

Enhanced Lattice-Based Signatures on Reconfigurable Hardware

Thomas Pöppelmann¹ Léo Ducas² Tim Güneysu¹

¹Horst Görtz Institute for IT-Security,
Ruhr-University Bochum, Germany

²University of California,
San-Diego, USA

September 2014, CHES Workshop

Enhanced Lattice-Based Signatures on Reconfigurable Hardware

- 1 Introduction
- 2 Algorithmic contributions
- 3 Implementation and Performances
- 4 Conclusion

Lattice Based Cryptography

Many theoretical advantages over ECC/RSA:

- Strong theoretical guarantee of hardness
- Resist known quantum algorithms
- Very versatile: PKE, Signatures, IBE, FHE ...
- Asymptotically efficient $O(n^2)$ or even $O(n \log n)$.

In practice ?

Simple and very fast PKE (NTRU-Encrypt, LWE Encryption).

Efficient signatures have been more problematic.

BLISS: An optimized Signature Scheme

Signature without Trapdoors (Fiat-Shamir transform).

[Lyu09] Fiat-Shamir with aborts ($\approx 50Kbits$)

[Lyu12] Abort Rate improved using Gaussians ($\approx 12Kbits$)

[DDLL13] Abort Rate improved using Bimodal Gaussians ($\approx 5Kbits$)

BLISS vs. ECDSA vs. RSA on Software

Scheme.	Security	Sign Size	Sign./s	Ver./s
BLISS-I	128 bits	5.5kbits	8k	33k
RSA 2048	112 bits	2kbits	0.8k	27k
RSA 4096	\geq 128 bits	4kbits	0.1k	7.5k
ECDSA 256	128 bits	512 bits	9.5k	2.5k

BLISS [DDLL13] compared to `openssl` implem. of RSA and ECDSA on x86-64.

Fiat-Shamir with aborts [Lyu09, Lyu12, DLLL13]

sk : $\mathbf{S} \in \mathbb{Z}_{2q}^{m \times k}$, short

pk : $\mathbf{A} \in \mathbb{Z}_{2q}^{n \times m}$, random
 $\mathbf{T} = \mathbf{AS} = q \mathbf{Id}$

Fiat-Shamir with aborts [Lyu09, Lyu12, DLLL13]

sk : $\mathbf{S} \in \mathbb{Z}_{2q}^{m \times k}$, short

pk : $\mathbf{A} \in \mathbb{Z}_{2q}^{n \times m}$, random
 $\mathbf{T} = \mathbf{AS} = q \mathbf{Id}$

Sample $\mathbf{y} \leftarrow D_{\mathbb{Z}, \sigma}^m$, short

$$\mathbf{w} = \mathbf{Ay} \in \mathbb{Z}^n$$



Fiat-Shamir with aborts [Lyu09, Lyu12, DLLL13]

sk : $\mathbf{S} \in \mathbb{Z}_{2q}^{m \times k}$, short

pk : $\mathbf{A} \in \mathbb{Z}_{2q}^{n \times m}$, random
 $\mathbf{T} = \mathbf{AS} = q\mathbf{Id}$

Sample $\mathbf{y} \leftarrow D_{\mathbb{Z}, \sigma}^m$, short

$$\mathbf{w} = \mathbf{Ay} \in \mathbb{Z}^n$$



Sample $\mathbf{c} \in \mathbb{Z}^k$, short



\mathbf{c}

Fiat-Shamir with aborts [Lyu09, Lyu12, DLLL13]

sk : $\mathbf{S} \in \mathbb{Z}_{2q}^{m \times k}$, short

pk : $\mathbf{A} \in \mathbb{Z}_{2q}^{n \times m}$, random
 $\mathbf{T} = \mathbf{AS} = q \mathbf{Id}$

Sample $\mathbf{y} \leftarrow D_{\mathbb{Z}, \sigma}^m$, short

$\mathbf{w} = \mathbf{Ay} \in \mathbb{Z}^n$

Sample $\mathbf{c} \in \mathbb{Z}^k$, short

$\mathbf{z} = \mathbf{y} \pm \mathbf{Sc}$

Abort with proba.

$D_{\mathbb{Z}, \sigma}^m(\mathbf{z}) / M \cdot D_{\mathbb{Z}, \sigma}^m(\mathbf{z} \pm \mathbf{Sc})$

\mathbf{z}

Rejection probability is such that $\mathbf{z} \sim D_{\mathbb{Z}, \sigma}^m$ is independent from \mathbf{S} .

Fiat-Shamir with aborts [Lyu09, Lyu12, DLLL13]

sk : $\mathbf{S} \in \mathbb{Z}_{2q}^{m \times k}$, short

pk : $\mathbf{A} \in \mathbb{Z}_{2q}^{n \times m}$, random
 $\mathbf{T} = \mathbf{AS} = q \mathbf{Id}$

Sample $\mathbf{y} \leftarrow D_{\mathbb{Z}, \sigma}^m$, short

$$\mathbf{w} = \mathbf{Ay} \in \mathbb{Z}^n$$

Sample $\mathbf{c} \in \mathbb{Z}^k$, short

$$\mathbf{z} = \mathbf{y} \pm \mathbf{Sc}$$

Abort with proba.

$$D_{\mathbb{Z}, \sigma}^m(\mathbf{z}) / M \cdot D_{\mathbb{Z}, \sigma}^m(\mathbf{z} \pm \mathbf{Sc})$$

Check that

$$\|\mathbf{z}\| \text{ is short, and } \mathbf{Az} = \mathbf{Tc} + \mathbf{w} \pmod{q}$$

Rejection probability is such that $\mathbf{z} \sim D_{\mathbb{Z}, \sigma}^m$ is independent from \mathbf{S} .

Hardware Implementation of BLISS

The two costly steps are:

Hardware Implementation of BLISS

The two costly steps are:

- Polynomial multiplications **Ay** in $\mathbb{Z}_q[X]/(\Phi_{2n}(X))$
Already addressed in [PG13] and [RVM⁺14] (next talk)

Hardware Implementation of BLISS

The two costly steps are:

- Polynomial multiplications $\mathbf{A}\mathbf{y}$ in $\mathbb{Z}_q[X]/(\Phi_{2n}(X))$
Already addressed in [PG13] and [RVM⁺14] (next talk)
- Discrete Gaussian Sampling $\mathbf{y} \leftarrow D_{\mathbb{Z},\sigma}^m$ for large σ
Needs high precision sampling (learning attacks)
 - Long Floating Points Arith. [GPV08]
 - Slow algorithms [DDLL13]
 - Large tables/trees [DG14]

Hardware Implementation of BLISS

The two costly steps are:

- Polynomial multiplications $\mathbf{A}\mathbf{y}$ in $\mathbb{Z}_q[X]/(\Phi_{2n}(X))$
Already addressed in [PG13] and [RVM⁺14] (next talk)
- Discrete Gaussian Sampling $\mathbf{y} \leftarrow D_{\mathbb{Z},\sigma}^m$ for large σ
Needs high precision sampling (learning attacks)
 - Long Floating Points Arith. [GPV08]
 - Slow algorithms [DDLL13]
 - Large tables/trees [DG14]

Most of our contributions deal with hardware implementation of **Discrete Gaussian Sampling**.

Algorithmic contributions to Gaussian sampling

We focus on the fastest method to sample Discrete Gaussian using **Cumulative Distribution Tables (CDT)**.

Operation: Binary search accelerated by guide tables

Problem: Naïve implementation would require *42KB*

Algorithmic contributions to Gaussian sampling

We focus on the fastest method to sample Discrete Gaussian using **Cumulative Distribution Tables (CDT)**.

Operation: Binary search accelerated by guide tables

Problem: Naïve implementation would require $42KB$

We introduce two new techniques:

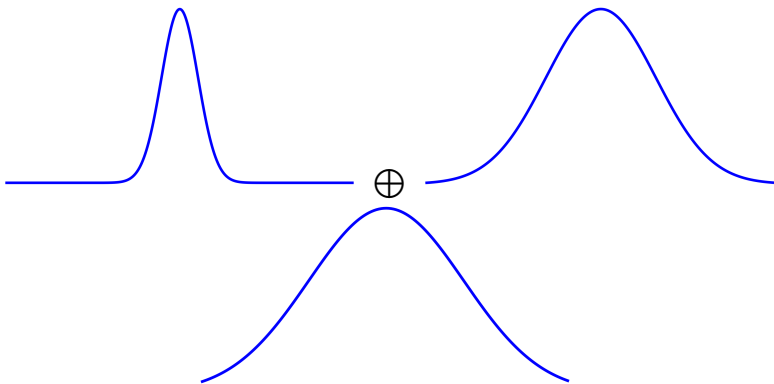
- Gaussian sampling by **convolution**
- **Kullback-Leibler-divergence** based security argument

Our algorithm requires a table of $2.1KB$, and is almost as fast

- ① Introduction
- ② Algorithmic contributions
- ③ Implementation and Performances
- ④ Conclusion


(Discrete) Gaussians convolutions

A convolution of gaussians is a gaussian.



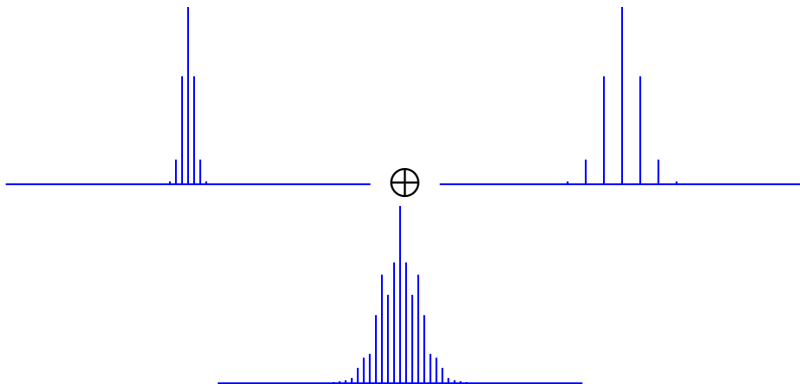
(Discrete) Gaussians convolutions

And what about discrete Gaussians ?




(Discrete) Gaussians convolutions

Well, depending on the parameter...



(Discrete) Gaussians convolutions

It may seem quite Gaussian.



Peikert's Convolution Theorem

Lemma (Adapted from [Pei10])

Let $x_1 \leftarrow D_{\mathbb{Z}, \sigma_1}$, $x_2 \leftarrow \cdot D_{\mathbb{Z}, \sigma_2}$ and set $\sigma_3^{-2} = \sigma_1^{-2} + \sigma_2^{-2}$, and $\sigma^2 = \sigma_1^2 + \sigma_2^2$. If $\sigma_1 \geq \omega(\sqrt{\log n})$ and $\sigma_3 \geq k \cdot \omega(\sqrt{\log n})$, then:

$$x_1 + kx_2 \simeq D_{\mathbb{Z}, \sigma}.$$

Application to BLISS-I: We can sample two variables x_1, x_2 of deviation $\sigma' = 19.5$ to obtain $x = x_1 + 11x_2$, of deviation $\sigma = 215$.

Impact: Size of table is reduced from 42KB to 4.5KB.

Running time is less than doubled (binary search is faster for σ').

Kullback-Leibler divergence

Definition (Kullback-Leibler Divergence)

Let \mathcal{P} and \mathcal{Q} be distribution over S . The KL divergence, of \mathcal{Q} from \mathcal{P} is defined as:


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

KL-divergence allows the same arguments as Statistical Distance:


Fact (Additivity and Data Processing inequality)

- $D_{\text{KL}}(\mathcal{P}_0 \times \mathcal{P}_1 \parallel \mathcal{Q}_0 \times \mathcal{Q}_1) = D_{\text{KL}}(\mathcal{P}_0 \parallel \mathcal{Q}_0) + D_{\text{KL}}(\mathcal{P}_1 \parallel \mathcal{Q}_1)$
- *for any function f : $D_{\text{KL}}(f(\mathcal{P}) \parallel f(\mathcal{Q})) \leq D_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q})$*


Kullback-Leibler vs. Statistical distance: Example



Kullback-Leibler vs. Statistical distance: Example



Kullback-Leibler vs. Statistical distance: Example




Statistical distance argument

$$\Delta(\mathcal{B}_{\frac{1}{2}}, \mathcal{B}_{\frac{1}{2}+\epsilon}) = \Theta(\epsilon), \quad \Delta(\mathcal{B}_0, \mathcal{B}_{\frac{1}{3}}) = \Theta(1)$$

$$n \geq \Theta(1/\epsilon)$$

Kullback-Leibler vs. Statistical distance: Example



KL-Divergence argument

$$D_{\text{KL}}(\mathcal{B}_{\frac{1}{2}} \parallel \mathcal{B}_{\frac{1}{2}+\epsilon}) = \Theta(\epsilon^2), \quad D_{\text{KL}}(\mathcal{B}_0 \parallel \mathcal{B}_{\frac{1}{3}}) = \Theta(1)$$

$$n \geq \Theta(1/\epsilon^2)$$

Kullback-Leibler vs. Statistical distance: Rule of Thumb

△ Averaged **absolute error**

$$\Delta(\mathcal{P}, \mathcal{Q}) = \frac{1}{2} \sum_i |\mathcal{P}(i) - \mathcal{Q}(i)|$$

D_{KL} Averaged squared **relative error**

$$D_{\text{KL}}(\mathcal{P} \parallel \mathcal{Q}) \leq 2 \sum_i \left| \frac{\mathcal{P}(i) - \mathcal{Q}(i)}{\mathcal{P}(i)} \right|^2 \mathcal{P}(i).$$

Limits of KL-divergence

- D_{KL} is not symmetric
- Can be worse than Δ (e.g. Tailcutting)
- Improvements only for reduction to **Search Problems**

Truncating Cumulative Distribution Table

Full CDT

```

1111110000101011100101101100001011011010001100101010000111001100111011101010000
10110001100110100001010001110011010010101110001010101101010100101111110000011001
01110001001010001100101100010001011100100100101111111101000001111011101011001101
01000000101100011101100111000110101110111110011100010101000111100010101110111011
0010000011111001111111010000001001010011101000011010101101011101111001010101111
00001110111010011000100111111100111010111000000101000100011010001001111001000011
000001011110101101001101100001000000100010110001000100010101110100111110101010
0000010000110010011101100010101011110000001100100011001101100110100011101000101
0000000010100110010010111100110101011010010110001111100100000000100100100110011
00000000010110100101010100111100000001001000001001100011000110110000110100000
000000000000101010111110110011101010101101100000011101101100100101100111110111
00000000000001000111100111011110110100010101010001000100111101100111100000001
000000000000000011010001001100111101011100000000101101011110100110101001000
000000000000000000010001000110000110100110111111001011000111000011001011101001
000000000000000000001001001101011101001100100011111100000101010011001010010100
0000000000000000000000001000111011001010010111101010100110110111111101101011
000000000000000000000000000111100010101000110110100000110011011001000011110111
0000000000000000000000000000010110001110000101110010110000110010000000111010
0000000000000000000000000000000000000000000011001000011011110101011110100010111
0000000000000000000000000000000000000000000111111110100000101010111111011001101
0000000000000000000000000000000000000000000111110001010011011101000011100111110
00000000000000000000000000000000000000000001101001011101011100110111010001
0000000000000000000000000000000000000000001001101111000011100000111011
00000000000000000000000000000000000000000001100100000000000000000000000000001100101
0000000000000000000000000000000000000000000000000000000000000000000000000011000000
000000000000000000000000000000000000000000000000000000000000000000000000001101101000110
000000000000000000000000000000000000000000000000000000000000000000000000000101111000

```

Truncating Cumulative Distribution Table

Truncating left-most zeros

```
1111110000101011100101101100001011011010001100101010000111001100111011101010000  
1011000110011010000101000111001101001010111000101010110101010010111110000011001  
01110001001010001100101100010001011100100100101111110100000111011101011001101  
01000000101100011101100111000110101110111100111000101010001111000101011011011  
001000001111100111111010000000100101001110100001101010110101110111100101010111  
0000111011101001100010011111100111010111000000101000100011010001001111001000011  
000010111101011010011011000010000000100010110001000100010101110100111110101010  
000001000011001001110110001010101110000001100100011001101100110100011101000101  
00000000101001100100101111001101010110100101100011111001000000000100100100110011  
0000000001011010010101010001111000000001001000001001100011000110110000110100000  
000000000000101010111111011001110101010111011000000111011011001000101100111110111  
000000000000010001111001110111101101000101010100010001001111011001111000000001  
000000000000000110100010011001110101111000000001011010111110011010100100100  
00000000000000000100001001011000011010011011111001011000011000011100011001001  
00000000000000000001001001101011101001100100011111100000101010011001010010100  
0000000000000000000001000111011001010010111101010100110110101111111101101011  
000000000000000000000000111100010101000110110100000110011011001000011110111  
000000000000000000000000000101100011100001011100101100001100100000000111010  
0000000000000000000000000000000011100100001101111010101111010001011110010111  
000000000000000000000000000000000011111111010000001010110111111110111001101  
0000000000000000000000000000000000111110001010011011101000011100111110  
000000000000000000000000000000000001101001011101011100110111010001  
000000000000000000000000000000000001001101111000011100000111011  
00000000000000000000000000000000000110010001000001100000110101  
00000000000000000000000000000000000001100100001000000000000011000000001010000  
00000000000000000000000000000000000001101101000110  
00000000000000000000000000000000000010111000
```

Truncating Cumulative Distribution Table

Truncating right-most bits using KL-divergence


```

111111000010101110010110110000101101101000110011001110111010110000
101100011001101000010100011100110100101011100010101010100101111110000011001
0111000100101000110010110001000101110010010011111110100000111011101011001101
0100000010110001110110011100011010111011111100111000101010110111011
001000001111100111111101000000010010100111010000110101011010111011100101010111
00001110111010011000100111111100111010111000000101000100011010001001111001000011
0000101111010110100110110000100000001000101100010001000101110100111110101010
0000010000110010011101100010101011110000001100100011001101100110100011101000101
0000000010100110010010111100110101011010010110001111100100000000100100100110011
00000000010110100101010100011110000000010010000010011000110011010000110100000
0000000000010101011111101100111010101011101100000011101101100100101100111110111
00000000000001000111100111011110110100010101010001000010011101100111100000001
000000000000000110100001001100111010111100000000010101011110000000010101001000
00000000000000000100001001011000011010011011111001011000101100011000011001001
00000000000000000001001001101011101001100100011111100000101010011001010010100
00000000000000000000010001110110010100101111010101001101101111111101101011
000000000000000000000011110001010100011011010100000110011011001000011110111
0000000000000000000000010110001110000101110010110000110010000000111010
000000000000000000000000000000000000000000000000000000000000000000000000000000011
0000000000000000000000000000000000000000000000000000000000000000000000000000000111111110100000010
000000000000000000000000000000000000000000000000000000000000000000000000000000011111000101001110
000000000000000000000000000000000000000000000000000000000000000000000000000000011010010111010
000000000000000000000000000000000000000000000000000000000000000000000000000000010011011110000111011
0000000000000000000000000000000000000000000000000000000000000000000000000000000110010000100010101
000000000000000000000000000000000000000000000000000000000000000000000000000000011000000001010000
0000000000000000000000000000000000000000000000000000000000000000000000000000000101101000110
000000000000000000000000000000000000000000000000000000000000000000000000000000010111000

```


- 1 Introduction
- 2 Algorithmic contributions
- 3 Implementation and Performances**
- 4 Conclusion

FPGA Implementation of BLISS-I : CDT Sampler



- Ring-buffer to store random numbers generated by Trivium instantiation
- Binary-search component operates on block RAM B
- **Biggest challenge:** Critical path of binary search

FPGA Implementation of BLISS-I : Signing



- Number theoretic transform (NTT) multiplier (\mathbf{ay}_1)
- Keccak-1600 hash function
- Fast sparse multiplier ($\mathbf{z}_1 = \mathbf{s}_1\mathbf{c} + \mathbf{y}_1, \mathbf{z}_2 = \mathbf{s}_2\mathbf{c} + \mathbf{y}_2$)

FPGA Implementation of BLISS-I : Results

Algorithm	LUT	FF	BRAM	DSP	OPs/s
BLISS-1[Sign]	7,491	7,033	7.5	6	7.9k signs/s
BLISS-1[Ver]	5,275	4,488	4.5	3	14,4k verifs/s
CDT-Sampler	928	1,121	1	0	17,4 M samp./s
Bernoulli Sampler	1,178	1,183	0	1	7,4 M samp./s

- Results are given for a 1024-bit message on Spartan6-LX25-3
- High-speed signing and verification
- DT sampler is twice as fast as Bernoulli sampler for similar resource consumption

FPGA Implementation of BLISS-I : Comparison

Algorithm	LUT	FF	BRAM	DSP	OPs/s
BLISS-1[Sign]	7,491	7,033	7.5	6	7,958 signs/sec
BLISS-1[Ver]	5,275	4,488	4.5	3	14,438 signs/sec
GLP-1[Sign/Ver]	6,088	6,804	19.5	4	1,627/7,438
RSA-2048 [Sign]	4190 slices		7	17	79
Curve25519	2783	3592	2	20	2518
ECDSA-256 [Sign/Ver]	32,299 LUT/FF pairs		0	0	139/110

- Implementation faster than RSA and prime curve ECC/ECDSA
- Faster, shorter and more secure than GLP lattice-signature
- Reasonable area consumption

Conclusion

Lattice-based Crypto is ready.

- BLISS compares to standardized signature schemes, in both Software and Hardware
- New geometric analysis will make it **even faster**
[D14, To appear, see Rump Session]

Time has come for **standardization** of lattice-based cryptography.

- Provide alternative/fallbacks to ECC/RSA
- Mostly **unpatented**
- Motivate more work:
Comprehensive Cryptanalysis, Improved Algorithms,
Lightweight Implementation, Side Channel Attacks, ...

Open Access, Open-Sources

Full and updated paper:

<http://eprint.iacr.org/2014/254>

Software implementation:

<http://bliss.di.ens.fr/>

Hardware implementation:

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

Thanks !



Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky.

Lattice signatures and bimodal gaussians.

In Ran Canetti and Juan A. Garay, editors, [CRYPTO 2013, Part I](#), volume 8042 of [LNCS](#), pages 40–56, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Berlin, Germany.



Nagarjun C. Dwarakanath and Steven D. Galbraith.

Sampling from discrete Gaussians for lattice-based cryptography on a constrained device.

[Applicable Algebra in Engineering, Communication and Computing](#), pages 1–22, 2014.



Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan.

Trapdoors for hard lattices and new cryptographic constructions.

In Richard E. Ladner and Cynthia Dwork, editors, [40th ACM STOC](#), pages 197–206, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press.



Vadim Lyubashevsky.

Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures.

In Mitsuru Matsui, editor, [ASIACRYPT 2009](#), volume 5912 of [LNCS](#), pages 598–616, Tokyo, Japan, December 6–10, 2009. Springer, Berlin, Germany.



Vadim Lyubashevsky.

Lattice signatures without trapdoors.

In David Pointcheval and Thomas Johansson, editors, EUROCRYPT 2012, volume 7237 of LNCS, pages 738–755, Cambridge, UK, April 15–19, 2012. Springer, Berlin, Germany.



Chris Peikert.

An efficient and parallel gaussian sampler for lattices.

In Tal Rabin, editor, CRYPTO 2010, volume 6223 of LNCS, pages 80–97, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Berlin, Germany.



Thomas Pöppelmann and Tim Güneysu.

Towards practical lattice-based public-key encryption on reconfigurable hardware.

In Tanja Lange, Kristin Lauter, and Petr Lisonek, editors, SAC 2013, volume 8282 of LNCS, pages 68–85, Burnaby, BC, Canada, August 14–16, 2013. Springer, Berlin, Germany.



Sujoy Sinha Roy, Frederik Vercauteren, Nele Mentens, Donald Donglong Chen, and Ingrid Verbauwhede.

Compact Ring-LWE based cryptoprocessor.

In CHES'14, 2014.