Enhanced Lattice-Based Signatures on Reconfigurable Hardware

Thomas Pöppelmann¹, Léo Ducas², and Tim Güneysu¹

¹ Horst Görtz Institute for IT-Security, Ruhr-University Bochum, Germany
thomas.poeppelmann@rub.de, tim.gueneysu@rub.de

² University of California, San-Diego, USA
lducas@eng.ucsd.edu

Abstract. The recent Bimodal Lattice Signature Scheme (BLISS) showed that lattice-based constructions have evolved to practical alternatives to RSA or ECC. Besides reasonably small signatures with 5600 bits for a 128-bit level of security, BLISS enables extremely fast signing and signature verification in software. However, due to the complex sampling of Gaussian noise with high precision, it is not clear whether this scheme can be mapped efficiently to embedded devices. Even though the authors of BLISS also proposed a new sampling algorithm using Bernoulli variables this approach is more complex than previous methods using large precomputed tables. The clear disadvantage of using large tables for high performance is that they cannot be used on constrained computing environments, such as FPGAs, with limited memory. In this work we thus present techniques for an efficient Cumulative Distribution Table (CDT) based Gaussian sampler on reconfigurable hardware involving Peikert’s convolution lemma and the Kullback-Leibler divergence. Based on our enhanced sampler design, we provide a first BLISS architecture for Xilinx Spartan-6 FPGAs that integrates fast FFT/NTT-based polynomial multiplication, sparse multiplication, and a Keccak hash function. Additionally, we compare the CDT with the Bernoulli approach and show that for the particular BLISS-I parameter set the improved CDT approach is faster with lower area consumption. Our core uses 2,431 slices, 7.5 BRAMs, and 6 DSPs and performs a signing operation in 126 μ s on average. Verification takes even less with 70 μ s.

Keywords: Ideal Lattices, Gaussian Sampling, Digital Signatures, FPGA.

1 Introduction and Motivation

Virtually all currently used digital signature schemes rely either on the factoring (RSA) or the discrete logarithm problem (DSA/ECDSA). However, with Shor’s algorithm [39] sufficiently large quantum computers can solve these problems in polynomial time which potentially puts billions of devices and users at risk. Although powerful quantum computers will certainly not become available soon, significant resources are definitely spent by various organizations to boost their

further development [35]. Also motivated by further advances in classical cryptanalysis (e.g., [4, 5, 20]), it is important to investigate potential alternatives now to have secure constructions and implementations at hand when they are finally needed.

In this work we deal with such a promising alternative, namely the Bimodal Lattice Signature Scheme (BLISS) [12], and specifically address implementation challenges for constrained devices and reconfigurable hardware. First efforts in this direction were made in 2012 by Güneysu et al. [16] (GLP). Their scheme was based on work by Lyubashevsky [26] and tuned for practicability and efficiency in embedded systems. This was achieved by a new signature compression mechanism, a more "aggressive", non-standard hardness assumption, and the decision to use uniform (as in [25]) instead of Gaussian noise to hide the secret key contained in each signature via rejection sampling. While GLP allows high performance on low-cost FPGAs [16] and CPUs [17] it later turned out that the scheme is suboptimal in terms of signature size and its claimed security level compared to BLISS. The main reason for this is that Gaussian noise, which is prevalent in almost all lattice-based constructions, allows more efficient, more secure, and also smaller signatures. However, while other techniques relevant for lattice-based cryptography, like fast polynomial arithmetic on ideal lattices received some attention [1, 32, 36], it is currently not clear how efficient Gaussian sampling can be done on reconfigurable and embedded hardware for large standard deviations. Results from electrical engineering (e.g., [19, 41]) are not directly applicable, as they target continuous Gaussians. Applying these algorithms for the discrete case is not trivial (see, e.g., [8] for a discrete version of the Ziggurat algorithm). First progress was recently made by Roy et al. [37] based on work by Galbraith and Dwarakanath [13] providing results for a Gaussian sampler in lattice-based encryption that requires low resources. We would also like to note that for lattice-based digital signature schemes large tables in performance optimized implementations might imply the impression that Gaussian-noise based schemes are a suboptimal choice on constrained embedded systems. A recent example is a microcontroller implementation of BLISS [7] that requires tables for the Gaussian sampler of roughly 40 to 50 KB on an ATxmega64A3. Other lattice-based signatures with explicit reductions to standard lattice problems [14, 24, 28] are also inefficient in terms of practical signature and public key sizes (see [3] for an implementation of [28]). Thus, despite the necessity of improving Gaussian sampling techniques (which is one contribution of this work) BLISS seems to be currently the most promising scheme with a signatures length of 5600 bit, equally large public keys, and 128-bit of equivalent symmetric security. There surely is some room for theoretical improvement, as suggested by the new compression ideas developed by Bai and Galbraith [2]; one can hope that all those techniques can be combined to further improve lattice-based signatures.

Contribution. One contribution of this work are improved techniques for efficient sampling of Gaussian noise that support parameters required for digital signature schemes such as BLISS and similar constructions. First, we detail how

to accelerate the binary search on a cumulative distribution table (CDT) using a shortcut table of intervals (also known as guide table [9, 11]) and develop an optimal data structure that saves roughly half of the table space by exploiting the properties of the Kullback-Leibler divergence. Furthermore, we apply a convolution lemma [29] for discrete Gaussians that results in even smaller tables of less than 2.1 KB for BLISS-I parameters. Based on these techniques we provide an implementation of the BLISS-I parameter set on reconfigurable hardware that is tweaked for performance and offers 128-bit of security. For practical evaluation we compare our improvements for the CDT-based Gaussian sampler to the Bernoulli approach presented in [12]. Our implementation includes an FFT/NTT-based polynomial multiplier (contrary to the schoolbook approach from [16]), more efficient sparse multiplication, and the KECCAK- f [1600] hash function to provide the full picture of the performance that can be achieved by employing latest lattice-based signature schemes on reconfigurable hardware. Our implementation on a Xilinx Spartan-6 FPGA supports up to 7958 signatures per second using 7,491 LUTs, 7,033 flip-flops, 6 DSPs, and 7.5 block RAMs and outperforms previous work [16] both in time and area.

In order to allow third-party evaluation of our results, source code, test-benches, and documentation is available on our website¹.

2 The Bimodal Lattice Signature Scheme

The most efficient instantiation of the BLISS signature scheme [12] is based on ideal-lattices [27] and operates on polynomials over the ring $\mathcal{R}_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$. For quick reference, the BLISS key generation, signing as well as verification algorithms are given in Figure 1 and implementation relevant parameters as well as achievable signature and key sizes are listed in Table 1. Note that for the remainder of this work, we will focus solely on BLISS-I. The BLISS key generation basically involves uniform sampling of two small and sparse polynomials $\mathbf{f}, \mathbf{g}$, computation of a certain rejection condition ($N_\kappa(\mathbf{S})$), and computation of an inverse. For signature generation two polynomials $\mathbf{y}_1, \mathbf{y}_2$ of length n are sampled from a discrete Gaussian distribution with standard deviation σ . Note that the computation of $\mathbf{a}\mathbf{y}_1$ can still be performed in the FFT-enabled ring $\mathcal{R}_q$ instead of $\mathcal{R}_{2q}$. The result $\mathbf{u}$ is then hashed with the message μ . The output of the hash function is interpreted as sparse polynomial $\mathbf{c}$. The polynomials $\mathbf{y}_{1,2}$ are then used to mask the secret key polynomials $\mathbf{s}_{1,2}$ which are multiplied with the polynomial $\mathbf{c}$ and thus "sign" the hash of the message. In order to prevent any leakage of information on the secret key, rejection sampling is performed and signing might restart. Finally, the signature is compressed and $(\mathbf{z}_1, \mathbf{z}_2^\dagger, \mathbf{c})$ returned. For verification the norms of the signature are first validated, then the input to the hash function is reconstructed and it is checked whether the corresponding hash output matches $\mathbf{c}$ from the signature.

¹ See <http://www.sha.rub.de/research/projects/lattice/>

Algorithm KeyGen()

1. Choose $\mathbf{f}, \mathbf{g}$ as uniform polynomials with exactly $d_1 = \lceil \delta_1 n \rceil$ entries in $\{\pm 1\}$ and $d_2 = \lceil \delta_2 n \rceil$ entries in $\{\pm 2\}$
2. $\mathbf{S} = (\mathbf{s}_1, \mathbf{s}_2)^t \leftarrow (\mathbf{f}, 2\mathbf{g} + 1)^t$
3. **if** $N_\kappa(\mathbf{S}) \geq C^2 \cdot 5 \cdot (\lceil \delta_1 n \rceil + 4\lceil \delta_2 n \rceil) \cdot \kappa$ **then restart**
4. $\mathbf{a}_q = (2\mathbf{g} + 1)/\mathbf{f} \bmod q$ (**restart** if $\mathbf{f}$ is not invertible)
5. **Return** ($pk = \mathbf{A}, sk = \mathbf{S}$) where $\mathbf{A} = (\mathbf{a}_1 = 2\mathbf{a}_q, q - 2) \bmod 2q$

Alg. Sign($\mu, pk = \mathbf{A}, sk = \mathbf{S}$)

1. $\mathbf{y}_1, \mathbf{y}_2 \leftarrow D_{\mathbb{Z}^n, \sigma}$
2. $\mathbf{u} = \zeta \cdot \mathbf{a}_1 \cdot \mathbf{y}_1 + \mathbf{y}_2 \bmod 2q$
3. $\mathbf{c} \leftarrow H(\lfloor \mathbf{u} \rfloor_d \bmod p, \mu)$
4. Choose a random bit b
5. $\mathbf{z}_1 \leftarrow \mathbf{y}_1 + (-1)^b \mathbf{s}_1 \mathbf{c}$
6. $\mathbf{z}_2 \leftarrow \mathbf{y}_2 + (-1)^b \mathbf{s}_2 \mathbf{c}$
7. **Continue** with probability $1 / \left(M \exp\left(-\frac{\|\mathbf{Sc}\|^2}{2\sigma^2}\right) \cosh\left(\frac{\langle \mathbf{z}, \mathbf{Sc} \rangle}{\sigma^2}\right) \right)$
otherwise **restart**
8. $\mathbf{z}_2^\dagger \leftarrow (\lfloor \mathbf{u} \rfloor_d - \lfloor \mathbf{u} - \mathbf{z}_2 \rfloor_d) \bmod p$
9. **Return** $(\mathbf{z}_1, \mathbf{z}_2^\dagger, \mathbf{c})$

Alg. Verify($\mu, pk = \mathbf{A}, (\mathbf{z}_1, \mathbf{z}_2^\dagger, \mathbf{c})$)

1. **if** $\|(\mathbf{z}_1|^{2d} \cdot \mathbf{z}_2^\dagger)\|_2 > B_2$ **then Reject**
2. **if** $\|(\mathbf{z}_1|^{2d} \cdot \mathbf{z}_2^\dagger)\|_\infty > B_\infty$ **then Reject**
3. **Accept** iff $\mathbf{c} = H(\lfloor \zeta \cdot \mathbf{a}_1 \cdot \mathbf{z}_1 + \zeta \cdot q \cdot \mathbf{c} \rfloor_d + \mathbf{z}_2^\dagger \bmod p, \mu)$

Fig. 1. The Bimodal Lattice Signature Scheme [12]

3 Improving Gaussian Sampling for Lattice-Based Digital Signatures

Target distribution. We recall that the centered discrete Gaussian distribution $D_{\mathbb{Z}, \sigma}$ is defined by a weight proportional to $\rho_\sigma(x) = \exp(-\frac{x^2}{2\sigma^2})$ for all integers x . Our goal is to efficiently sample from that distribution for a constant value $\sigma \approx 215.73$ as specified in BLISS-I (precisely $\sigma = 254 \cdot \sigma_{\text{bin}}$ where $\sigma_{\text{bin}} = \sqrt{1/(2 \ln 2)}$ is the parameter of the so-called binary-Gaussian; see [12]). This can easily be reduced to sampling from a distribution over $\mathbb{Z}^+$ proportional to $\rho(x)$ for all $x > 0$ and to $\rho(0)/2$ for $x = 0$.

Overview. Gaussian sampling using a large cumulative distribution table (CDT) has been shown to be an efficient strategy for the software implementation of BLISS given in [12]. In this section, we further enhance CDT-based Gaussian sampling for use on constrained devices. For simplicity, we explicitly refer to the parameter set BLISS-I although we remark that our enhancements can be transferred to any other parameter set as well. To increase performance, we first analyze and improve the binary search step to reduce the number of comparisons (cf. Section 3.1). Secondly, we decrease the size of the precomputed tables. In Section 3.3 we therefore apply a convolution lemma for discrete Gaussians adapted from [30] that enables the use of a sampler with much smaller standard deviation $\sigma' \approx \sigma/11$, reducing the table size by a factor 11. In Section 3.4 we

Table 1. Parameters proposals from [12]

Name of the scheme	BLISS-I	BLISS-II	BLISS-III	BLISS-IV
Security	128 bits	128 bits	160 bits	192 bits
(n, q)	(512,12289)	(512,12289)	(512,12289)	(512,12289)
Secret key densities δ_1, δ_2	0.3 , 0	0.3 , 0	0.42 , 0.03	0.45, 0.06
Gaussian std. dev. σ	215.73	107.86	250.54	271.93
Weight of the challenge κ	23	23	30	39
Verif. thresholds B_2, B_∞	12872, 2100	11074, 1563	10206,1760	9901, 1613
Repetition rate	1.6	7.4	2.8	5.2
Signature size	5.6kb	5kb	6kb	6.5kb
Secret key size	2kb	2kb	3kb	3kb
Public key size	7kb	7kb	7kb	7kb

finally reduce the size of the precomputed table further by roughly a factor of two using floating-point representation by introducing an *adaptive mantissa size*.

For those last two steps we require the “measure of distance”² for a distribution, called Kullback-Leibler divergence [10, 23], that offers tighter proofs than the usual statistical distance (cf. Section 3.2). Kullback-Leibler is a standard notion in information theory and already played a role in cryptography, mostly in the context of symmetric cryptanalysis [6, 42].

3.1 Binary Search with Shortcut Intervals

The CDT sampling algorithm uses a table $0 = T[0] \leq T[i] \leq \dots \leq T[S+1] = 1$ to sample from a uniform real $r \in [0, 1)$. The output x is the unique index satisfying $T[x] \leq r < T[x+1]$ and it is obtain via a binary search. Each output $x \in \{0 \dots S\}$ has a probability $T[x+1] - T[x]$. For BLISS-I we need a table with $S = 2891 \approx 13.4\sigma$ entries to dismiss only a portion of the tail less than 2^{-128} . As a result, the naive binary search would require $C \in [\lceil \log_2 S \rceil, \lceil \log_2 S \rceil] = [11, 12]$ comparisons on average.

As an improvement we propose to combine the binary search with a hash map based on the first bits of r to narrow down the search interval in a first step (an idea that is not exactly new [9, 11], also known as guide tables). For the given parameters and memory alignment reasons, we choose the first byte of r for this hash map: the unique $v \in \{0 \dots 255\}$ such that $v/256 \leq r < (v+1)/256$. This table I of intervals has length 256 and each entry $I[v]$ encodes the smallest interval (a_v, b_v) such that $T[a_v] \leq v/256$ and $T[b_v] \geq (v+1)/256$. With this approach, the search can be directly reduced to the interval (a_v, b_v) . By letting C denote the number of comparison on average, we have that $\sum_v \frac{\lceil \log_2(b_v - a_v) \rceil}{256} \leq C \leq \sum_v \frac{\lceil \log_2(b_v - a_v) \rceil}{256}$. For this distribution this would give $C \in [1.3, 1.7]$ comparisons on average.

² Technically, Kullback-Leibler divergence is not a distance; it is not even symmetric.

3.2 Preliminaries on the Kullback-Leibler Divergence

We now present the notion of Kullback-Leibler (KL) divergence that is later used to further reduce the table size. Detailed proofs of following lemmata are given in the full version [31].

Definition 1 (Kullback-Leibler Divergence). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions over a common countable set Ω , and let $S \subset \Omega$ be the strict support of $\mathcal{P}$ ($\mathcal{P}(i) > 0$ iff $i \in S$). The Kullback-Leibler divergence, noted D_{KL} of $\mathcal{Q}$ from $\mathcal{P}$ is defined as:*

$$D_{KL}(\mathcal{P} \parallel \mathcal{Q}) = \sum_{i \in S} \ln \left(\frac{\mathcal{P}(i)}{\mathcal{Q}(i)} \right) \mathcal{P}(i)$$

with the convention that $\ln(x/0) = +\infty$ for any $x > 0$.

The Kullback-Leibler divergence shares many useful properties with the more usual notion of statistical distance. First, it is additive so that $D_{KL}(\mathcal{P}_0 \times \mathcal{P}_1 \parallel \mathcal{Q}_0 \times \mathcal{Q}_1) = D_{KL}(\mathcal{P}_0 \parallel \mathcal{Q}_0) + D_{KL}(\mathcal{P}_1 \parallel \mathcal{Q}_1)$ and, second, non-increasing under any function $D_{KL}(f(\mathcal{P}) \parallel f(\mathcal{Q})) \leq D_{KL}(\mathcal{P} \parallel \mathcal{Q})$. An important difference though is that it is not symmetric. Choosing parameters so that the theoretical distribution $\mathcal{Q}$ is at KL-divergence about 2^{-128} from the actually sampled distribution $\mathcal{P}$, the next lemma will let us conclude the following³: if the ideal scheme $\mathcal{S}^{\mathcal{Q}}$ (i.e. BLISS with a perfect sampler) has about 128 bits of security, so has the implemented scheme $\mathcal{S}^{\mathcal{P}}$ (i.e. BLISS with our imperfect sampler).

Lemma 1 (Bounding Success Probability Variations). *Let $\mathcal{E}^{\mathcal{P}}$ be an algorithm making at most q queries to an oracle sampling from a distribution $\mathcal{P}$ and returning a bit. Let $\epsilon \geq 0$, and $\mathcal{Q}$ be a distribution such that $D_{KL}(\mathcal{P} \parallel \mathcal{Q}) \leq \epsilon$. Let x (resp. y) denote the probability that $\mathcal{E}^{\mathcal{P}}$ (resp. $\mathcal{E}^{\mathcal{Q}}$) outputs 1. Then, $|x - y| \leq \sqrt{q\epsilon/2}$.*

In certain cases, the KL-divergence can be as small as the square of the statistical distance. For example, noting $\mathcal{B}_c$ the Bernoulli variable that returns 1 with probability c , we have $D_{KL}(\mathcal{B}_{\frac{1-\epsilon}{2}} \parallel \mathcal{B}_{\frac{1}{2}}) \approx \epsilon^2/2$. In such a case, one requires $q = O(1/\epsilon^2)$ samples to distinguish those two distribution with constant advantage. Hence, we yield higher security using KL-divergence than statistical distance for which the typical argument would only prove security up to $q = O(1/\epsilon)$ queries. Intuitively, statistical distance is the sum of absolute errors, while KL-divergence is about the sum of squared relative errors.

Lemma 2 (Kullback-Leibler divergence for bounded relative error). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions of same countable support. Assume that for any $i \in S$, there exists some $\delta(i) \in (0, 1/4)$ such that we have the relative error bound $|\mathcal{P}(i) - \mathcal{Q}(i)| \leq \delta(i)\mathcal{P}(i)$. Then*

$$D_{KL}(\mathcal{P} \parallel \mathcal{Q}) \leq 2 \sum_{i \in S} \delta(i)^2 \mathcal{P}(i).$$

³ Apply the lemma to an attacker with success probability $3/4$ against $\mathcal{S}^{\mathcal{P}}$ and number of queries $< 2^{127}$ (amplifying success probability by repeating the attack if necessary), and deduce that it also succeeds against $\mathcal{S}^{\mathcal{Q}}$ with probability at least $1/4$.

Using floating-point representation, it seems now possible to halve the storage ensuring a relative precision of 64 bits instead of an absolute precision of 128 bits. Indeed, storing data with slightly more than of relative 64 bits of precision (that is, mantissa of 64 bits in floating-point format) one can reasonably hope to obtain relative errors $\delta(i) \leq 2^{-64}$ resulting in a KL-divergence less than 2^{-128} . We further exploit this idea in Section 3.4. But first, we will also use KL-divergence to improve the convolution Lemma of Peikert [30] and construct a sampler using convolutions.

3.3 Reducing Precomputed Data by Gaussian Convolution

Given that x_1, x_2 are variables from continuous Gaussian distributions with variances σ_1^2, σ_2^2 , then their combination $x_1 + cx_2$ is Gaussian with variance $\sigma_1^2 + c^2\sigma_2^2$ for any c . While this is not generally the case for discrete Gaussians, there exists similar convolution properties under some smoothing condition as proved in [29, 30]. Yet those lemmata were designed with asymptotic security in mind; for practical purpose it is in fact possible to improve the $O(\epsilon)$ statistical distance bound to a $O(\epsilon^2)$ KL-divergence bound. We refer to [30] for the formal definition of the smoothing parameter η ; for our purpose it only matters that $\eta_\epsilon(\mathbb{Z}) \leq \sqrt{\ln(2 + 2/\epsilon)/\pi}$ and thus our adapted lemma allows to decrease the smoothing condition by a factor of about $\sqrt{2}$.

Lemma 3 (Adapted from Thm. 3.1 from [30]). *Let $x_1 \leftarrow D_{\mathbb{Z}, \sigma_1}$, $x_2 \leftarrow D_{k\mathbb{Z}, \sigma_2}$ for some positive reals σ_1, σ_2 and let $\sigma_3^{-2} = \sigma_1^{-2} + \sigma_2^{-2}$, and $\sigma^2 = \sigma_1^2 + \sigma_2^2$. For any $\epsilon \in (0, 1/2)$ if $\sigma_1 \geq \eta_\epsilon(\mathbb{Z})/\sqrt{2\pi}$ and $\sigma_3 \geq \eta_\epsilon(k\mathbb{Z})/\sqrt{2\pi}$, then distribution $\mathcal{P}$ of $x_1 + x_2$ verifies*

$$D_{KL}(\mathcal{P} \| D_{\mathbb{Z}, \sigma}) \leq 2 \left(1 - \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^2 \right)^2 \approx 32\epsilon^2.$$

Remark. The factor $1/\sqrt{2\pi}$ in our version of this lemma is due to the fact that we use the standard deviation σ as the parameter of Gaussians and not the renormalized parameter $s = \sqrt{2\pi}\sigma$ often found in the literature.

Proof. The proof is similar to the one of [30], with $\Lambda_1 = \mathbb{Z}$, $\Lambda_2 = k\mathbb{Z}$, $\mathbf{c}_1 = \mathbf{c}_2 = \mathbf{0}$; but for the last argument of the proof where we replace statistical distance by KL-divergence. As in [30], we first establish that for any $\bar{x} \in \mathbb{Z}$ one has the following relative error bound

$$\mathbb{P}_{x \leftarrow \mathcal{P}}[x = \bar{x}] \in \left[\left(\frac{1 - \epsilon}{1 + \epsilon} \right)^2, \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^2 \right] \cdot \mathbb{P}_{x \leftarrow D_{\mathbb{Z}, \sigma}}[x = \bar{x}].$$

It remains to conclude using Lemma 2. □

To exploit this lemma, for BLISS-I we set $k = 11$, $\sigma' = \sigma/\sqrt{1 + k^2} \approx 19.53$, and sample $x = x_1 + kx'_2$ for $x_1, x'_2 \leftarrow D_{\mathbb{Z}, \sigma'}$ (equivalently $k \cdot x'_2 = x_2 \leftarrow D_{k\mathbb{Z}, k\sigma'}$). The smoothness conditions are verified for $\epsilon = \sqrt{2^{-128}/32}$ and $\eta_\epsilon(\mathbb{Z}) \leq 3.92$.

Due to usage of the much smaller σ' instead of σ the size of the precomputation table reduces by a factor of about $k = 11$ at the price of sampling twice. However, the running time does not double in practice since the enhancement based on the shortcut intervals reduces the number of necessary comparisons to $C \in [0.22, 0.25]$ on average. For a majority of first bytes v the interval length $b_v - a_v$ is reduced to 1 and x is determined without any comparison.

Asymptotics cost. If one considers the asymptotic costs in σ our methods allow one to sample using a table size of $\Theta(\sqrt{\sigma})$ rather than $\Theta(\sigma)$ by doubling the computation time. Actually, for much larger σ one could use $O(\log \sigma)$ samples of constant standard deviation and thus achieve a table size of $O(1)$ for computational cost in $O(\log \sigma)$.

3.4 CDT Sampling with Reduced Table Size

We recall that when doing floating-point error analysis, the relative error of a computed value v is defined as $|v - v_e|/v_e$ where v_e is the exact value that was meant to be computed. Using the table $0 = T[0] \leq T[i] \leq \dots \leq T[S+1] = 1$, the output of a CDT sampler follows the distribution $\mathcal{P}$ with $\mathcal{P}(i) = T[i+1] - T[i]$. When applying the results from KL-divergence obtained above, the relative error of $T[i+1] - T[i]$ might be significantly larger than the one of $T[i]$. This is particularly true for the tail, where $T[i] \approx 1$ but $\mathcal{P}(i)$ is very small. Intuitively, we would like the smallest probability to come first in the CDT. A simple workaround is to reverse the order of the table so that $1 = T[0] \geq T[i] \geq \dots \geq T[S+1] = 0$ with a slight modification of the algorithm so that $\mathcal{P}(i) = T[i] - T[i+1]$. With this trick, the subtraction only increase the relative error by a factor roughly σ . Indeed, leaving aside the details relative to discrete Gaussian, for $x \geq 0$ we have

$$\int_{y=x}^{\infty} \rho_s(y) dy / \rho_s(x) \leq \sigma \quad \text{whereas} \quad \int_{y=0}^x \rho_s(y) dy / \rho_s(x) \xrightarrow{x \rightarrow \infty} +\infty.$$

The left term is an estimation of the relative-error blow-up induced by the subtraction with the CDT in the reverse order and the right term the same estimation for the CDT in the natural order. We aim to have a variable precision in the table $T[i]$ so that $\delta(i)^2 \mathcal{P}(i)$ is about constant around $2^{-128}/|S|$ as suggested by Lemma 2 while $\delta(i)$ denotes the relative error $\delta(i) = |\mathcal{P}(i) - \mathcal{Q}(i)|/\mathcal{P}(i)$. As a trade-off between optimal variable precision and hardware efficiency, we propose the following data-structure. We define 9 tables $M_0 \dots M_8$ of bytes for the mantissa with respective lengths $\ell_0 \geq \ell_1 \geq \dots \geq \ell_8$ and another byte table E for exponents, of length ℓ_0 . The value $T[i]$ is defined as

$$T[i] = 256^{-E[i]} \cdot \sum_{k=0}^8 256^{-(k+1)} \cdot M_k[i]$$

where $M_k[i]$ is defined as 0 when the index is out of bound $i \geq \ell_k$. Thus, the value of $T[i]$ is stored with $p(i) = 9 - \min\{k | \ell_k > i\}$ bytes of precisions. More precisely,

Fig. 2. Data of our optimized CDT sampler for a discrete Gaussian of parameter $\sigma' \approx 19.53$

lengths are defined as $[\ell_0, \dots, \ell_8] = [262, 262, 235, 223, 202, 180, 157, 125, 86]$ so that we store at least two bytes for each entry up to $i < 262$, three bytes up to $i < 213$ and so forth. Note that no actual computation is involved in constructing $T[i]$ following the plain CDT algorithm.

For evaluation, we used the closed formula for KL-divergence and measured $D_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \leq 2^{-128}$. The storage requirements of this table is computed by $2\ell_0 + \ell_1 + \dots + \ell_8 \approx 2.1$ KB. The straightforward CDF approach requires each entry up to $i < 262$ to be stored with $128 + \log_2 \sigma$ bits of precisions and thus requires a total of at least 4.4 KB. The storage requirements are graphically depicted by the area under the curves in the top-right quadrant of Figure 2.

4 Implementation on Reconfigurable Hardware

In this section we provide details on our implementation of the BLISS-I signature scheme on a Xilinx Spartan-6 FPGA. We include the enhancements from the previous section to achieve a design that is tweaked for high-performance at moderate resource costs. For details on the implementation of the Bernoulli sampler proposed in [12] we refer to the full version [31].

4.1 Enhanced CDT Sampling.

Along the lines of the previous section our hardware implementation operates on bytes in order to use the 1024x8-bit mode of operation of the Spartan-6

Fig. 3. Block diagram of the CDT sampler which generates two samples x'_1, x'_2 of standard deviation $\sigma' \approx 19.53$ which are combined to a sample $x = x'_1 + 11x'_2$ with standard deviation $\sigma = 215.73$. The sampling is performed using binary search on the size optimized Table T .

block RAMs. The design of our CDT sampler is depicted in Figure 3 and uses the aforementioned convolution lemma. Thus two samples with $\sigma' \approx 19.53$ are combined into a sample with standard deviation $\sigma \approx 215.73$. The **BinSearch** component performs a binary search on the table T as described in Section 3.4 for a random byte vector r to find a c such that $T[c] \geq r > T[c+1]$. It accesses T byte-wise and thus $T_j[i] = M_{j-E[i]}[i]$ denotes the entry at index $i \in (0, 261)$ and byte j where $T_j[i] = 0$ when $j - E[i] < 0$ or $i \geq \ell_{j-E[i]}$. When a sampling operation is started in the **BinSearch** component we set $j = 0$ and initialize the pointer registers **min** and **max** with the values stored in the reverse interval table $I[r_0]$ where r_0 is the first random byte. The reverse interval table is realized as 256×15 -bit single port distributed ROM (6 bits for the minimum and 9 bits for the maximum). The index of the middle element of the search radius is $i = (\text{min} + \text{max})/2$. In case $T_j[i] > r_j$ we set $(\text{min} = i, i = (i + \text{max})/2, \text{max} = \text{max}, j = 0)$. Otherwise, for $T_j[i] < r_j$ we set $(i = (\text{min} + i)/2, \text{min} = \text{min}, \text{max} = i, j = 0)$ until $\text{max} - \text{min} < 2$. In case of $T_j[i] = r_j$ we increase $j = j + 1$ and thus compare the next byte. The actual entries of $M_0 \dots M_8$ are consecutively stored in block memory B and the address is computed as $a = S[j - E[i] + i]$ where we store the start addresses of each byte group in a small additional LUT-based table $S = [0, 262, 524, 759, 982, 1184, 1364, 1521, 1646]$. Some control logic takes care that all invalid/out of bound requests to S and B return a zero.

For random byte generation we use three instantiations of the Trivium stream cipher (each **Trivium** instantiation outputs one bit per clock cycle) to generate a uniformly random byte every third clock cycle and store spare bits in a **LIFO** for later use as sign bits. The random values r_j are stored in a 128×8 bit ring buffer realized as simple dual-port distributed RAM. The idea is that the sampler may request a large number of random bytes in the worst-case but usually finishes after one or two comparisons due to the lazy search. As the **BinSearch** component keeps track of the maximum number of accessed random bytes, it allows the **Uniform** sampler to refresh only the used $\text{max}(j) + 1$ bytes in the buffer. In case the buffer is empty, we stop the Gaussian sampler until a sufficient amount

of randomness becomes available. In order to compute the final sample x we determine sign bits of two samples x'_1, x'_2 and finally output $x = x'_1 + 11x'_2$.

To achieve a high clock frequency, a comparison in the binary search step could not be performed in one cycle due to the excessive number of tables and range checks involved. We therefore allow two cycles per search step which are carefully balanced. For example, we precompute the indices $i' = (\text{min} + i)/2$ and $i'' = (\text{max} + i)/2$ in the cycle prior to a comparison to relax the critical paths. We further merged the block memory B (port A) and the exponent table E (port B) into one 18k block memory and optimized the memory alignment accordingly. Note also that we are still accessing the two ports of the block RAM holding B and E only every two clock cycles which would enable another sampler to operate on the same table using time-multiplexing.

4.2 Signing and Verification Architecture

The architecture of our implementation of a high-speed BLISS signing engine is given in Figure 4. Similar to the GLP design [16] we implemented a two stage pipeline where the polynomial multiplication $\mathbf{a}_1 \mathbf{y}_1$ runs in parallel to the hashing $H(\lfloor \mathbf{u} \rfloor_d, \mu)$ and sparse multiplication $\mathbf{z}_{1,2} = \mathbf{s}_{1,2} \mathbf{c} + \mathbf{y}_{1,2}$ ⁴. For polynomial multiplication [1, 32, 36] of $\mathbf{a}_1 \mathbf{y}_1$ we rely on a publicly available FFT/NTT-based polynomial multiplier [33] (**PolyMul**). The public key $\mathbf{a}_1$ is stored already in NTT format so that only one forward and one backward transform is required. The multiplier also instantiates either the Bernoulli or the CDT Gaussian sampler (configurable by a VHDL **generic**) and an intermediate FIFO for buffering. When a new triple $(\mathbf{a}_1 \mathbf{y}_1, \mathbf{y}_1, \mathbf{y}_2)$ is available the data is transferred into the block memories **BRAM-U**, **BRAM-Y1** and **BRAM-Y2** and the small polynomial $\mathbf{u} = \zeta \mathbf{a}_1 \mathbf{y}_1 + \mathbf{y}_2$ is computed on-the-fly and stored in **BRAM-U** for later use. The lower order bits $\lfloor \mathbf{u} \rfloor_d \bmod p$ of $\mathbf{u}$ are saved in the **RAM-U**. As random oracle we have chosen the KECCAK- $f[1600]$ hash function for its security and speed in hardware [22, 38]. A configurable hardware implementation⁵ is provided by the KECCAK project and the **mid-range** core is parametrized so that the KECCAK state is split into 16 pieces ($Nb = 16$). To simplify control logic and padding we just hash multiples of 1024 bit blocks and rehash in case of a rejection. Storing the state of the hash function after hashing the message (and before hashing $\lfloor \mathbf{u} \rfloor_d \bmod p$) would be possible but is not done due to the state size of KECCAK. After hashing the **ExtractPos** component extracts the κ positions of $\mathbf{c}$ which are one from the binary hash output and stores them in the 23x9-bit memory **RAM-Pos**.

For the computation of $\mathbf{s}_1 \mathbf{c}$ and $\mathbf{s}_2 \mathbf{c}$ we then exploited that $\mathbf{c}$ has mainly zero coefficients and only $\kappa = 23$ coefficients set to one. Moreover, only $d_1 = \lceil \delta_1 n \rceil = 154$

⁴ Another option would be a three stage pipeline with an additional buffer between the hashing and sparse multiplication. As a tradeoff this would allow to use a slower and thus more area efficient hash function but also imply a longer delay and require pipeline flushes in case of an accepted signature.

⁵ See http://keccak.nokeon.org/mid_range_hw.html for more information on the core.

coefficients in $\mathbf{s}_1$ are ± 1 and $\mathbf{s}_2$ has d_1 entries in ± 2 where the first coefficient is from $\{-1, 1, 3\}$. The simplest and, in this case, also best suited algorithm for sparse polynomial multiplication is the row- or column-wise schoolbook algorithm. While row-wise multiplication would benefit from the sparsity of $\mathbf{s}_{1,2}$ and $\mathbf{c}$, more memory accesses are necessary to add and store inner products. Since memory that has more than two ports is extremely expensive, this also prevents or at least limits efficient and configurable parallelization. As a consequence, our implementation consists of a configurable number of cores (C) which perform column-wise multiplication to compute $\mathbf{z}_1$ and $\mathbf{z}_2$, respectively. Each core stores the secret key (either $\mathbf{s}_1$ or $\mathbf{s}_2$) efficiently in a distributed RAM and accumulates inner products in a small multiply-accumulate unit (MAC). Positions of $\mathbf{c}$ are fed simultaneously into the cores. Another advantage of our approach is that we can compute the norms and scalar products for rejection sampling parallel to the sparse multiplication. In Figure 4 a configuration with $C = 2$ is shown for simplicity but our experiments show that $C = 8$ leads to an optimal trade-off between speed and resource consumption. Our verification engine uses only the **PolyMul** (without a Gaussian sampler) and the **Hash** component and is thus much more lightweight compared to signing. The polynomial $\mathbf{c}$ stored as (unordered) positions is expanded into a 512x1-bit distributed RAM and the input to the hash function is computed in a pipelined manner when **PolyMul** outputs $\mathbf{a}_1\mathbf{y}_1$.

Fig. 4. Block diagram of the implemented BLISS-I signing engine

5 Results and Comparison

In this section we discuss our results which were obtained post place-and-route (PAR) on a Spartan-6 LX25 (speed grade -3) with Xilinx ISE 14.6.

Gaussian Sampling. Detailed results on area consumption and timing of the CDT and Bernoulli Gaussian sampler designs are given in Table 2. The results

show that the enhanced CDT sampler consumes less logic resources than the Bernoulli sampler, as described in the full version [31], at the cost of one 18k block memory to store the tables E and B . This is a significant improvement in terms of storage size compared to a naive implementation without the application of the Kullback-Leibler divergence and Gaussian convolution. A standard CDT implementation would require at least $\sigma\tau\lambda = 370$ kbits (that is about 23 many 18K block Rams) for the defined parameters matching a standard deviation $\sigma = 215.73$, tailcut $\tau = 13.4$ and precision $\lambda = 128$.

Regarding randomness consumption the CDT sampler needs on average 21 bits for one sample (using two smaller samples and the convolution theorem) which are generated by three instantiations of Trivium. The Bernoulli sampler on the other hand consumes 33 bits on average, generated by two instantiations of Trivium. With respect to the averaged performance, 7.4 and 18.5 cycles are required by the CDT and the Bernoulli sampler to provide one sample, respectively.

As a consequence, by combining the convolution lemma and KL-divergence we were able to maintain the advantage of the CDT, namely high speed and relative simple implementation, but significantly reduced the memory requirements (from ≈ 23 18K block RAMs to one 18K block RAM). The convolution lemma works especially well in combination with the reverse tables as the overall table sizes shrink and thus the number of comparisons is reduced. Thus, we do not expect a CTD sampler that samples directly from standard deviation σ to be significantly faster. Additionally, larger tables would require more complex address generation which might lower the achievable clock frequency. The Bernoulli approach on the other hand does not seem as suitable for an application of the convolution lemma as the CDT. The reason is that the tables are already very small and thus a reduction would not significantly reduce the area usage.

Previous implementations of Gaussian sampling for lattice-based public key encryption can be found in [34, 37]. However, both works target a smaller standard deviation of $\sigma = 3.3$. The work of Roy et al. [37] uses the Knuth-Yao algorithm (see [13] for more details), is very area-efficient (47 slices on a Virtex-5), and consumes few randomness but requires 17 clock cycles for one sample. In [34] Bernoulli sampling is used to optimize simple rejection sampling by using Bernoulli evaluation instead of computation of $\exp()$. However, without usage of the binary Gaussian distribution (see [12]) the rejection rate is high and one sample requires 96 random bits and 144 cycles. This is acceptable for a relatively slow encryption scheme and possible due to the high output rate (one bit per cycle) of the used stream cipher but not a suitable architecture for BLISS. The discrete Ziggurat [8] performs well in software and might also profit from the techniques introduced in this work but does not seem to be a good target for a hardware implementation due to its infrequent rejection sampling operations and its costly requirement on high precision floating point arithmetic.

BLISS Operations. Results for the BLISS signing and verification engine and sub-modules can be found in Table 2 including averaged cycle counts for

successfully producing a signature. Note that the final slice, LUT, and FF counts of the signing engine cannot directly be computed as the sum of the sub modules due to cross module optimizations, timing optimization, and additional control logic between modules. One signing attempt takes roughly 10k cycles and on average 1.6 trials are necessary using the BLISS-I parameter set. To evaluate the impact of the sampler used in the design, we instantiated two signing engines of which one employs a CDT sampler and the other one two Bernoulli samplers to match the speed of the multiplier. For a similar performance of roughly 8,000 signing operations per second, the signing instance based on the Bernoulli sampler has a significantly higher resource consumption (about 470 extra slices). Due to the two pipeline stages involved, the runtime of both instances is determined by $\max(\text{Cycles}(\text{PolyMul}), \text{Cycles}(\text{Hash})) + \text{Cycles}(\text{SparseMul})$ where the rejection sampling in **Compression** is performed in parallel. Further design space exploration (e.g., evaluating the impact of a different number of parallel sparse multiplication operations or a faster configuration of KECCAK) always identified the **PolyMul** component as performance bottleneck or did not provide significant savings in resources for reduced versions. In order to further increase the clock rate it would of course also be possible to instantiate the Gaussian sampler in a separate clock domain. The verification runtime is determined by $\text{Cycles}(\text{PolyMul}) + \text{Cycles}(\text{Hash})$ as no pipelining is used and **PolyMul** is slightly faster than for signing as no Gaussian sampling is needed.

Table 2. Performance and resource consumption of the full BLISS-I signing engine using the CDT sampler or two parallel Bernoulli samplers (Ber) on the Spartan-6 LX25 for a small 1024 bit message

Configuration and Operation	Slices/LUT/FF /BRAM/DSP	MHz	Cycles	Operations per second (output)
Sign-I (CDT, C=8)	2,431/7,491/7,033/7.5/6	129	$\approx 16,210$	≈ 7958 (signature)
Sign-I (Ber, C=8)	2,960/9,029/8,562/6.5/8	131	$\approx 16,210$	$\approx 8,081$ (signature)
Ver-I	1,727/5,275/4,488/4.5/3	142	9,835	14,438 (valid/invalid)
CDT sampler	299/928/1,121/1/0	129	≈ 7.4	$\approx 17,432,432$ (sample)
Bernoulli sampler	416/1,178/1,183/0/1	138	≈ 18.5	$\approx 7,459,459$ (sample)
PolyMul (CDT)	1,138/3,259/3,242/6/1	130	9,429	13,787 ($\mathbf{a} \cdot \mathbf{y}_1$)
Hash ($Nb = 16$)	752/2,461/2,134/0/0	149	1,931	77,162 ($\mathbf{c}$)
SparseMul ($C = 1$)	64/162/125/0/0	274	15,876	17,258 ($\mathbf{c} \cdot \mathbf{s}_{1,2}$)
SparseMul ($C = 8$)	308/918/459/0/0	267	2,436	109,605 ($\mathbf{c} \cdot \mathbf{s}_{1,2}$)
SparseMul ($C = 16$)	628/1847/810/0/0	254	1,476	172,086 ($\mathbf{c} \cdot \mathbf{s}_{1,2}$)
Compression	1,230/3,851/3,049/3/0	151	-	parallel to SparseMul

Comparison. In comparison with the GLP implementation from [16], the design of this work achieves higher throughput with a lower number of block RAMs

and DSPs. The structural advantage of BLISS is a smaller polynomial modulus (GLP: $q = 8383489$ /BLISS-I: $q = 12289$), less iterations necessary for a valid signature (GLP: 7/BLISS-I: 1.6), and a higher security level (GLP: 80 bit/BLISS-I: 128 bit). Furthermore and contrary to [16], we remark that our implementation takes the area costs and timings of a hash function (KECCAK) into account. In summary, our implementation of BLISS is superior to [16] in almost all aspects.

Table 3. Signing or verification speed of comparable signature scheme implementations. The GLP implementation was measured on a Spartan-6 device, the B-163 ECDSA one on a Cyclone II and the other implementations were done on Virtex-5.

Operation	Security	Algorithm	Resources	Ops/s
GLP [sign] [16]	80	GLP	7465 LUT/ 8993 FF/ 28 DSP/ 29.5 BRAM18	931
GLP [ver] [16]	80	GLP	6225 LUT/ 6663 FF/ 8 DSP/ 15 BRAM18	998
ECDSA [sign/ver] [21]	80	Full ECDSA; B-163	15,879 LE / 8,472 FF/ 36 M4K	1063/621
RSA [sign] [40]	103	RSA-2048; private key	3237 LS/ 17 DSPs	89
ECDSA [sign] [15]	128	Full ECDSA; secp256r1	32299 LUT/FF pairs	139
ECDSA [ver] [15]	128	Full ECDSA; secp256r1	32299 LUT/FF pairs	110

In addition to that Glas et al. [15] report a vehicle-to-X communication accelerator based on an ECDSA signature over 256-bit prime fields. With respect to this, our BLISS implementation shows higher performance at less resource cost. An ECDSA implementation on a binary curve for an 80-bit security level on an Altera FPGA is given in [21] and achieves similar speeds and area consumption compared to our work. Other ECC implementations over 256-bit prime or binary fields (e.g., such as [18] on a Xilinx Virtex-4) only implement the point multiplication operation and not the full ECDSA protocol. Finally, a fast RSA-2048 core was presented for Virtex-5 devices in [40] which requires more logic/DSPs and provides significantly lower performance (11.2 ms per operation) than our lattice-based signature instance.

Acknowledgment. We express our gratitude to David Xiao and Vadim Lyubashevsky for helpful conversations. We also thank the anonymous CHES’14 reviewers for detailed comments. This research was supported in part by the DARPA PROCEED program, NSF grant CNS-1117936, German Research Foundation (DFG), DFG Research Training Group GRK 1817/1 and the German Federal Ministry of Economics and Technology (Grant 01ME12025 SecMobil). Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA, NSF, DFG, or German Federal Ministry of Economics and Technology.

References

1. Aysu, A., Patterson, C., Schaumont, P.: Low-cost and area-efficient FPGA implementations of lattice-based cryptography. In: HOST, pp. 81–86. IEEE (2013)
2. Bai, S., Galbraith, S.D.: An improved compression technique for signatures based on learning with errors. In: Benaloh, J. (ed.) CT-RSA 2014. LNCS, vol. 8366, pp. 28–47. Springer, Heidelberg (2014)
3. Bansarkhani, R.E., Buchmann, J.: Improvement and efficient implementation of a lattice-based signature scheme. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) SAC 2013. LNCS, vol. 8282, pp. 48–67. Springer, Heidelberg (2014)
4. Barbulescu, R.: Selecting polynomials for the function field sieve. Cryptology ePrint Archive, Report 2013/200 (2013), <http://eprint.iacr.org/2013/200>
5. Barbulescu, R., Gaudry, P., Joux, A., Thomé, E.: A heuristic quasi-polynomial algorithm for discrete logarithm in finite fields of small characteristic. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 1–16. Springer, Heidelberg (2014), <http://arxiv.org/abs/1306.4244>
6. Blondeau, C., Gérard, B.: On the data complexity of statistical attacks against block ciphers (full version). Cryptology ePrint Archive, Report 2009/064 (2009), <http://eprint.iacr.org/2009/064>
7. Boorghany, A., Jalili, R.: Implementation and comparison of lattice-based identification protocols on smart cards and microcontrollers. Cryptology ePrint Archive, Report 2014/078 (2014), <http://eprint.iacr.org/2014/078>
8. Buchmann, J., Cabarcas, D., Göpfert, F., Hülsing, A., Weiden, P.: Discrete zigurat: A time-memory trade-off for sampling from a Gaussian distribution over the integers. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) SAC 2013, vol. 8282, pp. 402–417. Springer, Heidelberg (2014)
9. Chen, H.-C., Asau, Y.: On generating random variates from an empirical distribution. *AIIE Transactions* 6(2), 163–166 (1974)
10. Cover, T.M., Thomas, J.: Elements of Information Theory. Wiley (1991)
11. Devroye, L.: Non-Uniform Random Variate Generation. Springer-Verlag (1986), <http://luc.devroye.org/rnbookindex.html>
12. Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal gaussians. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 40–56. Springer, Heidelberg (2013)
13. Dwarakanath, N.C., Galbraith, S.D.: Sampling from discrete Gaussians for lattice-based cryptography on a constrained device. In: Applicable Algebra in Engineering, Communication and Computing, pp. 1–22 (2014)
14. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC, Victoria, British Columbia, Canada, May 17–20, pp. 197–206. ACM Press (2008)
15. Glas, B., Sander, O., Stuckert, V., Müller-Glaser, K.D., Becker, J.: Prime field ECDSA signature processing for reconfigurable embedded systems. *Int. J. Reconfig. Comp.* (2011)
16. Güneysu, T., Lyubashevsky, V., Pöppelmann, T.: Practical lattice-based cryptography: A signature scheme for embedded systems. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 530–547. Springer, Heidelberg (2012)
17. Güneysu, T., Oder, T., Pöppelmann, T., Schwabe, P.: Software speed records for lattice-based signatures. In: Gaborit, P. (ed.) PQCrypto 2013. LNCS, vol. 7932, pp. 67–82. Springer, Heidelberg (2013)

18. Güneysu, T., Paar, C.: Ultra high performance ECC over NIST primes on commercial FPGAs. In: Oswald, E., Rohatgi, P. (eds.) CHES 2008. LNCS, vol. 5154, pp. 62–78. Springer, Heidelberg (2008)
19. Gutierrez, R., Torres-Carot, V., Valls, J.: Hardware architecture of a Gaussian noise generator based on the inversion method. *IEEE Trans. on Circuits and Systems* 59-II(8), 501–505 (2012)
20. Joux, A.: A new index calculus algorithm with complexity $l(1/4 + o(1))$ in very small characteristic. *Cryptology ePrint Archive*, Report 2013/095 (2013), <http://eprint.iacr.org/2013/095>
21. Järvinen, T.M.K., Skyttä, J.: Final project report: Cryptoprocessor for elliptic curve digital signature algorithm, ECDSA (2007), http://www.altera.com/literature/dc/2007/in_2007_dig_signature.pdf
22. Jungk, B., Apfelbeck, J.: Area-efficient FPGA implementations of the SHA-3 finalists. In: Athanas, P.M., Becker, J., Cumplido, R. (eds.) ReConFig, pp. 235–241. IEEE Computer Society (2011)
23. Kullback, S., Leibler, R.A.: On information and sufficiency. *Ann. Math. Statist.* 22(1), 79–86 (1951)
24. 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)
25. 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)
26. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Heidelberg (2012)
27. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010)
28. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer, Heidelberg (2012)
29. Micciancio, D., Peikert, C.: Hardness of SIS and LWE with small parameters. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 21–39. Springer, Heidelberg (2013)
30. Peikert, C.: An efficient and parallel gaussian sampler for lattices. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 80–97. Springer, Heidelberg (2010)
31. Pöppelmann, T., Ducas, L., Güneysu, T.: Enhanced lattice-based signatures on reconfigurable hardware. *IACR Cryptology ePrint Archive*, 2014:254 (2014)
32. Pöppelmann, T., Güneysu, T.: Towards efficient arithmetic for lattice-based cryptography on reconfigurable hardware. In: Hevia, A., Neven, G. (eds.) LatinCrypt 2012. LNCS, vol. 7533, pp. 139–158. Springer, Heidelberg (2012)
33. T. Pöppelmann and T. Güneysu. Towards practical lattice-based public-key encryption on reconfigurable hardware. T. Lange, K. Lauter, and P. Lison?ek
34. Pöppelmann, T., Güneysu, T.: Area optimization of lightweight lattice-based encryption on reconfigurable hardware. In: ISCAS (to appear, 2014), http://www.sha.rub.de/media/sh/veroeffentlichungen/2014/03/23/iscas_web_version.pdf
35. Rich, S., Gellman, B.: NSA seeks quantum computer that could crack most codes. *The Washington Post* (2013), <http://wapo.st/19DycJT>

36. Roy, S.S., Vercauteren, F., Mentens, N., Chen, D.D., Verbaauwhede, I.: Compact hardware implementation of Ring-LWE cryptosystems. *IACR Cryptology ePrint Archive*, 2013:866 (2013)
37. Roy, S.S., Vercauteren, F., Verbaauwhede, I.: High precision discrete Gaussian sampling on FPGAs. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) *SAC 2013*. LNCS, vol. 8282, pp. 383–401. Springer, Heidelberg (2014)
38. Shahid, R., Sharif, M.U., Rogawski, M., Gaj, K.: Use of embedded FPGA resources in implementations of 14 round 2 SHA-3 candidates. In: Tessier, R. (ed.) *FPT*, pp. 1–9. IEEE (2011)
39. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In: *35th FOCS*, Santa Fe, New Mexico, November 20–22, pp. 124–134. IEEE Computer Society Press (1994)
40. Suzuki, D., Matsumoto, T.: How to maximize the potential of FPGA-based DSPs for modular exponentiation. *IEICE Transactions* 94-A(1), 211–222 (2011)
41. Thomas, D.B., Luk, W., Leong, P.H.W., Villasenor, J.D.: Gaussian random number generators. *ACM Comput. Surv.* 39(4) (2007)
42. Vaudenay, S.: Decorrelation: A theory for block cipher security. *Journal of Cryptology* 16(4), 249–286 (2003)