

Lattice-Based Blind Signatures

Markus Rückert*

Technische Universität Darmstadt
Department of Computer Science
Cryptography and Computeralgebra, Germany
`markus.rueckert@cased.de`

Abstract. Blind signatures (BS), introduced by Chaum, have become a cornerstone in privacy-oriented cryptography. Using hard lattice problems, such as the shortest vector problem, as the basis of security has advantages over using the factoring or discrete logarithm problems. For instance, lattice operations are more efficient than modular exponentiation and lattice problems remain hard for quantum and subexponential-time adversaries. Generally speaking, BS allow a signer to sign a message without seeing it, while retaining a certain amount of control over the process. In particular, the signer can control the number of issued signatures. For the receiver of the signature, this process provides perfect anonymity, e.g., his spendings remain anonymous when using BS for electronic money.

We provide a positive answer to the question of whether it is possible to implement BS based on lattice problems. More precisely, we show how to turn Lyubashevsky's identification scheme into a BS scheme, which has almost the same efficiency and security in the random oracle model. In particular, it offers quasi-linear complexity, statistical blindness, and its unforgeability is based on the hardness of worst-case lattice problems with an approximation factor of $\tilde{O}(n^5)$ in dimension n . Moreover, it is the first blind signature scheme that supports leakage-resilience, tolerating leakage of a $(1 - o(1))$ fraction of the secret key in a model that is inspired by Katz and Vaikuntanathan.

Keywords: Blind signatures, post-quantum, lattices, provable security, leakage resilience.

1 Introduction

Since Chaum proposed his idea of blind signatures [Cha82], it has become an important primitive for anonymous Internet banking, e-voting (e.g., [RHOAGZ07]), as well as for oblivious transfer [CNS07]. These applications will retain their importance in both, near and far future. As for the near future, we are convinced that current factoring and discrete logarithm based instantiations are efficient and secure. *But for how long?*

Today, when building provably secure cryptographic schemes, one also has to anticipate emerging technologies that may lead to new attacks. This is why we

* This work was supported by CASED (www.cased.de).

typically try to use the mildest possible assumptions. Let us consider the example of quantum computers as a metaphor for these future developments. In the quantum-age, the cryptographic assumptions change with the leap in computing power that quantum computers will provide. There are only a few cryptographic assumptions that are conjectured to be *post-quantum*, i.e., they are considered to withstand quantum computer attacks. One of those assumptions is the hardness of finding short vectors in a lattice. Even for today, there are benefits when building cryptography upon hard lattice problems because, unlike factoring, they withstand subexponential attacks and the best known algorithms, e.g., [AKS01], have an exponential complexity in the lattice dimension. Furthermore, lattice problems typically allow a worst-case to average-case reduction that goes back to Ajtai [Ajt96]. It states that a randomly chosen instance of a certain lattice problem is at least as hard as the worst-case instance of a related lattice problem. Thus, choosing secure keys is easy. This reduction was later on adapted to work with ideal lattices by Lyubashevsky and Micciancio [LM06] because ideal lattices offer a compact public-key representation and very efficient operations at the expense of a slightly stronger assumption.

The security model, mainly influenced by Juels, Luby, and Ostrovsky [JLO97] as well as Pointcheval and Stern [PS00], requires blind signature schemes to satisfy blindness and one-more unforgeability. Blindness states that the signer must not obtain any information on the signed messages and one-more unforgeability means that an adversary cannot obtain more signatures than there were interactions with the signer.

Our Contribution. We construct the first lattice-based blind signature scheme. It is inspired by Lyubashevsky’s ID scheme [Lyu08] in combination with the Fiat-Shamir paradigm [FS86]. It is unconditionally blind, selective-failure blind [CNS07], and one-more unforgeable in the random oracle model [BR93] if standard lattice problems in ideal lattices [LM06] are hard in the worst-case. With its four moves it is quite efficient. All operations have quasi-linear complexity and all keys and signatures require a quasi-linear amount of storage bits, with respect to the main parameter n . Moreover, it is leakage resilient according to a model inspired by Katz and Vaikuntanathan [KV09]. Let L be the bit-length of the secret key. Our scheme remains secure, even if the adversary obtains $L(1 - o(1))$ bits of the secret key via arbitrary side channels. This brings the security model closer to reality, where the adversary may obtain information about the secret key, e.g. via (remote) timing attacks or by having physical access to the signing device. When applied in e-voting or e-cash schemes, such a resilience also helps against insider attacks and may improve the trust that we are willing to grant these schemes. Another application of our construction is identity-based blind signatures, when combined with [Rüc10].

Our scheme is also the first leakage resilient blind signature scheme and our results in this respect are applicable to Lyubashevsky’s ID and signature schemes [Lyu08, Lyu09]. It may be possible to use an analogue of Pointcheval and Stern’s approach [PS00] to turn the leakage resilient variants [KV09, ADW09] of the Okamoto-Schnorr signature scheme [Sch91, Oka92] into blind signature schemes.

Table 1. Comparison of RSA, Okamoto-Schnorr, and our blind signature scheme

Scheme	Secure until	Security (bits)	Moves	KeyGen	Sign	Verify
RSA-1229	2012	Current (76)	2	95 ms	16 ms	5 ms
RSA-3313	2050	Medium (102)	2	1250 ms	46 ms	6 ms
RSA-15424	2282	Future (256)	2	251849 ms	2134 ms	20 ms
OS-1229	2012	Current (76)	3	16 ms	64 ms	24 ms
OS-3313	2050	Medium (102)	3	46 ms	184 ms	69 ms
OS-15424	2282	Future (256)	3	2134 ms	8536 ms	3201 ms
Section 3 ($n = 1024$)	2012	Current (76)	4	37 ms	220 ms	33 ms
Section 3 ($n = 2048$)	2050	Medium (102)	4	52 ms	283 ms	57 ms
Section 3 ($n = 8192$)	2282	Future (256)	4	305 ms	1175 ms	320 ms

The table compares our scheme with RSA and Okamoto-Schnorr for various moduli according to [Len05] (Current, Medium) and [ECR10] (Future). The bitlengths can be computed on www.keylength.com. For our blind signature scheme, we propose three *optimized* parameter sets for the same security levels based on [RS10], which provides a framework for choosing secure parameters for lattice-based cryptography. Note that the parameters for RSA and OS do *not* take potential quantum-computer attacks into account. All timings are averaged over 1000 random instances.

However, it is unclear whether this will actually work and whether it will be efficient.

Table 1 compares RSA and Okamoto-Schnorr (OS) blind signatures with our construction in terms of computational cost. For all schemes, we propose parameter sets for current, medium, and future security levels. We believe that RSA is a good basis for comparison because it is easy to understand and very efficient as signing only involves two modular exponentiations and verification can be done in a single one (small exponent). We do *not* count multiplications. As observed in [BNPS03], the security of the RSA blind signature scheme is based on a specially tailored interactive assumption that is stronger than the original RSA assumption [BMV08]. Taking all this into account, the timings observed for RSA provide an optimistic lower bound for current practical and secure schemes. The timings for OS are expected timings based on the number of modular exponentiations, *not* counting multiplications. We include OS because it follows the typical 3-move structure and is based on a standard assumption. It is therefore closer to our protocol. The timings were obtained on an AMD Opteron CPU, running at 2.3 GHz. For RSA and OS, we have used OpenSSL 0.9.8g, which is supposed to be very efficient. For our blind signature schemes, we did a straightforward implementation, which certainly leaves room for improvements. Here, the timings reflect the *full* scheme.

From Table 1, we clearly see that our scheme benefits from its quasi-linear complexity, especially in higher levels of security. In addition, for our scheme, we

can have various trade-offs between signature size and speed. For more details, refer to the full version [Rüc08]. There, we also show how to optimize the key and signature sizes, which are typically large in lattice-based constructions.

We believe that our work is an important contribution because the previous efficient constructions, such as [Cha82, PS97, PS00, Abe01, BNPS03, CKW04, Oka06], have one thing in common: they are built upon classic number theoretic assumptions, like the hardness of factoring large integers or computing discrete logarithms. The more recent approaches, e.g., by Boldyreva [Bol03] or Okamoto [Oka06], tend to use pairings that yield very elegant constructions. They, however, are again based on the discrete logarithm problem in this specific setting. None of the above schemes remains secure in the presence of reasonably large quantum computers, where both factoring and computing discrete logarithms become easy due to the seminal work of Shor [Sho97].

Main Obstacles. For every blind signature scheme, one has to overcome three basic obstacles. The scheme needs to be blind, one-more unforgeable, and at the same time complete. Blindness and unforgeability are already somewhat orthogonal because granting the user too much power to ensure blindness harms unforgeability and vice-versa. Since working with lattices, we do not have access to a cyclic group structure as in schemes that are based on the DDH or DL assumptions. There, blindness is typically easier to achieve by multiplying the message with a random group element. The result is again a random group element.

In lattices, we need to emulate this over an infinite structure via a filtering technique that is inspired by [Lyu08]. However, this technique introduces a completeness defect that even affects the interaction of an honest user with an honest signer. Thus, the protocol may need to be restarted. We show how this technique can be refined to allow a time-memory trade-off, reducing the number of expected restarts at the expense of only slightly larger signatures. When addressing this defect, we need additional means to ensure blindness over repetitions of the protocol. Our solution involves a statistically hiding commitment.

Similarly, the completeness defect has implications with respect to unforgeability as the user may claim that the protocol has failed, whereas it was indeed successful. Here, we extend the typical three-move structure to a four-move structure where the user needs to demonstrate that he or she could not obtain a valid signature. Such a last move, from user to signer, is highly unusual for blind signature schemes. We solve this issue by designing a special proof of failure and by employing a computationally binding commitment scheme.

All these issues, and the additional leakage resilience, need to be addressed simultaneously as they are interconnected. This leads to an intricate process of correctly setting up the numerous parameters and sets for our scheme in Table 2.

RSA-style Blind Signatures. One might think that RSA-style (hash \rightarrow blind \rightarrow invert \rightarrow unblind) lattice-based blind signatures can be implemented using the preimage sampleable trapdoor function $f : D \subset \mathbb{Z}^m \rightarrow \mathbb{Z}_q^n$ from [GPV08]. If certain lattice problems are hard, it is hard to sample preimages from D (small norm) unless one knows short vectors x such that $f(x) = 0$. The user

would hash the message M using a full-domain hash $h \leftarrow \mathsf{H}(M)$ and blind using $M^* \leftarrow h + f(\beta)$ for $\beta \in D$. The signer would *sample* from $f^{-1}(M^*) \cap D$ and return the result σ^* . The function is compressing, so there are no unique preimages. Using β and the fact that f is linear, the user can compute $\sigma \leftarrow \sigma^* - \beta$, which passes verification: $f(\sigma) = f(\sigma^*) - f(\beta) = \mathsf{H}(M^*)$. For the proof, one would rely on an interactive “one-more“ trapdoor inversion assumption akin to [BNPS03]. However, the adversary must never obtain a non-zero $x \in D$ such that $f(x) = 0$ because this would imply learning a piece of the secret key. Unfortunately, such an attack is easy: take $u \in D$ and send $M^* = f(u)$ to the signer, who returns σ^* . Now, $x = u - \sigma^*$ is small and $f(x) = 0$. Also, $x \neq 0$ with high probability because there are many preimages of $f(u)$.

Organization. After a brief preliminaries section, we propose our blind signature scheme in Section 3. There, we also provide a detailed analysis, including completeness, blindness, one-more unforgeability, and leakage resilience. The full version of the paper is [Rüc08]. There, we discuss how to choose practical parameters and prove all supporting lemmas for the theorems in Section 3.

2 Preliminaries

With n , we always denote the security parameter. The joint execution of two algorithms \mathcal{A} and \mathcal{B} in an interactive protocol with private inputs x to \mathcal{A} and y to \mathcal{B} is written as $(a, b) \leftarrow \langle \mathcal{A}(x), \mathcal{B}(y) \rangle$. The private outputs are a for \mathcal{A} and b for \mathcal{B} . Accordingly, $\langle \mathcal{A}(x), \mathcal{B}(y) \rangle^k$ means that the interaction can take place up to k times. The statement $x \leftarrow_{\mathfrak{s}} X$ means that x is chosen uniformly at random from the finite set X . Recall that the statistical distance of two random variables X, Y over a discrete domain D is defined as $\Delta(X, Y) = 1/2 \sum_{a \in D} |\text{Prob}[X = a] - \text{Prob}[Y = a]|$. A function is negligible if it vanishes faster than $1/p(n)$ for any polynomial p . All logarithms are base 2 and we identify $\{1, \dots, k\}$ with $[k]$.

We recall the definitions of blind signatures and commitments. Afterwards, we briefly recall some facts from lattice theory.

2.1 Blind Signatures

A blind signature scheme BS consists of three algorithms $(\text{Kg}, \text{Sign}, \text{Vf})$, where Sign is an interactive protocol between a signer \mathcal{S} and a user \mathcal{U} . The specification is as follows.

Key Generation. $\text{Kg}(1^n)$ outputs a private signing key sk and a public verification key pk .

Signature Protocol. $\text{Sign}(\text{sk}, M)$ describes the joint execution of \mathcal{S} and \mathcal{U} . The private output of \mathcal{S} is a view \mathcal{V} and the private output of \mathcal{U} is a signature \mathbf{s} on the message $M \in \mathcal{M}$ with message space \mathcal{M} under sk . Thus, we write $(\mathcal{V}, \mathbf{s}) \leftarrow \langle \mathcal{S}(\text{sk}), \mathcal{U}(\text{pk}, M) \rangle$.

Signature Verification. The algorithm $\text{Vf}(\text{pk}, \mathbf{s}, M)$ outputs 1 if \mathbf{s} is a valid signature on M under pk and otherwise 0.

Completeness is defined as with digital signature schemes, i.e., every honestly created signature for honestly created keys and for any messages $M \in \mathcal{M}$ has to be valid under this key. Views are interpreted as random variables, whose output is generated by subsequent executions of the respective protocol. Two views \mathcal{V}_1 and \mathcal{V}_2 are considered equal if they cannot be distinguished by any computationally unbounded algorithm with noticeable probability.

As for security, blind signatures have to satisfy two properties: blindness and one-more unforgeability [JLO97, PS00]. The notion of blindness is defined in an experiment $\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}$, where the adversarial signer \mathcal{S}^* works in three modes. In mode *find*, it chooses two messages M_0, M_1 and interacts with two users in mode *issue*. Depending on a coin flip b , the first (second) user obtains a blind signature for M_b (M_{1-b}). After seeing the unblinded signatures in the original order, with respect to M_0, M_1 , the signer has to guess the bit b in mode *guess*. If either of the user algorithms fails in outputting a valid signature, the signer is merely notified of the failure and does not get any signature. Below, we deal with aborts as an extension. Also note that we allow the adversary to keep a state that is fed back in subsequent calls. A scheme BS is (t, δ) -blind, if there is no adversary \mathcal{S}^* , running in time at most t , that wins the above experiment with advantage at least δ , where the advantage is defined as $\text{Adv}_{\mathcal{S}^*, \text{BS}}^{\text{blind}} = \left| \text{Prob} \left[\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}(n) = 1 \right] - \frac{1}{2} \right|$. A scheme is *statistically* blind if it is (∞, δ) -blind for a negligible δ . The second security property, one-more unforgeability, ensures that each completed interaction between signer and user yields at most one signature. It is formalized in the experiment $\text{Exp}_{\mathcal{U}^*, \text{BS}}^{\text{omf}}$, where an adversarial user tries to output j valid signatures after $\ell < j$ completed interactions with an honest signer. \mathcal{H} is a family of random oracles.

A signature scheme BS is $(t, q_{\text{Sign}}, q_{\text{H}}, \delta)$ -one-more unforgeable if there is no adversary \mathcal{A} , running in time at most t , making at most q_{Sign} signature queries and at most q_{H} hash oracle queries, that wins the above experiment with probability at least δ .

2.2 Extensions

We consider three extensions to the above security model for blind signatures: one deals with user aborts, the second with dishonestly chosen keys, and the third with leakage resilience.

Security Under Aborts. Blindness in the previous subsection does not cover the case where the protocol is aborted prematurely. There is the strengthened notion of selective failure blindness [CNS07], where the malicious signer may choose either M_0 or M_1 according to some secret distribution that makes the protocol fail. Preventing this generically is easy as was shown by Fischlin and Schröder in [FS09]. In the course of the discussion of our construction, we argue that it already is blind in this sense.

Adversely-chosen Keys. Consider the blindness experiment in [ANN06]. Instead of having the experiment select pk, sk , we can let the signer output pk .

Blindness may be harder to achieve in this setting. However, our construction remains blind in this stronger model as the proof does not exploit specifics about the key.

Leakage Resilience. Resilience to key leakage is a way to ensure security against side-channel attacks. In [KV09], Katz and Vaikuntanathan give a nice overview of past developments and the evolution of leakage resilience for authenticity and secrecy. Obviously, we are interested in authenticity in the special case of blind signatures. We model key leakage in the unforgeability experiment by adding a leakage oracle $\text{Leak}(\cdot)$ to $\text{Exp}_{\mathcal{U}^*, \text{BS}}^{\text{omf}}$. The adversary can adaptively query Leak with a series of functions f_i , $i \in \{1, \dots, \kappa\}$, and receives $f_i(\text{sk})$. The only restriction is that $\sum_{i=1}^{\kappa} |f_i(\text{sk})| < \lambda(|\text{sk}|)$, where the function λ determines the amount of leakage that we are willing to tolerate. Notice that the signer’s key does not have to evolve over time and its secret state consists of the secret key only. Furthermore, observe that this extension is only sensible as long as $\lambda(|\text{sk}|) < \min\{|\text{sk}|, |\text{s}|\}$, where $|\cdot|$ denotes bit-length and s is a signature. Otherwise, the adversary could easily obtain the entire secret key or a signature of its choice. See the full version [Rüc08] for the experiment. To demonstrate leakage resilience, one has to show that the conditional min-entropy $H_{\infty}(\text{sk} | \text{Leak}(\text{sk})) = \min_{\text{sk}'} \{-\log(\text{Prob}[\text{sk} = \text{sk}' | \text{Leak}(\text{sk})])\}$ of the secret key is still sufficiently large to prove security.

2.3 Commitments

Commitments typically work in two phases. First, one party publishes a commitment $C = \text{com}(M; r) \in \{0, 1\}^n$, $r \leftarrow_{\$} \{0, 1\}^n$, to a message $M \in \{0, 1\}^*$ without revealing any information about it. This is the “hiding” property of the commitment scheme. In the second phase, the party can prove that C actually corresponds to M by revealing r . It is important that no algorithm can find a second message M' and randomness r' such that $C = \text{com}(M'; r')$, i.e., break the “binding” property. As usual, these properties are defined for families of such commitment functions. A scheme is (t, δ) -hiding (-binding) if there is no algorithm running in time at most t that can break the hiding (binding) property with probability at least δ . Both properties can be satisfied computationally or unconditionally but there is no scheme that is unconditionally hiding *and* unconditionally binding [Gol04].

For our scheme, we assume a statistically $\delta_{\text{com}}^{(h)}$ -hiding and computationally $(t_{\text{com}}, \delta_{\text{com}}^{(b)})$ -binding commitment scheme. As we are interested in fully lattice-based schemes, we would like to point out that such commitment schemes can be built upon hard lattice problems [KTX08] but in practice, one rather uses cryptographic hash functions as a message authentication code. For example, using a lattice-based hash function [ADL⁺08].

2.4 Lattices

A lattice in \mathbb{R}^n is a discrete set $\Lambda = \{\sum_{i=1}^d x_i \mathbf{b}_i \mid x_i \in \mathbb{Z}\}$, where $\mathbf{b}_1, \dots, \mathbf{b}_d$ are linearly independent over \mathbb{R} . The matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_d]$ is a basis of the

lattice Λ and we write $\Lambda = \Lambda(\mathbf{B})$. The dimension of the lattice is d . The main computational problem in lattices is the shortest vector problem (SVP), where an algorithm is given a description, a basis, of a lattice Λ and is supposed to find the shortest vector $\mathbf{v} \in \Lambda \setminus \{\mathbf{0}\}$ with respect to a certain ℓ_p norm (up to an approximation factor). More precisely, find a vector $\mathbf{v} \in \Lambda \setminus \{\mathbf{0}\}$, such that $\|\mathbf{v}\|_p \leq \gamma \|\mathbf{w}\|_p$ for all $\mathbf{w} \in \Lambda \setminus \{\mathbf{0}\}$ for a fixed approximation factor $\gamma \geq 1$.

In this work, we are interested in a special family of lattices related to ideals in the ring $\mathbf{R} = \mathbb{Z}_q[X]/\langle \mathbf{g} \rangle$, where q is prime and $\mathbb{Z}_q = \{-(q-1)/2, \dots, (q-1)/2\}$. We focus on $\mathbf{g} = X^n + 1$ and $n = \text{“power of two”}$ for efficiency reasons but it may be replaced by any irreducible polynomial over \mathbb{Z} . Then, our scheme and the analysis become only slightly more involved. We identify $\mathbf{f} \in \mathbf{R}$ with its coefficient vector $\mathbf{f} = (f_0, \dots, f_{n-1}) \in \mathbb{Z}_q^n$. Furthermore, we denote elements of the \mathbf{R} -module \mathbf{R}^m with $\hat{\mathbf{a}} = (\mathbf{a}_0, \dots, \mathbf{a}_{m-1})$ or directly with $(a_0, \dots, a_{mn-1}) \in \mathbb{Z}_q^{mn}$. Consequently, we define $\|\mathbf{f}\|_\infty = \|(f_0, \dots, f_{n-1})\|_\infty$. The norm on \mathbf{R} is a slight abuse of notation, but it will only be used if \mathbf{f} has small coefficients over \mathbb{Z} . A lattice corresponds to an ideal $I \subset \mathbf{R}$ if and only if every lattice vector is the coefficient vector of a polynomial in I . The SVP problem easily translates to ideal lattices, where we call it ideal-SVP (ISVP).

The average-case hardness assumption for our construction relies on the problem of finding short vectors in the kernel of the family $\mathcal{H}(\mathbf{R}, m)$ of module homomorphisms $h_{\hat{\mathbf{a}} \in \mathbf{R}^m} : \mathbf{R}^m \rightarrow \mathbf{R}, \hat{\mathbf{x}} \mapsto \hat{\mathbf{a}} \otimes \hat{\mathbf{x}} = \sum_{i=0}^{m-1} \mathbf{a}_i \mathbf{x}_i$, when restricting the domain to $D' \subset \mathbf{R}$, i.e., restricting the coefficients in the input to $[-2d, 2d] \cap \mathbb{Z}$. This problem can be stated as the following collision problem [LM06].

Definition 1 (Collision Problem). *The collision problem $Col(\mathcal{H}(\mathbf{R}, m), D)$ asks to find a distinct pair $(\hat{\mathbf{x}}, \hat{\mathbf{x}}') \in D^m \times D^m$ such that $h(\hat{\mathbf{x}}) = h(\hat{\mathbf{x}}')$ for $h \leftarrow_{\S} \mathcal{H}(\mathbf{R}, m)$.*

Obviously, the function is linear over \mathbf{R}^m , i.e., $h(\mathbf{a}(\hat{\mathbf{x}} + \hat{\mathbf{y}})) = \mathbf{a}(h(\hat{\mathbf{x}}) + h(\hat{\mathbf{y}}))$ for all $\mathbf{a} \in \mathbf{R}, \hat{\mathbf{x}}, \hat{\mathbf{y}} \in \mathbf{R}^m$. In addition, solving $Col(\mathcal{H}(\mathbf{R}, m), D)$ implies being able to solve ISVP $^\infty$ in every lattice that corresponds to an ideal in \mathbf{R} by the following theorem.

Theorem 1 (Worst-case to Average-case, Theorem 2 in [LM06]). *Let $D = \{\mathbf{f} \in \mathbf{R} : \|\mathbf{f}\|_\infty \leq d\}$, $m > \log(q)/\log(2d)$, and $q \geq 4dmn\sqrt{n} \log(n)$. An adversary \mathcal{C} that solves the $Col(h, D)$ problem, i.e., finds distinct preimages $\hat{\mathbf{x}}, \hat{\mathbf{y}} \in D^m$ such that $h(\hat{\mathbf{x}}) = h(\hat{\mathbf{y}})$, can be used to solve ISVP $^\infty$ with approximation factors $\gamma \geq 16dmn \log^2(n)$ in the worst case.*

3 Blind Signatures from Ideal Lattices

We construct a lattice-based blind signature scheme. It is secure in the random oracle model under a worst-case assumption in ideal lattices and its time and space complexity is quasi-optimal, $\tilde{O}(n)$.

The road map for this section is as follows: We describe the 4-move blind signature scheme BS. Then, we prove completeness, blindness, and one-more unforgeability. Proving completeness is non-trivial as we need to address an

Table 2. Parameters for the security parameter n

Parameter	Value	Asymptotics	Usage
n	power of 2	-	main security parameter
d_s	positive integer constant $< q/(4n)$	$\mathcal{O}(1)$	secret key size, unforgeability
D_s	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq d_s\}$		set of secret keys
c_m	$> 1/\log(2d_s)$	$\tilde{\mathcal{O}}(1)$	witness indistinguishability, leakage resilience
m	$\lfloor c_m \log(q) \rfloor + 1$	$\Omega(\log(n))$	worst-case to average-case reduction
D_ϵ	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq 1 =: d_\epsilon\}$	$\mathcal{O}(1)$	hash output size
ϕ, ψ	positive integer constant ≥ 1	$\mathcal{O}(1)$	completeness, speed
D_α	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq \psi n d_\epsilon =: d_\alpha\}$	$\mathcal{O}(n)$	blindness
D_{ϵ^*}	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq d_\alpha - d_\epsilon =: d_{\epsilon^*}\}$	$\mathcal{O}(n)$	blindness
D_y	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq \phi m n^2 d_s d_{\epsilon^*} =: d_y\}$	$\tilde{\mathcal{O}}(n^3)$	witness indistinguishability
G_*	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq d_y - n d_s d_{\epsilon^*} =: d_{G_*}\}$	$\tilde{\mathcal{O}}(n^3)$	witness indistinguishability, completeness defect
D_β	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq \phi m n d_{G_*} =: d_\beta\}$	$\tilde{\mathcal{O}}(n^4)$	blindness
G	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq d_\beta - d_{G_*} =: d_G\}$	$\tilde{\mathcal{O}}(n^4)$	blindness, completeness defect
D	$\{\mathbf{f} \in \mathbf{R} : \ \mathbf{f}\ _\infty \leq d_{G_*} + d_\beta + n d_s d_\epsilon =: d_D\}$	$\tilde{\mathcal{O}}(n^4)$	collisions under h
q	$\geq 4mn\sqrt{n}\log(n)d_D$, prime	$\tilde{\Theta}(n^5\sqrt{n})$	worst-case to average-case reduction

The table defines all parameters and sets for our scheme. The sets are defined via a norm bound, for which we also state the asymptotic growth with respect to the security parameter n . The last column states the main usage for the individual parameter or set. Some sets introduce a completeness error to the scheme that can be reduced by increasing ϕ . Reducing this defect also significantly improves performance. All sets are subsets of the ring $\mathbf{R} = \mathbb{Z}_q[X]/(X^n + 1)$.

inevitable completeness defect. In the course of the discussion we show that it neither harms security nor efficiency. Afterwards, we prove that the scheme is statistically blind and that it is one-more unforgeable unless the collision problem $Col(\mathcal{H}(\mathbf{R}, m), D)$ is easy. In consequence, one-more unforgeability can be based on the worst-case hardness of the ISVP. After the main analysis, we prove that our scheme also supports leakage resilience.

Observe that the scheme requires lots of parameters that need to be carefully worked out. Their definition in Table 2 will be justified later in the analysis. We chose not to “unwind” the parameters d_s, d_ϵ , etc. because we need their relative size in the various lemmas below, making the proofs easier to understand. The asymptotics in the third column should help estimating their magnitude. The parameter d_ϵ is a constant 1 here but it can be increased if it is necessary to sign hash values of bit length $> n \log_2(3)$. The “usage” hint in the table points at the section, where they are most influential. As for selecting practical parameters, we refer the reader to the full version [Rüc08]. There, we propose secure parameter sets based on the analysis in [RS10]. The full version also includes a discussion on possible trade-offs for efficiency.

3.1 Our Construction

We construct our blind signature scheme $BS = (\text{Kg}, \text{Sign}, \text{Vf})$ as follows.




Fig. 1. Issue protocol of the blind signature scheme BS. All parameters and sets are defined in Table 2. Note that the signer implicitly verifies that the user’s protocol messages come from the correct domains.

Key Generation. $\text{BS.Kg}(1^n)$ selects a secret key $\hat{s} \leftarrow_{\mathcal{S}} D_s^m$, and a compression function $h \leftarrow_{\mathcal{S}} \mathcal{H}(\mathbf{R}, m)$. Let $\mathcal{C}(1^n)$ be a commitment scheme, mapping $\{0, 1\}^* \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. The algorithm chooses a function $\text{com} \leftarrow_{\mathcal{S}} \mathcal{C}(1^n)$ and, in addition, selects $\mathbf{H} \leftarrow_{\mathcal{S}} \mathcal{H}(1^n)$ mapping $\{0, 1\}^* \rightarrow D_\epsilon \subset D$.

Then, it computes the public key $\mathbf{S} \leftarrow h(\hat{s})$ and outputs (\hat{s}, \mathbf{S}) . For simplicity, we treat h , com , \mathbf{H} , and the parameters in Table 2 as globally known and implicit inputs to all algorithms. However, each signer may choose them individually and include them in the public key.

Signature Protocol. The signature issue protocol for messages $M \in \{0, 1\}^*$ is depicted in Figure 1. Eventually, the user outputs a message M and a signature (r, \hat{z}, ϵ) .

Notes: Upon a restart after Step 2, the user only selects a fresh $\alpha \leftarrow_{\mathcal{S}} D_\alpha$ and repeats the operations that involve α . Whenever the signer triggers a restart, the user chooses a fresh r in order to make the protocol execution independent of the previous ones. Therefore, we omit values from previous runs in the signer’s view. During Step 5, the signer can detect a cheating user that tries to trigger a restart, despite having received a valid signature. In this case, the signer can stop the protocol and assume that the user has obtained a valid signature.

Verification. $\text{BS.Vf}(\mathbf{S}, (r, \hat{z}, \epsilon), M)$ outputs 1 iff $\hat{z} \in G^m$ and $\mathbf{H}(h(\hat{z}) - \mathbf{S}\epsilon, \text{com}(M; r)) = \epsilon$.

3.2 Analysis and Security

In this section, we analyze our blind signature scheme with regard to completeness, blindness, one-more unforgeability, and leakage resilience. For each aspect, we prove a main theorem. Supporting lemmas are stated before the theorems and proven in the full version [Rüc08].

Completeness. Completeness of BS is a non-trivial issue due to the eventual restarts and the many parameters involved. The next lemma ensures that the number of restarts is small, effectively constant.

Lemma 1. *Let $k = \Omega(n)$, $\mathbf{a}, \mathbf{b} \in \mathbb{Z}^k$ with arbitrary $\mathbf{a} \in \{\mathbf{v} \in \mathbb{Z}^k : \|\mathbf{v}\|_\infty \leq A\}$ and random $\mathbf{b} \leftarrow_{\mathfrak{S}} \{\mathbf{v} \in \mathbb{Z}^k : \|\mathbf{v}\|_\infty \leq B\}$. Given $B \geq \phi k A$ for $\phi \in \mathbb{N}_{>0}$, we have $\text{Prob}_{\mathbf{b}}[\|\mathbf{a} - \mathbf{b}\|_\infty \leq B - A] > \frac{1}{e^{1/\phi}} - o(1)$.*

Theorem 2 (Completeness). *Let $g(n) = \omega(\log^2(n))$. The scheme BS is complete after at most $g(n)$ (or, an expected number of $e^{2/\phi}$) repetitions.*

See the full version [Rüc08] for the proof. There, we also argue that $\phi = 4$ is good choice to make the protocol more efficient in practice. Observe that in any case, all operations (including eventual restarts) in BS have $\tilde{O}(n)$ complexity and that private keys, public keys, and signatures have size $\tilde{O}(n)$.

Blindness. We prove that BS is statistically blind based on the observation that the signer only sees values that are independent of the message being signed. More precisely, the views generated by two different messages are indistinguishable. For this argument to work, we require a statistically hiding commitment scheme and carefully selected sets $D_\alpha, D_\beta, D_{\epsilon^*}$, and G . The following probabilistic lemma is crucial as it guarantees that the user’s message after Step 2 and the final output are independent of the message. In the context of $\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}$, this establishes a form of witness indistinguishability w.r.t. the messages that are chosen by the malicious signer.

Lemma 2. *Let $k \in \mathbb{N}$, $\mathbf{a}, \mathbf{a}', \mathbf{b} \in \mathbb{Z}^k$ with arbitrary $\mathbf{a}, \mathbf{a}' \in \{\mathbf{v} \in \mathbb{Z}^k : \|\mathbf{v}\|_\infty \leq A\}$, a random $\mathbf{b} \leftarrow_{\mathfrak{S}} \{\mathbf{v} \in \mathbb{Z}^k : \|\mathbf{v}\|_\infty \leq B\}$ for $B > A$. We define the random variables $\mathbf{c} \leftarrow \mathbf{a} - \mathbf{b}$ and $\mathbf{c}' \leftarrow \mathbf{a}' - \mathbf{b}$ if $\max\{\|\mathbf{a} - \mathbf{b}\|_\infty, \|\mathbf{a}' - \mathbf{b}\|_\infty\} \leq B - A$, otherwise, we resample \mathbf{b} . Then, $\Delta(\mathbf{c}, \mathbf{c}') = 0$.*

The role of `com` is to ensure that the signer can only obtain negligible information from restarts. Notice that BS is perfectly blind ($(\infty, 0)$ -blind) if the commitment scheme is perfect (0-hiding).

Theorem 3 (Blindness). *BS is $(\infty, \delta_{\text{com}}^{(h)})$ -blind if `com` is $\delta_{\text{com}}^{(h)}$ -hiding.*

Proof. As per experiment $\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}$, the adversarial signer outputs two messages M_0, M_1 and interacts with two users $\mathcal{U}(\mathbf{S}, M_b), \mathcal{U}(\mathbf{S}, M_{1-b})$ after a secret coin flip $b \leftarrow \{0, 1\}$. We show that these users do not leak any information about their respective message.

Technically, we establish that all protocol messages and the output, when interpreted as random variables, are distributed independently of the message being signed. This involves an analysis of ϵ^* , $\hat{\mathbf{z}}$, and eventual restarts. As for ϵ and r we need not worry. They are chosen uniformly at random.

Distribution of ϵ^* . Let $\epsilon_b^*, \epsilon_{1-b}^*$ be the first protocol messages of $\mathcal{U}(\text{pk}, M_b)$ resp. $\mathcal{U}(\text{pk}, M_{1-b})$. They are in D_{ϵ^*} and they are both of the form $\epsilon - \alpha$ with $\epsilon \in D_\epsilon$ and $\alpha \leftarrow_{\mathcal{S}} D_\alpha$. The statistical distance $\Delta(\epsilon_b^*, \epsilon_{1-b}^*)$ is 0 by Lemma 2 ($k = n, A = d_s, B = d_\alpha$) because the coefficients in D_{ϵ^*} are bounded by $B - A = d_\alpha - d_s$.

Distribution of $\hat{\mathbf{z}}$. Let $\hat{\mathbf{z}}_0, \hat{\mathbf{z}}_1$ be part of the final output of $\mathcal{U}(\text{pk}, M_0)$ resp. $\mathcal{U}(\text{pk}, M_1)$. Both are of the form $\hat{\mathbf{z}}^* - \hat{\beta}$ for $\hat{\mathbf{z}}^* \in G_*^m$ and $\hat{\beta} \leftarrow_{\mathcal{S}} D_\beta^m$. Furthermore, $\hat{\mathbf{z}}_0$ and $\hat{\mathbf{z}}_1$ are forced to be in G^m , having coefficients bounded by $d_\beta - d_{G_*}$. Hence, the statistical distance $\Delta(\hat{\mathbf{z}}_0, \hat{\mathbf{z}}_1)$ is 0 because of Lemma 2 ($k = mn, A = d_{G_*}, B = d_\beta$).

Restarts. Observe that each protocol run is statistically independent of the previous runs by the statistical hiding property of the commitment `com` and because the user selects fresh r, α, β after every restart. This is the reason why we inherit the statistical $\delta_{\text{com}}^{(h)}$ -hiding property to obtain $(\infty, \delta_{\text{com}}^{(h)})$ -blindness instead of perfect blindness. Finally, we need to argue about the restart after Step 4. The user sends $(C, \alpha, \beta, \epsilon)$ to the signer. These information allow the verification of the signature with respect to C . The message is still statistically hidden by the hiding property of `com` because the user never reveals the decommitment r .

Hence, the protocol hides the to-be-signed message and subsequent runs of the protocol for the same message are statistically independent. \square

Furthermore, our scheme already supports selective failure blindness as shown in [FS09] because we are signing commitments instead of the adversely chosen messages. Even the fourth move does not reveal any information about the message due to the hiding property of the commitment.

One-more Unforgeability. In this section, we show that `BS` is one-more unforgeable, provided that the collision problem $\text{Col}(\mathcal{H}(\mathbf{R}, m), D)$ is hard and the commitment scheme is binding. The main tool in the reduction is the Forking Lemma [PS00, BN06]. To simulate the environment, especially blind signature queries, for the attacker \mathcal{A} in the unforgeability experiment, we require that there are at least two possible secret keys for each public key \mathbf{S} (Lemma 3). Moreover, we need the signature protocol to be witness indistinguishable to prevent the attacker from learning the secret key (Lemma 4). The binding property of `com` is necessary to prevent an attacker from obtaining one signature that works for two messages by changing the message under the commitment. All other attackers output at least one signature that does not correspond to a completed interaction. Here, we apply the Forking Lemma to extract knowledge about the secret key that was used to compute the forgery. Using this knowledge the reduction can solve the collision problem. Finally, we need to deal with Step 5 in the

protocol. The adversary proves that it was unable to obtain a valid signature. We show that this is sufficient if Col is hard.

Since the function family $\mathcal{H}(\mathbf{R}, m)$ compresses the domain D_s^m , it is easy to show that all secret keys collide with at least one other secret key.

Lemma 3. *Let $h \in \mathcal{H}(\mathbf{R}, m)$. For every secret key $\hat{s} \leftarrow_{\S} D_s^m$, there is a second $\hat{s}' \in D_s^m \setminus \{\hat{s}\}$ with $h(\hat{s}) = h(\hat{s}')$ (with overwhelming probability).*

The next lemma establishes witness indistinguishability of the protocol. Witness indistinguishability ensures that the malicious verifier cannot distinguish whether the prover uses one of two possible secret keys $\hat{s}, \hat{s}' \in h^{-1}(\mathbf{S}) \cap D_s^m$. Basically, it can be interpreted as an application of Lemma 2 to $\hat{\mathbf{z}}^* = (\hat{s}\epsilon^*) + \hat{\mathbf{y}} \in G_*^m$ with some further observations. The choice of $\hat{\mathbf{y}} \leftarrow_{\S} D_y$ and the restriction “ $\in G_*^m$ ” hide the first summand.

Lemma 4. *Let $h \in \mathcal{H}(\mathbf{R}, m)$ and $\mathbf{S} \in \mathbf{R}$. For any message M and any two secret keys $\hat{s}, \hat{s}' \in D_s^m$ with $h(\hat{s}) = \mathbf{S} = h(\hat{s}')$, the resulting protocol views $(\mathbf{Y}, \epsilon^*, \hat{\mathbf{z}}^*)$ and $(\mathbf{Y}', \epsilon'^*, \hat{\mathbf{z}}'^*)$ are indistinguishable.*

Using lemmas 3 and 4, we can exploit witness indistinguishability to simulate all blind signature oracle queries with a secret key \hat{s} and at the same time expect the adversary to output a forgery that corresponds to a different secret key \hat{s}' with non-negligible probability or break the binding property of the commitment scheme. We apply the Forking Lemma to extract a solution to the $Col(\mathcal{H}(\mathbf{R}, m), D)$.

Theorem 4 (One-more unforgeability). *Let Sig be the signature oracle. Let T_{Sig} and T_H be the cost functions for simulating the oracles Sig and H , and let $c < 1$ be the probability for a restart in the protocol. BS is $(t, q_{Sig}, q_H, \delta)$ -one-more unforgeable if com is $(t', \delta/2)$ -binding and $Col(\mathcal{H}(\mathbf{R}, m), D)$ is $(t', \delta'/2)$ -hard with $t' = t + q_H^{q_{Sig}}(q_{Sig}T_{Sig} + q_H T_H)$ and non-negligible δ' if δ is non-negligible.*

The probability δ' depends on the number of issued signatures. It can be found at the end of the proof.

Proof. Towards contradiction, we assume that there exists a successful forger \mathcal{A} against one-more unforgeability of BS with non-negligible probability δ . Using \mathcal{A} , we construct an algorithm \mathcal{B} , such that it either solves the collision problem or breaks the binding property of com .

Setup. \mathcal{B} flips a coin $b \leftarrow_{\S} \{0, 1\}$. For $b = 0$, it selects $h \leftarrow_{\S} \mathcal{H}(\mathbf{R}, m)$. For $b = 1$, it gets the description of h as input. \mathcal{B} initializes a list $L_H \leftarrow \emptyset$ of query-hash pairs $(\mathbf{R} \times \{0, 1\}^*, D_\epsilon)$. It chooses $\hat{s} \leftarrow_{\S} D_s^m$ and sets $\mathbf{S} \leftarrow h(\hat{s})$. Furthermore, it randomly pre-selects random oracle answers $\mathbf{h}_1, \dots, \mathbf{h}_{q_H} \leftarrow_{\S} D_\epsilon$ and a random tape ρ . It runs $\mathcal{A}(\mathbf{S}; \rho)$ in a black-box simulation.

Random Oracle Queries. On input (\mathbf{u}, C) , \mathcal{B} looks up (\mathbf{u}, C) in L_H . If it finds corresponding hash value ϵ then it returns ϵ . Otherwise, \mathcal{B} selects the first unused ϵ from the list $\mathbf{h}_1, \dots, \mathbf{h}_{q_H}$, stores $((\mathbf{u}, C), \epsilon)$ in L_H , and returns ϵ .

Blind Signature Queries. \mathcal{B} acts according to the protocol in Figure 1.

Output. Eventually, \mathcal{A} stops and outputs $(M_1, (r_1, \hat{\mathbf{z}}_1, \epsilon_1), \dots, (M_j, (r_j, \hat{\mathbf{z}}_j, \epsilon_j)), q_{\text{Sign}} + 1 = j$, for distinct messages. If $b = 0$, the reduction looks for two pairs $(M_1^*, (r_1^*, \hat{\mathbf{z}}^*, \epsilon^*))$ and $(M_2^* \neq M_1^*, (r_2^*, \hat{\mathbf{z}}^*, \epsilon^*))$ and outputs $(M_1^*, r_1^*), (M_2^*, r_2^*)$ to break the binding property of com . If there is no such collision, \mathcal{B} aborts. If $b = 1$, the simulator \mathcal{B} guesses an index $k \leftarrow_{\mathcal{S}} [j]$ such that $h_i = \epsilon_k$ for some $i \in [q_{\text{H}}]$. Then, \mathcal{B} starts over, running $\mathcal{A}(\mathbf{S}; \rho)$ with random oracle answers $\mathbf{h}_1, \dots, \mathbf{h}_{i-1}, \mathbf{h}'_i, \dots, \mathbf{h}'_{q_{\text{H}}}$ for a fresh set $\mathbf{h}'_i, \dots, \mathbf{h}'_{q_{\text{H}}} \leftarrow_{\mathcal{S}} D_\epsilon$. Both \mathcal{A} and \mathcal{B} are run with the same random tape as in the first run. Among other values, \mathcal{A} outputs $(M'_k, (r'_k, \hat{\mathbf{z}}'_k, \epsilon'_k))$ and \mathcal{B} returns $(\hat{\mathbf{z}}_k - \hat{\mathbf{s}}\epsilon_k, \hat{\mathbf{z}}'_k - \hat{\mathbf{s}}\epsilon'_k)$ if $\epsilon'_k = \epsilon_k$ in an attempt to solve $\text{Col}(\mathcal{H}(\mathbf{R}, m), D)$. If $\epsilon'_k \neq \epsilon_k$, the reduction retries at most q_{H}^2 times with a different random tape and random oracle.

Analysis. \mathcal{A} 's environment is perfectly simulated. Especially, restarts happen with the same probability as in the original protocol. For $b = 0$, \mathcal{B} $(t', \delta/2)$ -breaks the binding property of com if \mathcal{A} breaks the binding property of com to break one-more unforgeability.

For $b = 1$, we assume that \mathcal{A} breaks one-more unforgeability without attacking com . So, at least one of the output signatures is not obtained via an interaction. The probability that \mathcal{B} guesses the index k of this signature correctly is at least $1/(q_{\text{Sign}} + 1)$. Observe that ϵ_k is a random oracle answer but with probability $1/|D_\epsilon|$. Furthermore, notice that with probability $1/2$, at least one of the re-runs of \mathcal{A} yields the same map $\{(i, k) : h_i = \epsilon_k\}$ as in the first run of \mathcal{A} . Thus, we consider the indices in both “interesting” replays to be constant.

Applying the Forking Lemma, we know that with probability $\delta_{\text{frk}} \geq (1 - c)(\delta - 1/|D_\epsilon|)((\delta - 1/|D_\epsilon|)/q_{\text{H}} - 1/|D_\epsilon|)$, \mathcal{A} is again successful in the one-more unforgeability experiment and outputs $(M'_k, (r'_k, \hat{\mathbf{z}}'_k, \epsilon'_k))$ using the *same random oracle query* as in the first run. The additional $(1 - c)$ factor takes a potential abort during the second run into account, which happen with probability at most c . Therefore, we know that $(h(\hat{\mathbf{z}}_k - \mathbf{S}\epsilon_k), \text{com}(M_k; r_k)) = (h(\hat{\mathbf{z}}'_k - \mathbf{S}\epsilon'_k), \text{com}(M'_k; r'_k))$.

Now, we turn to solving the collision problem. We have to show that $\hat{\mathbf{z}}_k - \hat{\mathbf{s}}\epsilon_k \neq \hat{\mathbf{z}}'_k - \hat{\mathbf{s}}\epsilon'_k$ and $h(\hat{\mathbf{z}}_k - \hat{\mathbf{s}}\epsilon_k) = h(\hat{\mathbf{z}}'_k - \hat{\mathbf{s}}\epsilon'_k)$. The second requirement follows directly from the previous paragraph. The first is more involved. Here, it is important that the protocol is witness indistinguishable (Lemma 4), i.e., the adversary does not recognize whether we have used one of at least two possible $\hat{\mathbf{s}}, \hat{\mathbf{s}}'$ (Lemma 3 with probability greater than $1/2$). Thus, with probability at least $1/2$ its output corresponds to $\hat{\mathbf{s}}'$. We show that either $\hat{\mathbf{z}}_k - \hat{\mathbf{s}}\epsilon_k \neq \hat{\mathbf{z}}'_k - \hat{\mathbf{s}}\epsilon'_k$ or $\hat{\mathbf{z}}_k - \hat{\mathbf{s}}'\epsilon_k \neq \hat{\mathbf{z}}'_k - \hat{\mathbf{s}}'\epsilon'_k$. Assuming both are equal, we subtract the equations and obtain $(\epsilon_k - \epsilon'_k)(\hat{\mathbf{s}}' - \hat{\mathbf{s}}) = \mathbf{0}$. We know that $\epsilon_k - \epsilon'_k \neq \mathbf{0}$. Now, $\|(\epsilon_k - \epsilon'_k)(\hat{\mathbf{s}}' - \hat{\mathbf{s}})\|_\infty \leq 2d_s n < q/2$ because $\|\epsilon_k - \epsilon'_k\|_\infty \leq 2$ and $\|\hat{\mathbf{s}}' - \hat{\mathbf{s}}\|_\infty \leq 2d_s$. Thus, $(\epsilon_k - \epsilon'_k)(\hat{\mathbf{s}}' - \hat{\mathbf{s}}) = \mathbf{0}$ over $\mathbb{Z}[X]/\langle X^n + 1 \rangle$, which is an integral domain. So, we have the contradiction $\hat{\mathbf{s}}' = \hat{\mathbf{s}}$ and a collision $(\hat{\mathbf{z}}_k - \hat{\mathbf{s}}\epsilon_k, \hat{\mathbf{z}}'_k - \hat{\mathbf{s}}\epsilon'_k) \in D \times D$. The success probability is at least $\delta_{\text{col}} \geq 1/4 \delta_{\text{frk}}/(q_{\text{Sign}} + 1)$, which is non-negligible if δ is non-negligible.

Concerning restarts, we argue that the user cannot obtain a valid signature out of an aborted interaction without solving the collision problem. In order to

trigger an abort after Step 4, it outputs $result = (C, \alpha, \hat{\beta}, \epsilon)$ which, together with $\hat{\mathbf{z}}^*, \hat{\mathbf{y}}, \epsilon^*$, satisfies all abort criteria:

$$\epsilon^* + \alpha = \epsilon = H(\mathbf{Y} - \mathbf{S}\alpha - h(\hat{\beta}), C) \tag{1}$$

$$\epsilon = H(h(\hat{\mathbf{z}}^* - \hat{\beta}) - \mathbf{S}\epsilon, C) \tag{2}$$

$$\hat{\mathbf{z}}^* - \hat{\beta} \notin G^m \tag{3}$$

Assume that it also obtains a valid signature $(r', \hat{\mathbf{z}}', \epsilon')$ from this interaction. If $\epsilon = \epsilon'$, then $h(\hat{\mathbf{z}}^* - \hat{\beta}^* - \hat{\mathbf{s}}\epsilon) = h(\hat{\mathbf{z}}' - \hat{\mathbf{s}}\epsilon)$ by (2). If the arguments under h are equal, we have $\hat{\mathbf{z}}^* - \hat{\beta} \in G^m$ — a contradiction with (3). If the arguments are distinct, we have a collision in D because $\|\hat{\mathbf{z}}' - \hat{\mathbf{s}}\epsilon\|_\infty \leq d_G < d_D$ and $\|\hat{\mathbf{z}}^* - \hat{\beta}^* - \hat{\mathbf{s}}\epsilon\|_\infty \leq d_{G_*} + d_\beta + nd_s d_\epsilon = d_D$.

The adversary may succeed by hiding $\epsilon' \neq \epsilon$ in ϵ^* . But then, we necessarily have $\epsilon^* = \epsilon - \alpha = \epsilon' - \alpha'$ by (1) for an $\alpha \neq \alpha'$ and we know that $\alpha = \epsilon - \epsilon' + \alpha'$. So, the adversary had to be able to predict the output of H to compute α .

To conclude, the probability that we can extract a collision from a cheating user during an abort is at least $\delta_{\text{abort}} \geq \delta(1 - 1/|D_\epsilon|)$, which is non-negligible if δ is non-negligible. Thus, the overall success probability of the reduction is $\delta' \geq \min(\delta_{\text{col}}, \delta_{\text{abort}})$ if the guess $b = 1$ was correct. \square

Hence, we require that $q_{\text{Sig}} = o(n)$ to be able to rely on the subexponential hardness of lattice problems. This constraint is an artifact of the proof technique as discussed in [PS00] and it is not at all unusual for efficient blind signature schemes. There, it was even required that $q_{\text{Sig}} \leq (\log(n))^{\mathcal{O}(1)}$ because they needed a polynomial-time reduction. In consequence, in our reduction, we greatly benefit from the subexponential hardness of the underlying lattice problem. Alternatively, we believe that the running time of the reduction can be significantly reduced to being polynomial in q_{Sig} by using techniques due to Pointcheval [Poi98].

By Theorem 1, we get the following strong worst-case security guarantees.

Corollary 1. *BS is one-more unforgeable if solving ISVP $^\infty$ is hard in the worst case for approximation factors $\gamma \geq 16d_D m n \log^2(n) = \tilde{\mathcal{O}}(n^5)$ in lattices that correspond to ideals in \mathbf{R} .*

Leakage Resilience. Using an additional restriction for one of the parameters, we can safely leak a $(1 - o(1))$ fraction of the secret key in the unforgeability experiment according to the definition in the full version [Rüc08]. Recall that $m = \lfloor c_m \log(q) \rfloor + 1$ for some $c_m = \tilde{\mathcal{O}}(1)$. Thus, it is possible to choose c_m , say $\log(n)$, without loosing the scheme’s quasi-optimal efficiency. The following theorem states that such a choice is sufficient to provide strong leakage resilience. The proof can be found in the full version [Rüc08].

Theorem 5 (Leakage Resilience). *Let $c_m = \omega(1)$ and let $L := \log(|D_s^m|) = mn \log(2d_s + 1)$ be the length of the secret key. The conditional min-entropy H_∞ of $\hat{\mathbf{s}}$, conditioned on $\mathbf{S} = h(\hat{\mathbf{s}})$ and a total secret-key leakage $f(\hat{\mathbf{s}})$ of $\lambda = \delta L = (1 - o(1))L$ bits, is positive with overwhelming probability.*

4 Conclusions

We have shown how to construct an efficient and provably secure blind signature scheme based on the hardness of worst-case lattice problems. Our scheme has four moves, offers quasi-optimal performance, and it is leakage resilient in an almost optimal sense. Therefore, we expect our construction to withstand even subexponential-time and quantum computer attacks, as well as limited side-channel attacks against the secret key.

Acknowledgments

The author thanks Özgür Dagdelen, Marc Fischlin, Tibor Jager, Vadim Lyubashevsky, Chris Peikert, Michael Schneider, and Dominique Schröder for reviewing parts of this work and for very helpful and encouraging discussions. He also thanks the anonymous reviewers of ASIACRYPT 2010 for their valuable input.

References

- [Abe01] Abe, M.: A secure three-move blind signature scheme for polynomially many signatures. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 136–151. Springer, Heidelberg (2001)
- [ADL⁺08] Arbitman, Y., Dogon, G., Lyubashevsky, V., Micciancio, D., Peikert, C., Rosen, A.: SWIFFTX: A proposal for the SHA-3 standard. In: The First SHA-3 Candidate Conference (2008)
- [ADW09] Alwen, J., Dodis, Y., Wichs, D.: Leakage-resilient public-key cryptography in the bounded-retrieval model. In: Halevi, S. (ed.) Advances in Cryptology - CRYPTO 2009. LNCS, vol. 5677, pp. 36–54. Springer, Heidelberg (2009)
- [Ajt96] Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: STOC, pp. 99–108. ACM, New York (1996)
- [AKS01] Ajtai, M., Kumar, R., Sivakumar, D.: A sieve algorithm for the shortest lattice vector problem. In: STOC, pp. 601–610. ACM, New York (2001)
- [ANN06] Abdalla, M., Namprempre, C., Neven, G.: On the (im)possibility of blind message authentication codes. In: Pointcheval, D. (ed.) CT-RSA 2006. LNCS, vol. 3860, pp. 262–279. Springer, Heidelberg (2006)
- [BMV08] Bresson, E., Monnerat, J., Vergnaud, D.: Separation results on the "one-more" computational problems. In: Malkin, T. (ed.) CT-RSA 2008. LNCS, vol. 4964, pp. 71–87. Springer, Heidelberg (2008)
- [BN06] Bellare, M., Neven, G.: Multi-signatures in the plain public-key model and a general forking lemma. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM Conference on Computer and Communications Security, pp. 390–399. ACM, New York (2006)
- [BNPS03] Bellare, M., Namprempre, C., Pointcheval, D., Semanko, M.: The one-more-rsa-inversion problems and the security of chaum's blind signature scheme. *J. Cryptology* 16(3), 185–215 (2003)
- [Bol03] Boldyreva, A.: Threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In: Desmedt, Y. (ed.) PKC 2003. LNCS, vol. 2567, pp. 31–46. Springer, Heidelberg (2002)

- [BR93] Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: CCS. ACM, New York (1993)
- [Cha82] Chaum, D.: Blind signatures for untraceable payments. In: CRYPTO, pp. 199–203 (1982)
- [CKW04] Camenisch, J., Koprowski, M., Warinschi, B.: Efficient blind signatures without random oracles. In: Blundo, C., Cimato, S. (eds.) SCN 2004. LNCS, vol. 3352, pp. 134–148. Springer, Heidelberg (2005)
- [CNS07] Camenisch, J., Neven, G., Shelat, A.: Simulatable adaptive oblivious transfer. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 573–590. Springer, Heidelberg (2007)
- [ECR10] ECRYPT2. Yearly report on algorithms and key sizes — report D.SPA.13 (2010), <http://www.ecrypt.eu.org/documents/D.SPA.13.pdf>
- [FS86] 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)
- [FS09] Fischlin, M., Schröder, D.: Security of blind signatures under aborts. In: Jarecki, S., Tsudik, G. (eds.) PKC 2009. LNCS, vol. 5443, pp. 297–316. Springer, Heidelberg (2009)
- [Gol04] Goldreich, O.: The Foundations of Cryptography, vol. 1. Cambridge University Press, Cambridge (2004)
- [GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) STOC, pp. 197–206. ACM, New York (2008)
- [JLO97] Juels, A., Luby, M., Ostrovsky, R.: Security of blind digital signatures (extended abstract). In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 150–164. Springer, Heidelberg (1997)
- [KTX08] Kawachi, A., Tanaka, K., Xagawa, K.: Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 372–389. Springer, Heidelberg (2008)
- [KV09] Katz, J., Vaikuntanathan, V.: Signature schemes with bounded leakage resilience. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 703–720. Springer, Heidelberg (2009)
- [Len05] Lenstra, A.: The Handbook of Information Security. Key Lengths, ch. 14. Wiley, Chichester (2005), http://www.keylength.com/biblio/Handbook_of_Information_Security_-_Keylength.pdf
- [LM06] Lyubashevsky, V., Micciancio, D.: Generalized compact knapsacks are collision resistant. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 144–155. Springer, Heidelberg (2006)
- [Lyu08] Lyubashevsky, V.: Lattice-based identification schemes secure under active attacks. In: Cramer, R. (ed.) PKC 2008. LNCS, vol. 4939, pp. 162–179. Springer, Heidelberg (2008)
- [Lyu09] 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)
- [Mat09] Matsui, M. (ed.): ASIACRYPT 2009. LNCS, vol. 5912. Springer, Heidelberg (2009)

- [Oka92] Okamoto, T.: Provably secure and practical identification schemes and corresponding signature schemes. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 31–53. Springer, Heidelberg (1993)
- [Oka06] Okamoto, T.: Efficient blind and partially blind signatures without random oracles. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 80–99. Springer, Heidelberg (2006)
- [Poi98] Pointcheval, D.: Strengthened security for blind signatures. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 391–405. Springer, Heidelberg (1998)
- [PS97] Pointcheval, D., Stern, J.: New blind signatures equivalent to factorization (extended abstract). In: ACM Conference on Computer and Communications Security, pp. 92–99 (1997)
- [PS00] Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. *J. Cryptology* 13(3), 361–396 (2000)
- [RHOAGZ07] Rodríguez-Henríquez, F., Ortiz-Arroyo, D., García-Zamora, C.: Yet another improvement over the mu-varadharajan e-voting protocol. *Comput. Stand. Interfaces* 29(4), 471–480 (2007)
- [RS10] Rückert, M., Schneider, M.: Selecting secure parameters for lattice-based cryptography. *Cryptology ePrint Archive*, Report 2010/137 (2010), <http://eprint.iacr.org/>
- [Rüc08] Rückert, M.: Lattice-based blind signatures. *Cryptology ePrint Archive*, Report 2008/322 (2008), <http://eprint.iacr.org/>
- [Rüc10] Rückert, M.: Adaptively secure identity-based identification from lattices without random oracles. In: Garay, J.A., De Prisco, R. (eds.) SCN 2010. LNCS, vol. 6280, pp. 345–362. Springer, Heidelberg (2010), <http://dblp.uni-trier.de/rec/bibtex/conf/scn/Ruckert10>
- [Sch91] Schnorr, C.P.: Efficient signature generation by smart cards. *J. Cryptology* 4, 161–174 (1991)
- [Sho97] Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* 26(5), 1484–1509 (1997)