, we choose $m = 1$ and $m = 3$ for a practical scheme instantiation. Note that the choice of e.g., $m = 1$ implies an instantiation that is based on RLWE rather than RSIS and hence, it is not covered by the security proof as indicated in [4]. This is because the secret key has insufficient entropy or does not satisfy the condition $(m+1)\log(2d) > \log(q)$ (see the proof of Theorem 3). It seems that using RLWE does not reduce the security of the scheme, but rather using RSIS appears to be more an artifact of the proof technique.

Table 3. Parameters for BLAZE⁺ and BLAZE targeting approximately 128 bits of security. The performance, sizes, and communication complexity corresponding to these parameters are given in Table 1.

Scheme	Parameters															
	δ_{abort}	ℓ	h	n	m	q	σ	γ	κ	α^*	α	s^*	s	M_S	M_U	M
BLAZE ⁺	2^{-128}	71	7	1024	1	$\approx 2^{31}$	0.5	1.01	16	19	33	1736.9	12450734	1.9	1	1.9
BLAZE ⁺	2^{-40}	32	5	1024	1	$\approx 2^{31}$	0.5	1.01	16	28	22	2559.6	12232099	1.5	1	1.5
BLAZE ⁺	2^{-10}	8	3	1024	1	$\approx 2^{31}$	0.5	1.01	16	28	22	2559.6	12232099	1.5	1	1.5
BLAZE	0.38	1	0	1024	1	$\approx 2^{31}$	0.5	1.2	16	20	25	2172.2	11796306	1.8	1.6	2.9
BLAZE ⁺	2^{-128}	71	7	1024	3	$\approx 2^{31}$	9.6	1.01	16	19	33	47161.3	478102394	1.9	1	1.9
BLAZE ⁺	2^{-40}	32	5	1024	3	$\approx 2^{31}$	9.6	1.01	16	28	22	69500.9	469714882	1.5	1	1.5
BLAZE ⁺	2^{-10}	8	3	1024	3	$\approx 2^{31}$	9.6	1.01	16	28	22	69500.9	469714882	1.5	1	1.5
BLAZE	0.38	1	0	1024	3	$\approx 2^{31}$	9.6	1.2	16	20	25	54067.2	380633088	1.8	1.6	2.9

Acknowledgements. We thank the anonymous reviewers of ACISP’20 for their valuable comments. This work has been partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1119 – 236615297.

References

1. Abdalla, M., An, J.H., Bellare, M., Namprempre, C.: From identification to signatures via the fiat-shamir transform: minimizing assumptions for security and forward-security. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 418–433. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_28
2. Ajtai, M.: Generating hard instances of lattice problems. In: ACM Symposium on Theory of Computing - STOC 1996, pp. 99–108. ACM (1996)
3. Alkeilani Alkadri, N., El Bansarkhani, R., Buchmann, J.: On lattice-based interactive protocols: an approach with less or no aborts. Cryptology ePrint Archive, Report 2020/007 (2020). <http://eprint.iacr.org/2020/007>. Full version of this paper
4. Alkeilani Alkadri, N., El Bansarkhani, R., Buchmann, J.: BLAZE: practical lattice-based blind signatures for privacy-preserving applications. In: Financial Cryptography and Data Security - FC 2020. Springer, Cham (2020). <http://eprint.iacr.org/2019/1167>
5. Bai, S., Galbraith, S.D.: An improved compression technique for signatures based on learning with errors. In: Benaloh, J. (ed.) CT-RSA 2014. LNCS, vol. 8366, pp. 28–47. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-04852-9_2
6. Baum, C., Damgård, I., Lyubashevsky, V., Oechsner, S., Peikert, C.: More efficient commitments from structured lattice assumptions. In: Catalano, D., De Prisco, R. (eds.) SCN 2018. LNCS, vol. 11035, pp. 368–385. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-98113-0_20
7. Baum, C., Lin, H., Oechsner, S.: Towards practical lattice-based one-time linkable ring signatures. In: Naccache, D., Xu, S., Qing, S., Samarati, P., Blanc, G., Lu, R., Zhang, Z., Meddahi, A. (eds.) ICICS 2018. LNCS, vol. 11149, pp. 303–322. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01950-1_18

8. Bellare, M., Neven, G.: Multi-signatures in the plain public-key model and a general forking lemma. In: ACM Conference on Computer and Communications Security - CCS 2006, pp. 390–399. ACM (2006)
9. Benhamouda, F., Camenisch, J., Krenn, S., Lyubashevsky, V., Neven, G.: Better zero-knowledge proofs for lattice encryption and their application to group signatures. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8873, pp. 551–572. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45611-8_29
10. Bourse, F., Del Pino, R., Minelli, M., Wee, H.: FHE circuit privacy almost for free. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9815, pp. 62–89. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_3
11. Brakerski, Z., Perlman, R.: Order-LWE and the hardness of Ring-LWE with entropic secrets. Cryptology ePrint Archive, Report 2018/494 (2018). <https://eprint.iacr.org/2018/494>
12. Ducas, L., et al.: CRYSTALS-Dilithium: a lattice-based digital signature scheme. *Trans. Cryptogr. Hardware Embed. Syst. TCHES* **2018**(1), 238–268 (2018)
13. Ducas, L., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., Stehle, D.: CRYSTALS-Dilithium: digital signatures from module lattices. *Cryptology ePrint Archive, Report 2017/633* (2017). Version: 20170627:201152. <http://eprint.iacr.org/2017/633>
14. El Bansarkhani, R., Sturm, J.: An efficient lattice-based multisignature scheme with applications to bitcoins. In: Foresti, S., Persiano, G. (eds.) CANS 2016. LNCS, vol. 10052, pp. 140–155. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48965-0_9
15. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_12
16. Fischlin, M., Schröder, D.: On the impossibility of three-move blind signature schemes. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 197–215. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_10
17. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: ACM Symposium on Theory of Computing - STOC 2009, pp. 169–178. ACM (2009)
18. Hülsing, A., Butin, D., Gazdag, S., Rijneveld, J., Mohaisen, A.: XMSS: eXtended Merkle Signature Scheme. RFC 8391, May 2018
19. Juels, A., Luby, M., Ostrovsky, R.: Security of blind digital signatures. In: Kaliski, B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 150–164. Springer, Heidelberg (1997). <https://doi.org/10.1007/BFb0052233>
20. Katz, J., Kolesnikov, V., Wang, X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. In: ACM Conference on Computer and Communications Security - CCS 2018, pp. 525–537. ACM (2018)
21. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. *Des. Codes Crypt.* **75**(3), 565–599 (2014). <https://doi.org/10.1007/s10623-014-9938-4>
22. Lyubashevsky, V.: Fiat-Shamir with aborts: applications to lattice and factoring-based signatures. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_35
23. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_43

24. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_1
25. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 35–54. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_3
26. Merkle, R.C.: A certified digital signature. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 218–238. Springer, New York (1990). https://doi.org/10.1007/0-387-34805-0_21
27. Micciancio, D.: Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In: Proceedings of the 43rd Symposium on Foundations of Computer Science FOCS, pp. 356–365. IEEE (2002)
28. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.* **37**(1), 267–302 (2007)
29. von Neumann, J.: Various techniques used in connection with random digits. In: Monte Carlo Method. National Bureau of Standards Applied Mathematics Series, vol. 12, pp. 36–38 (1951)
30. del Pino, R., Lyubashevsky, V.: Amortization with fewer equations for proving knowledge of small secrets. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10403, pp. 365–394. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_13
31. Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. *J. Cryptol.* **13**(3), 361–396 (2000)
32. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: ACM Symposium on Theory of Computing, pp. 84–93. ACM (2005)
33. Rückert, M.: Lattice-based blind signatures. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 413–430. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_24
34. Schnorr, C.P.: Efficient signature generation by smart cards. *J. Cryptol.* **4**(3), 161–174 (1991). <https://doi.org/10.1007/BF00196725>
35. Alberto Torres, W.A., et al.: Post-quantum one-time linkable ring signature and application to ring confidential transactions in blockchain (lattice RingCT v1.0). In: Susilo, W., Yang, G. (eds.) ACISP 2018. LNCS, vol. 10946, pp. 558–576. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93638-3_32