

Efficient Quantum-immune Blind Signatures

— preliminary version —

Markus Rückert
rueckert@cdc.informatik.tu-darmstadt.de

Cryptography and Computeralgebra
Department of Computer Science
TU Darmstadt

October 28, 2008

Keywords Quantum-immune, blind signatures, lattice-based cryptography

Abstract We present the first quantum-immune blind signature scheme. Our scheme is provably secure in the random oracle model, efficient, and round-optimal. The basis of security is a new computational problem related to finding short lattice vectors. The underlying signature scheme is stateful and its basis of security is the problem of finding short vectors in a lattice.

1 Introduction

Since 1982, when David Chaum proposed his idea of blind signatures and a, by now classic, application in the context of digital payments, numerous blind signature schemes and other privacy-enhanced signature schemes have been developed. According to the security model, mainly influenced by Juels, Luby, and Ostrovsky [20] as well as Pointcheval and Stern [32], blind signature schemes have 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 enforces that an adversarial user cannot obtain more signatures than there were interactions with the signer.

Today, when building provably secure signature schemes, one has to keep emerging technologies and especially quantum computers in mind. In the quantum-age, the cryptographic assumptions change with the leap in computing power that quantum computers will provide.

To date, there are only a few cryptographic assumptions that are conjectured to be *quantum-immune*, i.e. they are considered to be able to withstand quantum computer attacks. One of those assumptions is the hardness of approximating shortest vectors in a lattice. Although the works of Ludwig [26] and Regev [33] suggest that today’s lattice reduction algorithms can benefit from the intrinsic parallelism in quantum computation, this does not invalidate the assumption. Slightly larger security parameters appear to be a sufficient countermeasure.

Our Contribution and related work. Using a problem that is related — yet not equivalent — to the problem of approximating the shortest vectors in a lattice (SVP) as our security assumption, we construct the first quantum-immune blind signature scheme. In order to prove unforgeability of our blind signature scheme, we introduce a new computational problem similar to the “one-more” problems [8] in the RSA context. As for its efficiency, we state that it is almost as efficient as the underlying signature scheme proposed by Gentry, Peikert, and Vaikuntanathan (GPV) [17] and with its two rounds, it is even *round-optimal*. Note that the GPV signature scheme is stateful and

that the stateless modification in [17] cannot be directly applied in the context of blind signatures, which is why we stick to the stateful variant.

The security of both, GPV signature scheme and our blind signature scheme, is proven in the random oracle model and, due to Ajtai’s result [2], is based on the worst-case hardness of the SVP. Ajtai showed that solving the average-case SVP is at least as hard as solving a related problem in the worst-case in lattices of a certain smaller dimension. The works of Micciancio and Regev [27] and [17] improve the tightness of the worst-case to average-case reduction.

Despite the uninstantiability result of Canetti, Goldreich, and Halevi [11], we believe that our construction is an important contribution and that we solve a longstanding problem because the previous efficient constructions [12], [31], [32], [1], [8], [13], [23], and [30] have one thing in common: they are built upon number theoretic assumptions, like the hardness of factoring large integers or computing discrete logarithms. The newer approaches of Boldyreva [9] and Okamoto [30] tend to use pairings and bilinear maps that yield very elegant constructions. They, however, are again based on the discrete logarithm problem in this specific setting.

None of the above schemes remain secure in the presence of reasonably powerful quantum computers, where both factoring and computing discrete logarithms becomes easy due to the seminal work of Shor [36]. Until the recent works of Lyubashevsky and Micciancio [25] and Gentry, Peikert, and Vaikuntanathan [17], it was not even clear that lattice problems can give rise to provably secure signature schemes. The earlier approaches of Goldreich, Goldwasser, and Halevi (GGH) [16] and Hoffstein, Howgrave-Graham, Pipher, Silverman, and Whyte (NTRUSign) [18] came without security proofs and were successfully broken by Nguyen [28] and Nguyen and Regev [29].

Finally, we would like to mention that there are also (inefficient) instantiations from general assumptions. The one of Juels, Luby, and Ostrovsky [20], e.g., can be used to construct a quantum-immune scheme because it is, like our construction, based on trapdoor permutations. In addition, there are the blind signature schemes of Fischlin [15] and of Hazay, Katz, Koo, and Lindell [19] that are built upon general assumptions as well. Whether they are quantum-immune, largely depends on the exact realization of primitives.

Organization. After a preliminaries section, we present our construction in Section 3. There, we also prove that our scheme has the well-established security properties. In Section 4, we discuss the details and the realization of the underlying trapdoor function.

2 Preliminaries

With n , we always denote the security parameter. $(a, b) \leftarrow \langle \mathcal{A}(x), \mathcal{B}(y) \rangle$ denotes 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} . The private outputs are a for \mathcal{A} and b for \mathcal{B} .

In the following, we recall the definitions of digital signature schemes and of blind signature schemes along with their respective security models, followed by a brief introduction to lattice theory.

Digital signatures. A digital signature scheme DS is a triple $(\text{Kg}, \text{Sig}, \text{Vf})$ where

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

Signature Generation. $\text{Sig}(sk, M)$ outputs a signature σ on a message M from the message space \mathcal{M} under sk .

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

Signature schemes are complete if for all $(sk, pk) \leftarrow \text{Kg}(n)$, all messages $M \in \mathcal{M}$, and any $\sigma \leftarrow \text{Sig}(sk, M)$, we have $\text{Vf}(pk, \sigma, M) = 1$.

Security of digital signature schemes is typically proven against existential forgery under a chosen message attack (EU-CMA), where an adversary wins if it outputs a signature on a *new* message M^* after accessing a signature oracle on a polynomial number of different messages. For our construction, we need the notion of strong unforgeability under a chosen message attack (SU-CMA), where the adversary even wins if it is able to output a *new pair* (M^*, σ^*) , i.e. it is not forced to output a signature on a new message. In the random oracle model, the adversary has access to a hash oracle H that is chosen from the family of all collision resistant hash functions $\mathcal{H}(n)$. The described concept is formalized in the following experiment.

Experiment $\text{Exp}_{\mathcal{A}, \text{DS}}^{\text{su-cma}}(n)$

$\text{H} \leftarrow \mathcal{H}(n)$

$(sk, pk) \leftarrow \text{Kg}(n)$

$(M^*, \sigma^*) \leftarrow \mathcal{A}^{\text{H}(\cdot), \text{Sig}(sk, \cdot)}(pk)$

let (M_i, σ_i) be the answer returned by $\text{Sig}(sk, \cdot)$ on input M_i , for $i = 1, \dots, k$.

Return 1 iff $\text{Vf}(pk, M^*, \sigma^*) = 1$ and $(M^*, \sigma^*) \notin \{(M_1, \sigma_1), \dots, (M_k, \sigma_k)\}$.

The scheme DS is called $(t, q_{\text{Sig}}, q_{\text{H}}, \epsilon)$ -strongly unforgeable if there is no adversary, running in time at most t while making at most q_{Sig} queries to the oracle $\text{Sig}(sk, \cdot)$ and at most q_{H} queries to the hash oracle, that succeeds in the above experiment with probability at least ϵ .

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

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

Signature Generation. $\text{Sig}(sk, M)$ 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 on the message M under sk . Thus, we write $(\mathcal{V}, \sigma) \leftarrow \langle \mathcal{S}(sk), \mathcal{U}(pk, M) \rangle$.

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

Completeness is defined as with digital signature schemes. 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 [20,32]. The notion of blindness is defined in the following experiment $\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}$, where the adversarial signer \mathcal{S}^* chooses two messages M_0, M_1 and interacts with two users who obtain blind signatures for the two messages in random order. After seeing the unblinded signatures in the original order, with respect to M_0, M_1 , the signer has to guess the message that has been signed for the first user.

Experiment $\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}(n)$

$b \leftarrow \{0, 1\}$

$(pk, sk) \leftarrow \text{BS.Kg}(n)$
 $(M_0, M_1) \leftarrow \mathcal{S}^*(pk, sk)$
 $(\mathcal{V}_0, \sigma_b) \leftarrow \langle \mathcal{S}^*(sk), \mathcal{U}_0(pk, m_b) \rangle$
 $(\mathcal{V}_1, \sigma_{1-b}) \leftarrow \langle \mathcal{S}^*(sk), \mathcal{U}_1(pk, m_{1-b}) \rangle$
 Either of the signatures might equal fail.
 If $\sigma_0 \neq \text{fail}$ and $\sigma_1 \neq \text{fail}$
 $d \leftarrow \mathcal{S}^*(sk, pk, \sigma_0, \sigma_1)$
 Else
 $d \leftarrow \mathcal{S}^*(sk, pk, \text{fail}, \text{fail})$
 Return 1 iff $d = b$

A signature scheme BS is (t, ϵ) -blind, if there is no adversary \mathcal{S}^* , running in time at most t , that wins the above experiment with advantage at least ϵ , where

$$\text{Adv}_{\mathcal{S}^*, \text{BS}}^{\text{blind}} = \Pr[\text{Exp}_{\mathcal{S}^*, \text{BS}}^{\text{blind}}(n) = 1] - \frac{1}{2}.$$

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 following 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.

Experiment $\text{Exp}_{\mathcal{U}^*, \text{BS}}^{\text{omf}}(n)$

$H \leftarrow \mathcal{H}(n)$
 $(pk, sk) \leftarrow \text{BS.Kg}(n)$
 $\{(M_1, \sigma_1), \dots, (M_j, \sigma_j)\} \leftarrow \mathcal{U}^{*H(\cdot), \langle \mathcal{S}(sk), \cdot \rangle}(pk)$
 Let ℓ be the number of (complete) interaction between \mathcal{U}^* and the signer.
 Return 1 iff

1. $m_i \neq m_j$ for all $1 \leq i < j \leq j$
2. $\text{BS.Vf}(pk, \sigma_i, M_i) = 1$ for all $i = 1, \dots, j$
3. $\ell < j$.

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

Lattices. A lattice in \mathbb{R}^n is a 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 called a *basis* of the lattice Λ and we write $\Lambda = \Lambda(\mathbf{B})$. The number of linearly independent vectors in the basis is the dimension of the lattice. Now, consider *modular lattices* as a special form of lattices. Given a modulus q , a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and the equation

$$\mathbf{A} \mathbf{v} \equiv \mathbf{0} \pmod{q},$$

then the set of all vectors $\mathbf{v} \in \mathbb{Z}_q^m$ that satisfy the above equation is a lattice. Lattices of this form are denoted with $\Lambda_q^\perp(\mathbf{A})$.

The main computational problem in lattices is the (approximate) 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

$$\|\mathbf{v}\|_p \leq \gamma \|\mathbf{w}\|_p \text{ for all } \mathbf{w} \in \Lambda \setminus \{\mathbf{0}\}$$

for a fixed approximation factor $\gamma \geq 1$. This problem is known to be \mathcal{NP} -hard for all ℓ_p norms [14,35,21] with a constant approximation factor. For exponential (in the lattice dimension) approximation factors, the problem is solvable in polynomial time by the famous LLL algorithm by Lenstra, Lenstra, and Lovász [24]. We refer the interested reader to a recent survey [34] by Regev for the currently known ‘‘approximability’’ and ‘‘inapproximability’’ results.

In the special case of modular lattices, there is also a special version of the SVP, named short integer solution problem (SIS). There, an algorithm is given a basis of $\Lambda_q^\perp(\mathbf{A})$ and is supposed to output a non-zero solution $\mathbf{v} \in \mathbb{Z}_q^m$ to the above equation. The algorithm succeeds if $\|\mathbf{v}\|_2 \leq \nu$ for a given norm bound ν . The SIS was, in principle, introduced by Ajtai [2] and its hardness is analyzed in [27] and [17]. The latter work also explicitly deals with the ℓ_∞ norm, which we will use in our security proofs.

3 Our Construction

In this section, we describe the construction of our blind signature scheme and prove its security in terms of *blindness* and *one-more unforgeability*. We start with the specification of the required trapdoor function, whose realization is given in the STOC 2008 paper of Gentry, Peikert, and Vaikuntanathan (GPV) [17]. Then, we recall the GPV signature scheme, also presented in [17], and show how it can be used to implement blind signatures. For both, trapdoor function and signature scheme, we need some modifications that have been clearly marked below. The details are given in Section 4.

Our blind signature scheme is built upon a family of trapdoor functions, which are almost as good as trapdoor permutations. The family is described via a triple $(\text{TrapGen}, \text{SampleDom}, \text{SamplePre})$ and has the following properties.

Function generation. There is an efficient algorithm TrapGen that outputs $(a, t) \leftarrow \text{TrapGen}(n)$, where a fully defines the function f_a and the trapdoor t is used to sample from the inverse $f_t^{-1}(\cdot)$, which is defined as $\text{SamplePre}(t, \cdot)$.

Efficiency. The function $f_a : D_n \rightarrow R_n$ is efficiently computable. Furthermore, the three finite sets R_n, D_n, D_n^* are efficiently recognizable and R_n is closed under addition. Furthermore, let $D_n^* \subseteq D_n$, such that $x_1 \pm x_2 \in D_n$ for all $x_1, x_2 \in D_n^*$.

One-wayness. Computing the function $f_t^{-1} : R_n \rightarrow D_n^*$, is infeasible without the trapdoor t .

Domain sampling with uniform output. $\text{SampleDom}(n)$ samples from some distribution over D_n^* , such that their images under f_a are uniformly distributed over R_n .

Preimage sampling. Let $y \in R_n$. $f_t^{-1}(y)$ samples $x \leftarrow \text{SampleDom}(n)$ under the condition that $f_a(x) = y$. The entropy of x is at least $\omega(\log(n))$.

Linearity. Let $x_1 + x_2 \in D_n$. $f_a(x_1 + x_2) = f_a(x_1) + f_a(x_2)$.

Collision resistance. There is no algorithm $\mathcal{A}(n, a)$ that outputs a pair $(x, x') \in D_n^2$, such that $x \neq x'$ and $f_a(x) = f_a(x')$, in time polynomial in n with noticeable probability.

Note that we slightly modified the original construction regarding the sets D_n, D_n^* . In [17], it is always the same, whereas we have introduced different D_n, D_n^* for trapdoor evaluation and preimage

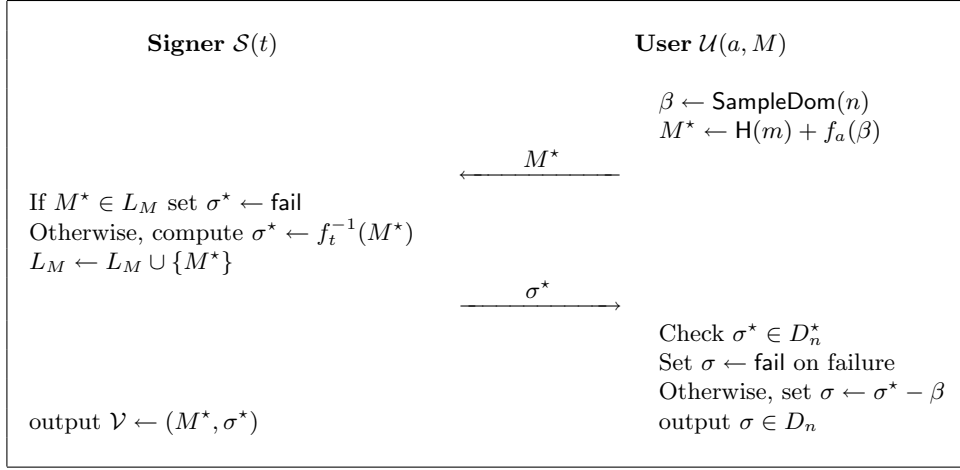


Fig. 1. Issue protocol of the blind signature scheme BS

sampling, respectively. As in the original work, we will always assume that the above properties, especially the statistical distributions, hold for f_a in a perfect sense.

In addition to the above trapdoor function, Gentry, Peikert, and Vaikuntanathan use the “hash-then-sign” paradigm with a full-domain hash function (cf. [10]) $\text{H} \leftarrow \mathcal{H}(n)$, where $\text{H} : \{0, 1\}^* \rightarrow R_n$ and \mathcal{H} is a family of collision resistant hash functions. In this setting, the GPV signature scheme is strongly unforgeable under a chosen message attack (cf. [17]). In Section 4, we show that this still holds with our modification.

With the modification $D_n^* \subseteq D_n$, the GPV signature scheme is a tuple $\text{GPV} = (\text{Kg}, \text{Sig}, \text{Vf})$, where:

Key generation. $\text{GPV.Kg}(n)$ outputs $(a, t) \leftarrow \text{TrapGen}(n)$.

Signature issue. Let $M \in \{0, 1\}^*$ be a message. $\text{GPV.Sig}(t, M)$ checks whether M has been signed before and, if so, outputs the same signature. Otherwise, it computes $\sigma \leftarrow f_t^{-1}(\text{H}(M))$, stores (M, σ) , and returns $\sigma \in D_n^*$.

Verification. Given a signature σ . $\text{GPV.Vf}(a, \sigma, M)$ returns 1 iff $\sigma \in D_n$ and $f_a(\sigma) = \text{H}(M)$.

Using a slight relaxation of the above signature scheme, we construct an equally efficient and provably secure blind signature scheme $\text{BS} = (\text{Kg}, \text{Sig}, \text{Vf})$ as follows.

Key generation. $\text{BS.Kg}(n)$ outputs $(a, t) \leftarrow \text{TrapGen}(n)$, where a is the public verification key and t is the secret signing key, and sets up a list of already signed messages $L_M = \{0\}$.

Signature protocol. The signature issue protocol for messages $M \in R_n$ is shown in Figure 1.

Note that the blind signature scheme is again stateful, i.e. the signer does not sign a blinded message twice and it does not sign $M^* = 0$ in particular. The returned signature is $\sigma^* \in D_n^*$.

Verification. $\text{BS.Vf}(a, \sigma, M)$ outputs 1 iff $\sigma \in D_n$ and $f_a(\sigma) = \text{H}(M)$.

If the user outputs fail, it may be that the signer is dishonest. In the special setting of e-cash, if the obtained signature σ is not in the domain of f_a , the process has to be repeated with a different M and the receiver of the signature has to reveal β to prove to the signer that the signature is literally worthless. For the moment, we assume that $\sigma^* - \beta \in D_n$. In Section 4, it becomes obvious that this always holds if both parties are honest.

Completeness. The scheme BS is complete because for all honestly generated key pairs (a, t) , all messages $m \in \{0, 1\}^*$, all outputs (β, M^*) of $\text{BS.Blind}(a, M)$, and all signatures $\sigma^* \leftarrow \text{BS.Sig}(t, M^*)$ we have

$$\sigma \leftarrow \sigma^* - \beta \in D_n$$

and

$$f_a(\sigma) = f_a(\sigma^* - \beta) = f_a(\sigma^*) - f_a(\beta) = f_a(f_t^{-1}(\mathbf{H}(M) + f_a(\beta))) - f_a(\beta) = \mathbf{H}(M).$$

Therefore, $\text{BS.Vf}(a, \sigma, M) = 1$.

In the following, we prove the security of our blind signature scheme. A blind signature scheme is called secure if it satisfies *blindness* and *one-more unforgeability* as described in Section 2.

Blindness. We prove that, like Chaum's blind signature scheme [12], BS is unconditionally blind, i.e. $(\infty, 0)$ -blind. The intuition is that the signer only sees random elements from R_n after the user has applied a random blinding value.

Theorem 1 (Blindness). *The blind signature scheme BS is $(\infty, 0)$ -blind.*

The idea of the proof is that, given the signer's views $\mathcal{V}_0, \mathcal{V}_1$ in the experiment $\text{Exp}_{S^*, \text{BS}}^{\text{blind}}$, there are always two equally distributed pairs of blinding values $(f_a(\beta_0), f_a(\beta_1))$ and $(f_a(\beta'_0), f_a(\beta'_1))$, such that both generate the same views, $(f_a(\beta_0), f_a(\beta_1))$ is a witness for b , and $(f_a(\beta'_0), f_a(\beta'_1))$ is a witness for $b' = 1 - b$. From this, we infer that the signer can only guess the correct b with probability $1/2$.

Proof. Let M_0, M_1 be the messages obtained from the signer. Assume that b is fixed and $\beta_0, \beta_1 \leftarrow \text{SampleDom}(n)$ were chosen by the users $\mathcal{U}_0, \mathcal{U}_1$, who compute

$$\begin{aligned} M_0^* &\leftarrow \mathbf{H}(M_b) + f_a(\beta_0), \\ M_1^* &\leftarrow \mathbf{H}(M_{1-b}) + f_a(\beta_1), \end{aligned}$$

and receive σ_0^* and σ_1^* , such that σ_0^* contains a signature on M_b and σ_1^* contains a signature on M_{1-b} . Then, the signer's views are

$$\begin{aligned} \mathcal{V}_0 &= (M_0^*, \sigma_0^*) \\ \text{and } \mathcal{V}_1 &= (M_1^*, \sigma_1^*). \end{aligned}$$

Now, we show that for the given choice of blinding values $(f_a(\beta_0), f_a(\beta_1))$, there is exactly one pair (B_0, B_1) that results in the same views while reversing the order, in which the messages are signed. Let

$$\begin{aligned} B_0 &\leftarrow \mathbf{H}(M_b) - \mathbf{H}(M_{1-b}) + f_a(\beta_0) \\ \text{and } B_1 &\leftarrow \mathbf{H}(M_{1-b}) - \mathbf{H}(M_b) + f_a(\beta_1) \end{aligned}$$

and observe that there are $\beta'_0, \beta'_1 \in D_n^*$ with $B_0 = f_a(\beta'_0)$ and $B_1 = f_a(\beta'_1)$. Since \mathbf{H} is a random oracle and $f_a(\beta_0), f_a(\beta_1)$ are distributed uniformly at random over R_n , so are the blinding values B_0 and B_1 . Assuming $b' = 1 - b$, the blinding values B_0, B_1 yield the following blinded messages:

$$\begin{aligned} \mathcal{U}_0 &\text{ computes } M_0^{*'} \leftarrow \mathbf{H}(M_{b'}) + B_0 = \mathbf{H}(M_{1-b}) + B_0 = \mathbf{H}(M_b) + f_a(\beta_0) = M_0^*; \\ \mathcal{U}_1 &\text{ computes } M_1^{*'} \leftarrow \mathbf{H}(M_{1-b'}) + B_1 = \mathbf{H}(M_b) + B_1 = \mathbf{H}(M_{1-b}) + f_a(\beta_1) = M_1^*. \end{aligned}$$

The resulting views \mathcal{V}'_0 and \mathcal{V}'_1 of the signer are equal to \mathcal{V}_0 and \mathcal{V}_1 , respectively. Therefore, there are indistinguishable witnesses $(f_a(\beta_0), f_a(\beta_1))$ for b and $(f_a(\beta'_0), f_a(\beta'_1))$ for $b' = 1 - b$, which concludes the proof. \square

One-more unforgeability. We prove that our blind signature scheme is unforgeable under a reasonable assumption, namely that the following “one-more trapdoor inversion problem” is hard.

Definition 1 (Chosen target trapdoor inversion problem (CTTI)). *The chosen target trapdoor inversion problem is defined via the following experiment, where the adversary \mathcal{A} has access to a challenge oracle O_{R_n} and to an inversion oracle f_t^{-1} . The adversary wins, if it outputs j preimages for challenges obtained from O_{R_n} , while making only $\iota < j$ queries to f_t^{-1} . The oracle f_t^{-1} does not answer queries twice and its does not invert 0.*

Experiment $\text{Exp}_{\mathcal{A}}^{\text{ctti}}(n)$

$(a, t) \leftarrow \text{TrapGen}(n)$

$(\pi, x_1, \dots, x_j) \leftarrow \mathcal{A}^{\mathsf{O}_{R_n}, f_t^{-1}(\cdot)}(n, a)$

Let y_1, \dots, y_ℓ be the challenges returned by O_{R_n} .

Let ι be the number of queries to f_t^{-1} .

Return 1 iff

1. $\pi : \{1, \dots, j\} \rightarrow \{1, \dots, \ell\}$ is injective and
2. $x_i \in D_n$ and $f_a(x_i) = y_{\pi(i)}$ for all $i = 1, \dots, j$ and
3. $\iota < j$.

The problem is (t, q_1, q_0, ϵ) -hard if there is no algorithm \mathcal{A} , running in time at most t , making at most q_1 inversion queries, and at most q_0 queries to O_{R_n} , which wins the above experiment with probability at least ϵ . The one-wayness of f_a gives us $(\text{poly}(n), 0, 1, \epsilon)$ -hardness, which we will extend to $(\text{poly}(n), \text{poly}(n), \text{poly}(n), \epsilon')$ -hardness for a negligible ϵ' . With our definition and this assumption, we follow the line of thought of Bellare, Namprempre, Pointcheval, and Semanko in [8]. They define a collection of “one-more” problems in the RSA context, which are perfectly tailored for proving one-more unforgeability. In [7], Bresson, Monnerat, and Vergnaud give a separation result on these “one-more” problems, showing that they cannot be proven equivalent to “simple” RSA inversion. The same dilemma seems to apply here as we are not able to prove equivalence of CTTI and simple trapdoor inversion.

In the following, we will assume $(\text{poly}(n), \text{poly}(n), \text{poly}(n), \epsilon)$ -hardness of CTTI on the grounds that it is similar to the provably hard problem of forging GPV signatures. In both cases, one has to find a solution $x \in D_n$ to the equation $f_a(x) = y$ for a given y , while knowing polynomially many distinct preimage-image pairs. Furthermore, we argue that an adversary cannot successfully recombine answers of f_t^{-1} in order to produce more preimages than there were oracle queries because adding or subtracting more than two of the oracle’s answers yields an invalid preimage $\notin D_n$. This will become obvious in Section 4, where we show how f_a is realized using lattices.

Theorem 2 (One-more unforgeability). *The BS blind signature scheme is $(t, q_{\text{Sig}}, q_{\text{H}}, \epsilon)$ -one-more unforgeable if the CTTI is $(t, q_{\text{Sig}}, q_{\text{H}}, \epsilon)$ -hard.*

Proof. Towards contradiction, we assume that there exists a successful forger \mathcal{A} against one-more unforgeability of BS. Using \mathcal{A} , we construct an algorithm \mathcal{B} via a black-box simulation, such that \mathcal{B} solves the respective instance of the CTTI. The simulation works as follows.

Setup. \mathcal{B} gets as input the public trapdoor parameter a and has access to the challenge oracle O_{R_n} and to a trapdoor inversion oracle f_t^{-1} . \mathcal{B} initializes a list $L_{\text{H}} \leftarrow \emptyset$ of pairs (M, c) , indexed by M , a list $L_1 \leftarrow \emptyset$ of pairs (M^*, σ^*) , indexed by M^* , and two counters $\ell \leftarrow 0$, $\iota \leftarrow 0$. It runs \mathcal{A} on input a in a black-box simulation.

Random oracle queries. On input M , \mathcal{B} looks up M in L_H . If it finds a pair (M, c) then it returns c . Otherwise, \mathcal{B} increments ι , chooses a new $c_\iota \leftarrow \mathcal{O}_{R_n}$, stores $(M_\iota \leftarrow M, c_\iota)$ in L_H . Afterwards, \mathcal{B} returns c_ι .

Blind signature queries. On input M^* , algorithm \mathcal{B} searches a pair (M^*, σ^*) in L_1 . If it exists, \mathcal{B} returns σ^* . Otherwise, algorithm \mathcal{B} increments ℓ , queries its inversion oracle $\sigma_\ell^* \leftarrow f_t^{-1}(M^*)$, stores $(M_\ell^* \leftarrow M^*, \sigma_\ell^*)$ in L_1 , and returns σ_ℓ^* .

Output. Finally, \mathcal{A} stops and outputs $((M_1, \sigma_1), \dots, (M_j, \sigma_j))$, $\ell < j$, for distinct messages. W.l.o.g., assume that $(M_i, c_i) \in L_H$, for all $i = 1, \dots, j$. Algorithm \mathcal{B} sets

$$\pi = \{(i, j) : f_a(\sigma_i) = c_j\}$$

and outputs $(\pi, \sigma_1, \dots, \sigma_j)$.

Analysis. First, observe that all of \mathcal{A} 's oracles are perfectly simulated. When \mathcal{A} calls H , algorithm \mathcal{B} draws a new challenge from its challenge oracle. Whenever \mathcal{A} queries its signature oracle on a new blinded message, \mathcal{B} calls its inversion oracle. Therefore, when \mathcal{A} outputs a one-more forgery, \mathcal{B} can use it to solve the CTTI. \mathcal{B} 's output is valid in the CTTI experiment because all preimages $(\sigma_i \in D_n)$ evaluate to challenges received from \mathcal{O}_{R_n} and the number of output inversions j is greater than the number of inversion queries ℓ . As for the map π , we state that it is injective. Otherwise, there would be a pair $\sigma \neq \sigma'$ in \mathcal{A} 's output with $f_a(\sigma) = f_a(\sigma') = H(M_i)$, which contradicts the collision resistance of f_a . Thus, \mathcal{B} is successful if \mathcal{A} is. \square

4 Realization

The underlying signature scheme was developed by Gentry, Peikert, and Vaikuntanathan (GPV) and presented at STOC 2008 [17]. It uses a modified Babai nearest plane algorithm [4,22] and two famous results by Ajtai [2,3] in order to build a trapdoor function that is arguably ‘‘as good as’’ a trapdoor permutation. Its security is proven in the random oracle model and reduces to the collision resistance of f_a , which in turn reduces to the hardness of finding short vectors in a lattice. We refer the reader to [17] and to [27] for further details and a comprehensive discussion of the involved lattice problems and on Gaussians in the lattice context. The practical hardness of these lattice problems is analyzed in [5] and subsequently in [6].

GPV trapdoor function. The trapdoor function from [17] is defined as follows.

Public parameters. Depending on the security parameter n , the other parameters in [17] can be chosen as

Modulus	$q = n^3$,
Domain dimension	$m = 5n \log(q)$,
Basis length bound	$L = m^{1+\epsilon}$, $\epsilon > 0$,
Gaussian parameter	$s = L\omega(\sqrt{\log(m)})$.

Spaces. The range is

$$R_n = \mathbb{Z}_q^n,$$

the (modified) domain D_n is

$$D_n = \{\mathbf{e} \in \mathbb{Z}^m : \|\mathbf{e}\|_\infty \leq 2s\omega(\sqrt{\log(m)})\},$$

and the range of the function `SamplePre` is

$$D_n^* = \{\mathbf{e} \in \mathbb{Z}^m : \|\mathbf{e}\|_\infty \leq s\omega(\sqrt{\log(m)})\}.$$

Trapdoor description. The public trapdoor key a describes the above public parameters and the public matrix

$$\mathbf{A} \in \mathbb{Z}_q^{n \times m}.$$

The set

$$\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} \equiv \mathbf{0} \pmod{q}\}$$

describes a lattice, for which the secret trapdoor parameter t describes a basis \mathbf{T} , such that

$$\|\tilde{\mathbf{T}}\| \leq L.$$

Here, $\tilde{\mathbf{T}}$ is the Gram-Schmidt orthogonalization of \mathbf{T} and $\|\cdot\|$ is defined as

$$\|\mathbf{X}\| = \left\| \begin{pmatrix} | & & | \\ \mathbf{x}_1 & \cdots & \mathbf{x}_c \\ | & & | \end{pmatrix} \right\| = \max_{i=1,\dots,c} \|\mathbf{x}_i\|_2.$$

Trapdoor evaluation. On input \mathbf{x} , the trapdoor function $f_a(\mathbf{x})$ evaluates to

$$\mathbf{y} \leftarrow \mathbf{A}\mathbf{x} \pmod{q}.$$

Preimage sampling. Sampling from f_t^{-1} is performed via a modified Babai nearest plane algorithm. The algorithm explicitly uses \mathbf{T} and relies on its short length. On input \mathbf{y} , it performs the following steps.

1. Compute $\mathbf{t} \in \mathbb{Z}_q^m$, such that $\mathbf{A}\mathbf{t} \equiv \mathbf{y} \pmod{q}$. This is done by linear algebra and most likely yields a $\mathbf{t} \notin D_n^*$.
2. Use the trapdoor basis \mathbf{T} to sample a vector \mathbf{v} from a Gaussian distribution around $-\mathbf{t}$ with standard deviation s , conditioned with $\mathbf{A}\mathbf{v} \equiv \mathbf{0} \pmod{q}$, and output $\mathbf{x} = \mathbf{t} + \mathbf{v} \in D_n^*$.

The described trapdoor function has all the properties mentioned in Section 3. As for the required linearity in our blind signature scheme, note that f_a is linear in the sense that for all $\mathbf{x}_1 + \mathbf{x}_2 \in D_n$:

$$f_a(\mathbf{x}_1 + \mathbf{x}_2) \equiv \mathbf{A}(\mathbf{x}_1 + \mathbf{x}_2) \equiv \mathbf{A}\mathbf{x}_1 + \mathbf{A}\mathbf{x}_2 \equiv f_a(\mathbf{x}_1) + f_a(\mathbf{x}_2) \pmod{q}.$$

Therefore, all computations in D_n , D_n^* , and R_n have to be performed modulo q .

Concerning security of the GPV signature scheme, [17] states that it is strongly unforgeable if the problem of finding short integer solutions $\mathbf{v} \in \mathbb{Z}^m$, $\|\mathbf{v}\|_\infty \leq 2s\omega(\sqrt{\log(m)})$, of the equation

$$\mathbf{A}\mathbf{v} \equiv \mathbf{0} \pmod{q},$$

i.e. the respective SIS problem, is hard. As for our modified setting, with D_n^* and D_n , we need a slightly stronger assumption, i.e. the above problem has to be hard with $\|\mathbf{v}\|_\infty \leq 3s\omega(\sqrt{\log(m)})$. Furthermore, we claim that this special setting cannot be exploited to forge a signature $\sigma' \in D_n$ from two valid signatures $\sigma_1, \sigma_2 \in D_n^*$ by simply adding them as $\sigma' \leftarrow \sigma_1 + \sigma_2$ because of the collision resistance of the full-domain hash H . We support this intuition by a modified security proof for the GPV signature scheme.

Unforgeability of the modified GPV signature In the following, we adapt the proof from [17] (Proposition 6.1 in the extended version) to the modified signature scheme in Section 3.

Theorem 3. *Let the parameters $n, m, q, L, s, \Lambda_q^\perp(\mathbf{A}), a$ be as defined above and let $T_{\text{SampleDom}}(n), T_{f_a}(n)$ be the cost functions for domain sampling and trapdoor evaluation. The modified GPV signature scheme is $(t, q_{\text{Sig}}, q_{\text{H}}, \epsilon)$ -strongly unforgeable if finding a vector $\mathbf{v} \in \Lambda_q^\perp(\mathbf{A}) \setminus \{\mathbf{0}\}$ with*

$$\|\mathbf{v}\|_\infty \leq 3s\omega(\sqrt{\log(m)})$$

is (t', ϵ') -hard with

$$t' = t + (q_{\text{Sig}} + q_{\text{H}})(T_{\text{SampleDom}}(n) + T_{f_a}(n)) \quad \text{and} \quad \epsilon' = \epsilon - 2^{-\omega(\log(n))}.$$

Proof. Given a successful adversary \mathcal{A} against strong unforgeability with success probability ϵ , we build an algorithm \mathcal{B} that finds a collision in f_a and, with that, a short vector in $\Lambda_q^\perp(\mathbf{A})$. Algorithm \mathcal{B} runs \mathcal{A} in a black-box simulation and uses random oracle techniques to simulate the signature oracle.

Setup. Algorithm \mathcal{B} receives the public parameter a of $\Lambda_q^\perp(\mathbf{A})$ as input sets up a list $L_{\text{H}} \leftarrow \emptyset$ of triples (M, h, σ) in order to simulate H and f_t^{-1} consistently. It runs \mathcal{A} on input a .

Random oracle H. When queried with $M \in \{0, 1\}^*$, algorithm \mathcal{B} looks for a triple $(M, h, \sigma) \in L_{\text{H}}$. If it exists, \mathcal{B} returns h . Otherwise, the simulator chooses $\sigma \leftarrow \text{SampleDom}(n)$, sets $h \leftarrow f_a(\sigma)$, stores (M, h, σ) in L_{H} , and returns h .

Signature Queries. On input $M \in \{0, 1\}^*$, algorithm \mathcal{B} runs $\text{H}(M)$, yielding a triple $(M, h, \sigma) \in L_{\text{H}}$. The simulator returns $\sigma \in D_n^*$.

Output. Finally, \mathcal{A} stops and returns a valid forgery (M^*, σ^*) with $\text{H}(M^*) = h^*$ and $\sigma^* \in D_n$. W.l.o.g., there is a triple $(M^*, h^*, \sigma) \in L_{\text{H}}$ with $\sigma \in D_n^*$. Algorithm \mathcal{B} outputs $\sigma^* - \sigma$.

Analysis. Observe that \mathcal{B} simulates the random oracle and the signature oracle perfectly and consistently. As for the output of \mathcal{B} , we have to show that $\sigma^* - \sigma \neq \mathbf{0}$ holds but with negligible probability. We have to distinguish three cases:

1. If $\sigma^* \in D_n \setminus D_n^*$, the condition trivially holds.
2. The adversary \mathcal{A} outputs a forgery in the strong sense, i.e. it has previously queried its signature oracle on M^* . Then, we have $\sigma^* - \sigma \neq \mathbf{0}$ by definition.
3. Algorithm \mathcal{A} has not queried its signature oracle on M^* . W.l.o.g., it has queried H on M^* and \mathcal{B} has a triple $(M^*, h^*, \sigma) \in L_{\text{H}}$. By the minimum conditional entropy $\omega(\log(n))$ of σ , we infer that $\sigma^* = \sigma$ with probability at most $2^{-\omega(\log(n))}$, which is still negligible.

Since σ^* and σ are valid signatures on M^* , we have

$$f_a(\sigma^*) = \text{H}(M^*) = f_a(\sigma)$$

and therefore a non-trivial solution to the characteristic equation of $\Lambda_q^\perp(\mathbf{A})$:

$$\mathbf{A}(\sigma^* - \sigma) \equiv \mathbf{0} \pmod{q}.$$

In consequence, algorithm \mathcal{B} has learned a lattice vector of norm

$$\|\sigma^* - \sigma\|_\infty \leq 2s\omega(\sqrt{\log(m)}) + s\omega(\sqrt{\log(m)}) \leq 3s\omega(\sqrt{\log(m)}).$$

The overhead of the reduction is dominated by the computational cost for domain sampling and trapdoor evaluation. In the worst-case, the adversary \mathcal{A} never queries H and Sig with the same message, which is why the overhead is $(q_{\text{Sig}} + q_{\text{H}})(T_{\text{SampleDom}}(n) + T_{f_a}(n))$. \square

5 Conclusions

We have shown how to construct a provably secure blind signature scheme from a lattice based signature scheme in the random oracle model. The conjectured quantum-immunity of the involved lattice problem makes our scheme secure even in the presence quantum computers. However, it is still an open question whether there exists an efficient quantum-immune blind signature scheme in the standard model. Furthermore, as our basis of security, we have introduced and analyzed the chosen target trapdoor inversion problem (CTTI), whose precise relation to SIS is unknown. We hope that this problem will be useful for other applications as well and eventually be proven equivalent to one of the well-established lattice problems.

Acknowledgments

The author would like to thank Dominique Schröder for numerous fruitful discussions on the subject. Furthermore, the author thanks Michael Schneider and Richard Lindner for reviewing parts of this work.

References

1. Masayuki Abe. *A Secure Three-Move Blind Signature Scheme for Polynomially Many Signatures*. Advances in Cryptology — Eurocrypt 2001, Lecture Notes in Computer Science, pages 136–151. Springer-Verlag, 2001.
2. Miklós Ajtai. *Generating Hard Instances of Lattice Problems (Extended Abstract)*. Proceedings of the Annual Symposium on the Theory of Computing (STOC) 1996, Lecture Notes in Computer Science, pages 99–108. Springer-Verlag, 1996.
3. Miklós Ajtai. *Generating Hard Instances of the Short Basis Problem*. ICALP 1999, Lecture Notes in Computer Science, pages 1–9. Springer-Verlag, 1999.
4. László Babai. *On Lovász’ lattice reduction and the nearest lattice point problem*. *Combinatorica*, 6(1):1–13, 1986.
5. Johannes Buchmann, Richard Lindner, and Markus Rückert. *Explicit hard instances of the shortest vector problem*. Post-Quantum Cryptography — PQCrypto 2008, Lecture Notes in Computer Science, pages 79–94. Springer-Verlag, 2008.
6. Johannes Buchmann, Richard Lindner, Markus Rückert, and Michael Schneider. *Explicit hard instances of the shortest vector problem (extended version)*. Number 2008/333 in Cryptology eprint archive. eprint.iacr.org, 2008.
7. Emmanuel Bresson, Jean Monnerat, and Damien Vergnaud. *Separation Results on the “One-More” Computational Problems*. Topics in Cryptology — Cryptographer’s Track, RSA Conference (CT-RSA)2008, Lecture Notes in Computer Science, pages 71–87. Springer-Verlag, 2008.
8. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. *The One-More-RSA-Inversion Problems and the Security of Chaum’s Blind Signature Scheme*. *Journal of Cryptology*, 16(3):185–215, 2003.
9. Alexandra Boldyreva. *Threshold Signatures, Multisignatures and Blind Signatures Based on the Gap-Diffie-Hellman-Group Signature Scheme*. Public-Key Cryptography (PKC) 2003, Volume 2567 of Lecture Notes in Computer Science, pages 31–46. Springer-Verlag, 2003.
10. Mihir Bellare and Pil Rogaway. *Random Oracles are Practical: A Paradigm for Designing Efficient Protocols*. Proceedings of the Annual Conference on Computer and Communications Security (CCS). ACM Press, 1993.
11. Ran Canetti, Oded Goldreich, and Shai Halevi. *The random oracle methodology, revisited*. *J. ACM*, 51(4):557–594, 2004.
12. David Chaum. *Blind Signatures for Untraceable Payments*. Advances in Cryptology — Crypto 1982, pages 199–203. Plenum, New York, 1983.
13. Jan Camenisch, Maciej Koprowski, and Bogdan Warinschi. *Efficient Blind Signatures Without Random Oracles*. Security in Communication Networks, Volume 3352 of Lecture Notes in Computer Science, pages 134–148. Springer-Verlag, 2004.
14. Irit Dinur. *Approximating SVP_{∞} to within almost-polynomial factors is NP-hard*. *Theoretical Computer Science*, 285(1):55–71, 2002.

15. Marc Fischlin. *Round-Optimal Composable Blind Signatures in the Common Reference String Model*. Advances in Cryptology — Crypto 2006, Volume 4117 of Lecture Notes in Computer Science, pages 60–77. Springer-Verlag, 2006.
16. O. Goldreich, S. Goldwasser, and S. Halevi. *Public-key cryptosystems from lattice reduction problems*. Advances in Cryptology — Crypto 1997, Volume 1294 of Lecture Notes in Computer Science, pages 112–131. Springer-Verlag, 1997.
17. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. *Trapdoors for hard lattices and new cryptographic constructions*. Proceedings of the Annual Symposium on the Theory of Computing (STOC) 2008, pages 197–206. ACM Press, 2008.
18. Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. *NTRUSIGN: Digital signatures using the NTRU lattice*. Topics in Cryptology — Cryptographer’s Track, RSA Conference (CT-RSA) 2003, Lecture Notes in Computer Science, pages 122–140. Springer-Verlag, 2003.
19. Carmit Hazay, Jonathan Katz, Chiu-Yuen Koo, and Yehuda Lindell. *Concurrently-Secure Blind Signatures Without Random Oracles or Setup Assumptions*. Theory of Cryptography Conference (TCC) 2007, Volume 4392 of Lecture Notes in Computer Science, pages 323–341. Springer-Verlag, 2007.
20. Ari Juels, Michael Luby, and Rafail Ostrovsky. *Security of Blind Digital Signatures*. Advances in Cryptology — Crypto 1997, Volume 1294 of Lecture Notes in Computer Science, pages 150–164. Springer-Verlag, 1997.
21. Subhash Khot. *Hardness of approximating the shortest vector problem in lattices*. *J. ACM*, 52(5):789–808, 2005.
22. Philip N. Klein. *Finding the closest lattice vector when it’s unusually close*. Proceedings of the Annual Symposium on Discrete Algorithms (SODA) 2000, pages 937–941. ACM Press, 2000.
23. Aggelos Kiayias and Hong-Sheng Zhou. *Two-Round Concurrent Blind Signatures without Random Oracles*. Number 2005/435 in Cryptology eprint archive. eprint.iacr.org, 2005.
24. A. Lenstra, H. Lenstra, and L. Lovász. *Factoring polynomials with rational coefficients*. *Mathematische Annalen*, 261(4):515–534, 1982.
25. Vadim Lyubashevsky and Daniele Micciancio. *Asymptotically Efficient Lattice-Based Digital Signatures*. Theory of Cryptography Conference (TCC) 2008, Lecture Notes in Computer Science, pages 37–54. Springer-Verlag, 2008.
26. Christoph Ludwig. *A Faster Lattice Reduction Method Using Quantum Search*. ISAAC 2003, Volume 2906 of Lecture Notes in Computer Science, pages 199–208. Springer-Verlag, 2003.
27. Daniele Micciancio and Oded Regev. *Worst-Case to Average-Case Reductions Based on Gaussian Measures*. *SIAM Journal on Computing*, 37(1):267–302, 2007.
28. P.Q. Nguyen. *Cryptanalysis of the Goldreich-Goldwasser-Halevi cryptosystem from Crypto’97*. Advances in Cryptology — Crypto 1999, Volume 1666 of Lecture Notes in Computer Science, pages 288–304. Springer-Verlag, 1999.
29. P. Q. Nguyen and O. Regev. *Learning a Parallelepiped: Cryptanalysis of GGH and NTRU Signatures*. Advances in Cryptology — Eurocrypt 2006, Volume 4004 of Lecture Notes in Computer Science, pages 215–233. Springer-Verlag, 2006.
30. Tatsuaki Okamoto. *Efficient Blind and Partially Blind Signatures Without Random Oracles*. Theory of Cryptography Conference (TCC) 2006, Volume 3876 of Lecture Notes in Computer Science, pages 80–99. Springer-Verlag, 2006.
31. David Pointcheval and Jacques Stern. *New Blind Signatures Equivalent to Factorization (extended abstract)*. ACM Conference on Computer and Communications Security, pages 92–99. ACM Press, 1997.
32. David Pointcheval and Jacques Stern. *Security Arguments for Digital Signatures and Blind Signatures*. *Journal of Cryptology*, 13(3):361–396, 2000.
33. Oded Regev. *Quantum Computation and Lattice Problems*. Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS) 2002, pages 520–529. IEEE Computer Society, 2002.
34. Oded Regev. *On the Complexity of Lattice Problems with Polynomial Approximation Factors*, 2007. A survey for the LLL+25 conference.
35. Oded Regev and Ricky Rosen. *Lattice problems and norm embeddings*. Proceedings of the Annual Symposium on the Theory of Computing (STOC) 2006, pages 447–456. ACM Press, 2006.
36. Peter W. Shor. *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.