

Efficient Quantum-immune Blind Signatures

— preliminary version —

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

Cryptography and Computeralgebra
Department of Computer Science
TU Darmstadt

July 25, 2008

Keywords Quantum-immune, blind signatures

Abstract We present an idea for the the first quantum-immune blind signature scheme. Our scheme is provably secure, efficient, and round-optimal.

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.

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 (SVP) in a lattice. Although the work of Ludwig [10] suggests that today's lattice reduction algorithms can benefit from the intrinsic parallelity in quantum computation, this does not invalidate the assumption. Slightly larger security parameters are considered to be a sufficient countermeasure.

Using the SVP as our security assumption, we construct the first quantum-immune blind signature scheme. As for its efficiency, we state that it is almost as efficient as the underlying signature scheme proposed by Gentry, Peikert, and Vaikuntanathan (GPV) [8]. With its two rounds, it is even *round-optimal*. 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, is based on the worst case hardness of the SVP.

All previous constructions have one thing in common. They are built upon number theoretic assumptions, like the hardness of factoring large integers or computing discrete logarithms. Newer approaches, like that of [5], 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 assumptions hold in the presence of quantum computers, where both factoring and computing discrete logarithms becomes easy due to the seminal work of Peter Shor [11].

Despite the uninstantiability result of Canetti, Goldreich, and Halevi [7], we believe that our construction is an important step towards quantum-immune blind signature schemes. A security proof in the standard model, however, remains an open problem.

Organization. After a brief preliminaries section, we present our construction in Section 3. There, we also prove that our scheme has the well-established security properties and point out the open problems. In Section 4, we discuss the details and the realization of the underlying trapdoor permutation. Finally, in Section 5, we propose reasonable parameters that lead to secure and efficient instantiations of our scheme.

2 Preliminaries

With n , we always denote the security parameter. $\langle \cdot \Leftarrow \cdot \rangle$ denotes the protocol view generated by two entities, i.e. the messages they exchange. 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 equal if they cannot be distinguished by any computationally unbounded algorithm with non-negligible probability.

3 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*.

The underlying signature scheme was developed by Gentry, Peikert, and Vaikuntanathan (GPV) and presented at STOC 2008 [8]. It is built upon a family of trapdoor functions, which are arguably as good as trapdoor permutations. The family is described via a triple $(\text{TrapGen}, \text{SampleDom}, \text{SamplePre})$ and has, among others, 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 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 two sets R_n, D_n are efficiently recognizable and R_n is closed under addition.

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 values from some distribution over D_n , such that their images under f_a are uniformly distributed over R_n .

Pre-image sampling. Let $y \in R_n$. $f_t^{-1}(y)$ samples $x \leftarrow \text{SampleDom}(n)$ under the condition that $f_a(x) = y$.

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 exists 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 non-negligible probability.

In addition to the above trapdoor function, Gentry, Peikert, and Vaikuntanathan use the “hash-then-sign” paradigm with a full-domain hash function (cf. [6]) $H \leftarrow \mathcal{H}(n)$, where $H : \{0, 1\}^* \rightarrow R_n$ and \mathcal{H} is a family of collision-resistant hash functions implementing the random oracle. In this setting, the GPV signature scheme is strongly unforgeable under a chosen message attack. Strong unforgeability of digital signatures means that an adversary is allowed to adaptively query a signature oracle on chosen messages. The adversary wins if it is able to output a new message-signature pair (m, σ) , in the sense that the signature oracle has never answered with σ on the query m .

The GPV signature scheme is a tuple $\text{GPV} = (\text{Kg}, \text{Sig}, \text{Vf})$, where

Key generation. $\text{GPV.Kg}(1^n)$ outputs $(a, t) \leftarrow \text{TrapGen}(1^n)$.

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 $\sigma \leftarrow \text{SamplePre}(t, H(m))$, stores (m, σ) , and returns σ .

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

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

Blinding. Let $m \in \{0, 1\}^*$ be a message. $\text{BS.Blind}(a, m)$ chooses a blinding value $\beta \leftarrow \text{SampleDom}(n)$ and computes $m^* \leftarrow H(m) + f_a(\beta)$. The output is (β, m^*) .

Signature issue. Let m^* be a blinded message. $\text{BS.Sig}(t, m^*)$ computes $\sigma^* \leftarrow f_t^{-1}(m^*)$ and returns σ^* .

Unblinding. Let σ^* be a blinded signature for the message m and the blinding value β . $\text{BS.Unblind}(a, m, \beta, \sigma^*)$ computes $\sigma \leftarrow \sigma^* - \beta$. After a postprocessing step on σ , explained in Section 4, it checks whether $\sigma \in D_n$ and $f_a(\sigma) = H(m)$. If either of the conditions is violated, the algorithm aborts with fail.

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

If **BS.Unblind** aborts, 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 $\sigma^* - \beta \in D_n$.

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 **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, $\text{BS.Vf}(a, \sigma, m) = 1$.

Open problem. The unblinding step uses postprocessing to “optimize” the obtained signature in order to fit it into D_n . It is unclear whether this is possible in all cases. A different, and likely more fruitful, approach might be to accept signatures that are slightly outside the trapdoor’s domain. We assume that this problem is efficiently solvable.

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 defined by Juels, Luby, and Ostrovsky in [9].

Blindness. The notion of blindness is defined in the following experiment, 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, according to m_0, m_1 , the signer has to guess the message that has been signed for the first user.

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

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

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

$(m_0, m_1) \leftarrow \mathcal{S}^*(n, pk, sk)$

Setup users $\mathcal{U}_0(n, pk, m_b), \mathcal{U}_1(n, pk, m_{1-b})$

$\mathcal{U}_0, \mathcal{U}_1$ interact with \mathcal{S}^* using the blind signature protocol of **BS**.

$\mathcal{U}_0, \mathcal{U}_1$ output signatures σ_b on m_b and σ_{1-b} on m_{1-b} ,

where either of them might equal fail.

$d \leftarrow \mathcal{S}^*(n, sk, pk, \sigma_0, \sigma_1)$

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 next theorem proves that BS is computationally blind, i.e. $(\text{poly}(n), \epsilon)$ -blind for a negligible ϵ .

Theorem 1 (Blindness). *The blind signature scheme BS is $(\text{poly}(n), \epsilon)$ -blind.*

The proof uses random oracle techniques in order to construct a pair of blinding values for any given pair (β_0, β_1) that generates the same view while reversing the order in which the signatures are obtained from the signer. The actual proof largely depends on the solution of the above mentioned open problems.

One-more unforgeability. Unforgeability in the context of blind signatures is defined in the experiment $\text{Exp}_{\mathcal{A}, \text{BS}}^{\text{omf}}$. There, a malicious user \mathcal{A} is successful if it is possible to obtain $k + 1$ distinct signatures on $k + 1$ messages from k interactions with the signer. More formally:

Experiment $\text{Exp}_{\mathcal{A}, \text{BS}}^{\text{omf}}(n)$
 $b \leftarrow \{0, 1\}$
 $\text{H} \leftarrow \mathcal{H}(n)$
 $(pk, sk) \leftarrow \text{BS.Kg}(n)$
 $\{(m_1, \sigma_1), \dots, (m_j, \sigma_j)\} \leftarrow \mathcal{A}^{\text{H}(\cdot), \text{BS.Sig}(sk, \cdot)}(n, pk)$
Let ℓ be the number of interaction between \mathcal{A} and the signer.
Return 1 iff $\text{BS.Vf}(pk, \sigma_i, m_i) = 1$ for all $i = 1, \dots, j$ and $\ell < j$.

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 .

We prove that our blind signature scheme is provably secure 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 $\ell < j$ queries to f_t^{-1} .*

Experiment $\text{Exp}_{\mathcal{A}}^{\text{ctti}}(n)$
 $(a, t) \leftarrow \text{TrapGen}(n)$
 $(\pi, x_1, \dots, x_j) \leftarrow \mathcal{A}^{\text{O}_{R_n}, f_t^{-1}(\cdot)}(n, a)$
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. $f_a(x_i) = y_{\pi(i)}$ 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 \mathcal{O}_{R_n} , which wins the above experiment with probability larger than 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 [4]. They define a collection of “one-more” problems in the RSA context.

As for its hardness, we show that it is as hard as forging GPV signatures.

Theorem 2. *The CTTI is (t, q_1, q_0, ϵ) -hard if and only if the GPV signature is (t, q_1, q_0, ϵ) -strongly unforgeable.*

Proof. We show both directions separately.

CTTI \Rightarrow GPV: Let’s assume that GPV is not (t, q_1, q_0, ϵ) -strongly unforgeable. Thus, there exists a forger \mathcal{A} against strong unforgeability. Using \mathcal{A} , we construct an adversary \mathcal{B} that solves the CTTI. The adversary \mathcal{B} works as follows.

Setup. \mathcal{B} sets up a list $L_{\mathbb{H}} \leftarrow \emptyset$ of triples (m, c, s) , which is indexed by the first component, and a counter $\ell \leftarrow 0$. It gets as input the public trapdoor parameter a and executes \mathcal{A} on input a in a black-box simulation. \mathcal{B} has access to \mathcal{O}_{R_n} and $f_t^{-1}(\cdot)$.

Random oracle \mathbb{H} . For each query m of \mathcal{A} to the random oracle \mathbb{H} , algorithm \mathcal{B} searches $L_{\mathbb{H}}$ for a triple $(m, c, *)$. If it exists, \mathcal{B} outputs c . Otherwise, \mathcal{B} increments ℓ , queries its challenge oracle $c_\ell \leftarrow \mathcal{O}_{R_n}$, stores $(m_\ell \leftarrow m, c_\ell, \square)$ in $L_{\mathbb{H}}$, and outputs c_ℓ . \square serves as a placeholder for “uninitialized”.

Signature queries. When \mathcal{A} queries its signature oracle on m , algorithm \mathcal{B} searches $L_{\mathbb{H}}$ for a triple (m_i, c_i, s_i) . If it exists, \mathcal{B} outputs s_i . Otherwise, \mathcal{B} queries f_t^{-1} with $\mathbb{H}(m_i)$, receives s_i , stores (m_i, c_i, s_i) in $L_{\mathbb{H}}$, and returns s_i to \mathcal{A} .

Output. When \mathcal{A} stops, it outputs a forgery (m^*, σ^*) . Assume $m^* = m_j$. Let $L'_{\mathbb{H}} = \{(m^{(1)}, c^{(1)}, s^{(1)}), \dots, (m^{(q_1)}, c^{(q_1)}, s^{(q_1)})\}$ be the set of all triples in $L_{\mathbb{H}}$, excluding those of form $(*, *, \square)$. \mathcal{B} sets

$$\pi = \{(i, j) : \exists a^{(i)} \in L'_{\mathbb{H}} \exists b_j \in L_{\mathbb{H}} : a^{(i)} = b_j\} \cup \{(q_1 + 1, j)\}$$

and outputs $(\pi, s^{(1)}, \dots, s^{(q_1)}, \sigma^*)$.

Analysis. Note that \mathcal{B} perfectly simulates \mathcal{A} ’s environment. Since f_a is collision resistant, we can safely assume that \mathcal{A} outputs a forgery on a message m^* that has never been sent to the signing oracle. Thus, \mathcal{B} has not queried f_t^{-1} on $\mathbb{H}(m^*)$. Therefore, \mathcal{B} makes $\iota = q_1$ queries to f_t^{-1} and outputs $q_1 + 1$ preimages along with an injective map π . Thus, the first and last requirements in the CTTI experiment are met. As for the second requirement, we state that $f_a(s_i) = \mathbb{H}(m_{\pi(i)})$ for all $i \neq j$ and $f_a(\sigma^*) = \mathbb{H}(m^*) = \mathbb{H}(m_{\pi(j)})$. Therefore, \mathcal{B} is successful whenever \mathcal{A} is.

CTTI \Leftarrow GPV: Now, assume that the CTTI is not (t, q_I, q_O, ϵ) -hard, i.e. there exists an adversary \mathcal{A} that efficiently solves the problem. We show that \mathcal{A} can be used to break strong unforgeability of GPV. We construct a forger \mathcal{B} as follows.

Setup. \mathcal{B} gets as input the public trapdoor parameter a . It sets up a list $L_H \leftarrow \emptyset$ of triples (m, c, x) , indexed by m . Furthermore it initializes a counter $\ell \leftarrow 0$. It runs a black-box simulation of \mathcal{A} on input a .

Random oracle H. On input m , \mathcal{B} searches L_H for a triple $(m, c, *)$. If it exists, it outputs c . Otherwise, \mathcal{B} increases ℓ , chooses a new $c_\ell \leftarrow R_n$, and stores $(m_\ell \leftarrow m, c_\ell, \square)$ in L_H , where “ \square ” denotes “uninitialized”. Finally, \mathcal{B} returns c_ℓ .

Challenge oracle queries. \mathcal{B} chooses a new $m \leftarrow D_n$, computes $c \leftarrow H(m)$, and returns c .

Inversion queries. On input c , \mathcal{B} searches L_H for a triple (m_i, c, x_i) . If it exists and $x_i \neq \square$ then \mathcal{B} outputs x_i . If it does not exist then \mathcal{B} increments ℓ , sets $i \leftarrow \ell$, chooses a new $m_\ell \leftarrow D_n$, and adds (m_ℓ, c, \square) to L_H . Finally, \mathcal{B} queries $x_\ell \leftarrow f_t^{-1}(c)$, stores (m_i, c, x_i) , and returns x_i .

Output. When \mathcal{A} stops, it outputs (π, x_1, \dots, x_j) . Algorithm \mathcal{B} searches the lowest index i , for which $(m_{\pi(i)}, c_{\pi(i)}, \square) \in L_H$. It outputs the forgery $(m_{\pi(i)}, x_i)$.

Analysis. First of all, note that \mathcal{B} perfectly simulates all of \mathcal{A} 's oracles. Since \mathcal{A} is a successful chosen target trapdoor inverter, there is an index i with $f_a(x_i) = c_{\pi(i)}$, such that \mathcal{A} never queried the inversion oracle on $c_{\pi(i)}$. Therefore, \mathcal{B} has never queried its signature oracle on $H(m_{\pi(i)}) = c_{\pi(i)}$ and x_i is a valid forgery on the message $m_{\pi(i)}$.

In both proofs, the number of inversion queries equals the number of signature queries and the number of challenge oracle queries equals the number of queries to the random oracle. The overhead of handling \mathcal{A} 's queries is minimal and consists mainly of list operations that can be neglected because they are essentially the same in both reductions. This concludes the proof. \square

Using the last theorem, we can now prove one-more unforgeability of our blind signature scheme.

Theorem 3. *The BS blind signature scheme is $(t, q_{\text{Sig}}, q_H, \epsilon)$ -one-more unforgeable if the CTTI is $(t, q_{\text{Sig}}, q_H, \epsilon)$ -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-block 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 $L_H \leftarrow \emptyset$ of pairs (m, c) , indexed by m , a list $L_I \leftarrow \emptyset$ 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 i , chooses a new c_i , stores $(m_i \leftarrow m, c_i)$ in L_H . Afterwards, \mathcal{B} returns c_i .

Blind signature queries. On input m^* , algorithm \mathcal{B} searches a pair (m^*, σ^*) in L_I . 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_I , and returns σ_ℓ^* .

Output. Finally, \mathcal{A} stops and outputs $((m_1, \sigma_1), \dots, (m_j, \sigma_j))$, $\ell < j$. W.l.o.g., assume that $(m_i, c_i) \in L_H$, for all $i = 1, \dots, j$. Algorithms \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 evaluate to challenges received from 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 when \mathcal{A} is.

Again, the overhead of handling \mathcal{A} 's queries is dominated by simple list processing and can be neglected. \square

Together with Theorem 2, our construction is one-more unforgeable if the GPV signature is strongly unforgeable.

4 Realization

The underlying signature scheme was developed by Gentry, Peikert, and Vaikuntanathan (GPV) and presented at STOC 2008 [8]. It uses a modified Babai nearest plane algorithm [3] and two famous results by Ajtai [1,2] in order to build a trapdoor function that is arguably “as good as” a trapdoor permutation. It’s 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.

GPV trapdoor function. The trapdoor function from [8] is defined as follows.

Parameters. Depending on the security parameter n , the other parameters in [8] are the following.

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

The above parameters will be made explicit in Section 5.

Spaces. The range is

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

and the domain is the set

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

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

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

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 orthogonalized matrix \mathbf{T} and the norm of a matrix 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 $\mathbf{y} \notin D_n$.
2. Use the trapdoor basis \mathbf{T} to sample a vector \mathbf{v} from a gaussian distribution around $-\mathbf{t}$ and output $\mathbf{x} = \mathbf{t} + \mathbf{v}$

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 for all $\mathbf{x}_1 + \mathbf{x}_2 \in D_n$:

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

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

Postprocessing. When unblinding a blind signature on a message m , the user obtains a vector

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

which is in the lattice

$$L_{H(m)} = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} \equiv H(m) \pmod{q}\}.$$

In order to make σ a valid signature, i.e. $\sigma \in D_n$, the user has to use a public basis of $L_{H(m)}$ and the short vector σ to find a slightly shorter vector $\sigma' \in L_{H(m)}$ with $\sigma' \in D_n$.

It is still an open problem, if this is attainable in all cases. As mentioned before, a different approach is to define D_n in terms of the Euclidean norm and accept signatures that are slightly larger than the bound given in [8].

5 Parameters

In this section, we analyze the choice of parameters in the GPV signature scheme and show how to apply their choice to our blind signature scheme. Then, we assess the practical efficiency of our construction.

This section will appear in the final version.

References

1. 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.
2. Miklós Ajtai. *Generating Hard Instances of the Short Basis Problem*. 1999, Lecture Notes in Computer Science, pages 1–9. Springer-Verlag, 1999.
3. László Babai. *On Lovász' lattice reduction and the nearest lattice point problem*. *Combinatorica*, 6(1):1–13, 1986.
4. 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.
5. 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.
6. 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.
7. Ran Canetti, Oded Goldreich, and Shai Halevi. *The random oracle methodology, revisited*. *J. ACM*, 51(4):557–594, 2004.
8. 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, Lecture Notes in Computer Science, pages 197–206. Springer-Verlag, 2008.

9. 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.
10. 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.
11. 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.