Sharper Bounds in Lattice-Based cryptography using the Rényi Divergence

Thomas Prest

Thales Communications & Security

Introduction

What is the Rényi divergence and why should we use it?

How to do security proofs involving distributions:

- **The standard approach:** using the statistical distance Δ.
  - Take a hard problem relying on some ideal distribution $\mathcal{Q}$,
  - Replace $\mathcal{Q}$ by a "real-life" distribution $\mathcal{P}$,
  - If $\Delta(\mathcal{P}, \mathcal{Q})$ is small enough, we win: the problem is still hard.
- **In lattice-based cryptography:** It is often relevant to replace the statistical distance with the *Rényi divergence*.
  - More efficient in many cases [LSS14, LPSS14, BLL+, BGM+],
  - But trickier to use.

Motivation of this work:

1. Formalize and optimize the use of the Rényi divergence in security proofs
   $\Rightarrow$ Section 2.
2. Five more applications of the Rényi divergence to lattice-based cryptography
   $\Rightarrow$ Section 3.

The Rényi Divergence

Definition. For $a \in (1, +\infty)$, the Rényi divergence between two distributions $\mathcal{P}, \mathcal{Q}$ is

$$R_a(\mathcal{P}\|\mathcal{Q}) = \left(\sum_{x \in Supp(\mathcal{P})} \frac{\mathcal{P}(x)^a}{\mathcal{Q}(x)^{a-1}} \right)^{\frac{1}{a-1}}$$

Motivation. We consider a cryptographic scheme doing q queries to a distribution $\mathcal{D}_i$ ($i \in \{0, 1\}$), we note ε_i the probability of an event breaking the scheme.

- With the statistical distance:

$$\varepsilon_0 \geq \varepsilon_1 - q\Delta(\mathcal{D}_1, \mathcal{D}_0) \qquad \boxed{\Delta \leq 2^{-\lambda} \Rightarrow \text{we win}}$$

- With the Rényi divergence:

$$\varepsilon_0 \geq \varepsilon_1^{\frac{a}{a-1}} / R_a(\mathcal{D}_1\|\mathcal{D}_0)^q \qquad \boxed{\log R_a \leq 2^{-q} \Rightarrow \text{we win}}$$

Observation. For "equal" values ($\log R_a \approx \Delta$), Rényi divergence is more interesting when $q \ll 2^\lambda$ [BLL+].

- Typically, $\lambda \in \{128, 192, 256\}$.
- In the NIST call for post-quantum schemes, $q = 2^{64}$.

Theory

The first and second lemmas

1. **Tailcut.** Let $\delta > 0$ such that $\frac{\mathcal{D}_\delta}{\mathcal{D}} \leq 1 + \delta$. For $a \in (1, \infty]$:
   - $R_a(\mathcal{D}_\delta || \mathcal{D}) \leq (1+\delta)^{a/a-1}$

   Example: $\mathcal{D}_\delta$ is a tailcut of $\mathcal{D}$ (discard a set S such that $\mathcal{D}(S) \leq \delta$).

2. **Relative error.** Suppose $\text{Supp}(\mathcal{D}_\delta) = \text{Supp}(\mathcal{D})$.
   Let $\delta > 0$ such that $1 - \delta \leq \frac{\mathcal{D}_\delta}{\mathcal{D}} \leq 1 + \delta$. For $a \in (1, \infty)$:
   - $R_a(\mathcal{D}_\delta || \mathcal{D}) \leq \left(1 + \frac{a(a-1)\delta^2}{2(1-\delta)^{a+1}}\right)^{\frac{1}{a-1}} \underset{\delta \to 0}{\sim} 1 + \frac{a\delta^2}{2}$

   Example: $\mathcal{D}_\delta$ implements $\mathcal{D}$ with finite precision (relative error δ).

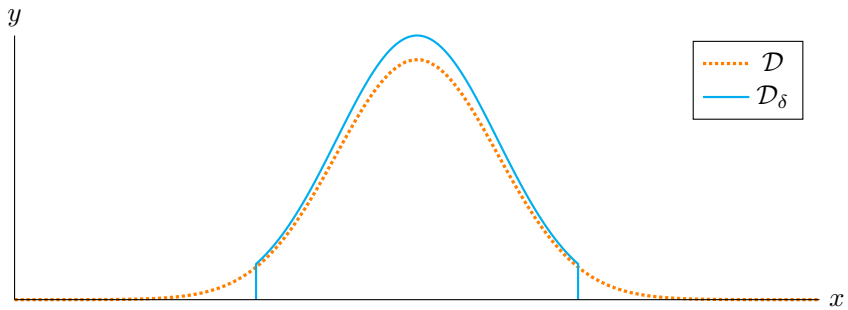

Tailcut lemma usecase

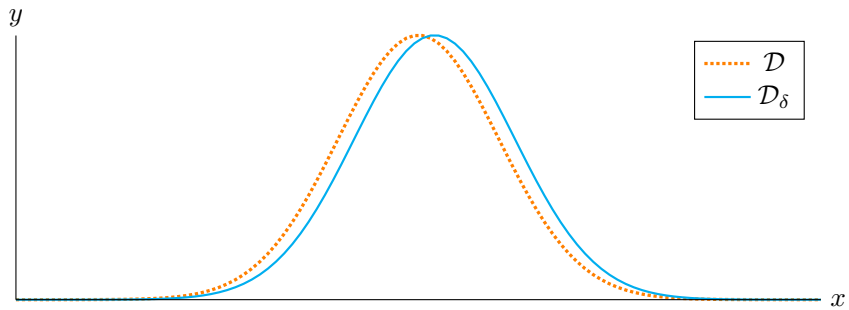

Relative error lemma usecase

The third lemma

The max-log distance. Introduced in [MW17].[1]
For two distributions $\mathcal{P}$ and $\mathcal{Q}$ over the same support S:

$$\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q}) = \max_{x \in S} |\log \mathcal{P}(x) - \log \mathcal{Q}(x)|$$

Unlike the Rényi divergence, it is a distance, so it verifies the:

- Triangle inequality: $\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{R}) \leq \Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q}) + \Delta_{\mathsf{ML}}(\mathcal{Q}, \mathcal{R})$
- Symmetry: $\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q}) = \Delta_{\mathsf{ML}}(\mathcal{Q}, \mathcal{P})$

[MW17] essentially states that $\boxed{\Delta_{\mathsf{ML}} \leq 2^{-\lambda/2} \Rightarrow \text{we win}}$.

[1] Actually similar to the differential privacy.

The third lemma

The max-log distance. Introduced in [MW17].[1]
For two distributions $\mathcal{P}$ and $\mathcal{Q}$ over the same support S:

$$\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q}) = \max_{x \in S} |\log \mathcal{P}(x) - \log \mathcal{Q}(x)|$$

Unlike the Rényi divergence, it is a distance, so it verifies the:

- Triangle inequality: $\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{R}) \leq \Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q}) + \Delta_{\mathsf{ML}}(\mathcal{Q}, \mathcal{R})$
- Symmetry: $\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q}) = \Delta_{\mathsf{ML}}(\mathcal{Q}, \mathcal{P})$

[MW17] essentially states that $\boxed{\Delta_{\mathsf{ML}} \leq 2^{-\lambda/2} \Rightarrow \text{we win}}$.

3. **A reverse Pinsker inequality.** For two distributions $\mathcal{P}, \mathcal{Q}$ of common support, we have:

$$R_a(\mathcal{P}||\mathcal{Q}) \leq \left(1 + \frac{a(a-1)(e^{\Delta_{\mathsf{ML}}(\mathcal{P},\mathcal{Q})} - 1)^2}{2(2 - e^{\Delta_{\mathsf{ML}}(\mathcal{P},\mathcal{Q})})^{a+1}}\right)^{\frac{1}{a-1}} \underset{\Delta_{\mathsf{ML}} \to 0}{\sim} 1 + \frac{a\Delta_{\mathsf{ML}}(\mathcal{P}, \mathcal{Q})^2}{2}$$

Consequence: Instead of $\boxed{\Delta_{\mathsf{ML}} \leq 2^{-\lambda/2}}$, we only need $\boxed{\Delta_{\mathsf{ML}} \leq \frac{1}{\sqrt{a}} 2^{-q/2}}$.

[1] Actually similar to the differential privacy.

Framework for using the Rényi Divergence

1. Take your favourite scheme
2. Set more aggressive parameters:
   1. First, try to apply the relative error lemma (the most powerful)
   2. Wherever it doesn't work, apply either the tailcut lemma or the reverse Pinsker's inequality

   ! Taking $R_a \leq 1 + 2^{-q}$ is sufficient.

   ! Taking $a = 2\lambda$ gives tight, efficient proofs.
3. Write an article

! *These arguments are only valid for search problems!*
For decision problems, achieving the same efficiency is still open.

! *In the rest of this presentation, we assume that we have less than* 2^{64} *queries.*

Practice

Security of a Sampler from [MW17]

Context. A new sampler over $\mathbb{Z}$ was introduced in [MW17].

Previous works. [MW17] perform a max-log distance-based analysis of the sampler. They find that

$$\text{64 bits of precision} \Rightarrow \Delta_{\mathsf{ML}} \leq 2^{-50} \Rightarrow \text{About 100 bits of security}$$

This work. We use the reverse Pinsker's inequality:

$$\begin{aligned} \text{64 bits of precision} \quad &\Rightarrow \quad \Delta_{\mathsf{ML}} \leq 2^{-50} \\ &\Rightarrow \quad R_a \leq 1 + 2^{-96} \\ &\Rightarrow \quad \text{256 bits of security, even with up to } 2^{94} \text{ queries} \end{aligned}$$

We gain this much security *for free*.
No knowledge about the sampler is required.

Revisiting the Table Approach

Context. We study the use of precomputed tables for sampling discrete distributions – typically, (pseudo)Gaussians.

Previous works. Existing approachs [Pei10, PDG14, DG14] require high precision ($\geq \lambda/2$) and/or floating-point arithmetic.

This work. We propose a simple approach which requires less than 64 bits of *fixed* precision in practice.

The classical CDF-table approach

Let $\mathcal{D}$ be a distribution over $\mathbb{N}$ that we want to sample from.
We suppose we have a precomputed table of $\mathrm{CDF}_{\mathcal{D}}$ defined over $\mathbb{N}$ by:

$$\mathrm{CDF}_{\mathcal{D}}(z) = \sum_{i \leq z} \mathcal{D}(i)$$

Algorithm 1 CDF sampler

Require: A precomputed table of $\mathrm{CDF}_{\mathcal{D}}$

$z \leftarrow 0$
$u \leftarrow [0, 1]$ uniformly
while $u \geq \mathrm{CDF}_{\mathcal{D}}(z)$ **do**
    $z \leftarrow z + 1$
Return z

Suppose we want to sample a half-Gaussian D_{σ}^{+}.

- *Statistical distance-based analysis.* We need to store about:
  - $\sigma \cdot \sqrt{2\lambda}$ values,
  - With a precision λ.
- *Rényi Divergence-based analysis.* We need to store about:
  - $\sigma \cdot \sqrt{2q}$ values,
  - With a precision λ. But we prefer/expect $\log_2(q)$ or $\log_2(q)/2$!

The CoDF sampler

Our solution. We use a "Rényi divergence-friendly" table. This requires a different algorithm. We define the conditional density function of $\mathcal{D}$ by:

$$\mathrm{CoDF}_{\mathcal{D}}(z) = \mathcal{D}(z) / \sum_{i \geq z} \mathcal{D}(i)$$

Algorithm 2 CoDF sampler

Require: A precomputed table of $\mathrm{CoDF}_{\mathcal{D}}$
Ensure: $z \leftarrow \mathcal{D}$
  $z \leftarrow 0$
  $u \leftarrow [0, 1]$ uniformly
  while $u \geq \mathrm{CoDF}_{\mathcal{D}}(z)$ **do**
    $z \leftarrow z + 1$
    $u \leftarrow [0, 1]$ uniformly
  Return z

Suppose we want to sample a half-Gaussian D_{σ}^{+}.

- *Rényi Divergence-based analysis.* We need to store about:
  - $\sigma \cdot \sqrt{2q}$ values,
  - With a precision $\log_2(q)/2$!

Example and Conclusion

A practical example: the distribution $D^{+}_{\mathbb{Z},0.85\ldots}$ from [DDLL13].

- CDF+SD approach: 20 elements of 266 bits each $\Rightarrow$ $\approx$ 5 300 bits.
- CoDF+RD approach: 11 elements of 53 bits each $\Rightarrow$ $\approx$ 600 bits.

Conclusion:

- Both in theory and practice, we gain an order of magnitude.
- Requires only standard (64 bits) fixed-point arithmetic.
- HIghly composable with other table-based techniques.

Context. Trapdoor sampling allows to sample a discrete Gaussian $D_{\Lambda(\mathbf{B}),\sigma,\mathbf{c}}$.

- Allows hash-and-sign, IBE [GPV08], standard model signatures [CHKP10, Boy10], hierarchical IBE [CHKP10, ABB10a, ABB10b], attribute-based encryption [Boy13, BGG+14] and so on.
- Current algorithms [Kle00, GPV08, Pei10, MP12, DP16] heavily rely on floating-point arithmetic.

This work. Two axes of improvement for trapdoor samplers:

1. Squeezing the standard deviation
2. Reducing the required precision

Our test subject: Klein's sampler

Algorithm 3 $\text{Klein}_{\mathbf{L},\sigma}(\mathbf{t})$

Require: $\sigma \geq \eta_\epsilon(\mathbb{Z}^n) \cdot \|\mathbf{B}\|_{\text{GS}}$, the Gram-Schmidt orthogonalization $\mathbf{B} = \mathbf{L} \cdot \tilde{\mathbf{B}}$, the values $\sigma_j = \sigma / \|\tilde{\mathbf{b}}_j\|$ and a target $\mathbf{t}$
Ensure: A vector $\mathbf{z}$ such that $\mathbf{zB} \leftarrow D_{\Lambda(\mathbf{B}),\sigma,\mathbf{tB}}$
  for $j = n, \ldots, 1$ **do**
    $c_j \leftarrow t_j + \sum_{i>j}(t_j - z_j)L_{ij}$
    $z_j \leftarrow D_{\mathbb{Z},\sigma_j,c_j}$
  return z

There are two operational constraints over the standard deviation σ:

1. σ too large $\Rightarrow$ $\text{Klein}_{\mathbf{L},\sigma}$ does not solve a hard problem and is useless in a cryptographic context.
2. σ too small $\Rightarrow$ $\text{Klein}_{\mathbf{L},\sigma}$ does not behave like a perfect Gaussian oracle anymore.

So the standard deviation must be small but the output of $\text{Klein}_{\mathbf{L},\sigma}$ must still look like a Gaussian distribution.

Trapdoor Sampling

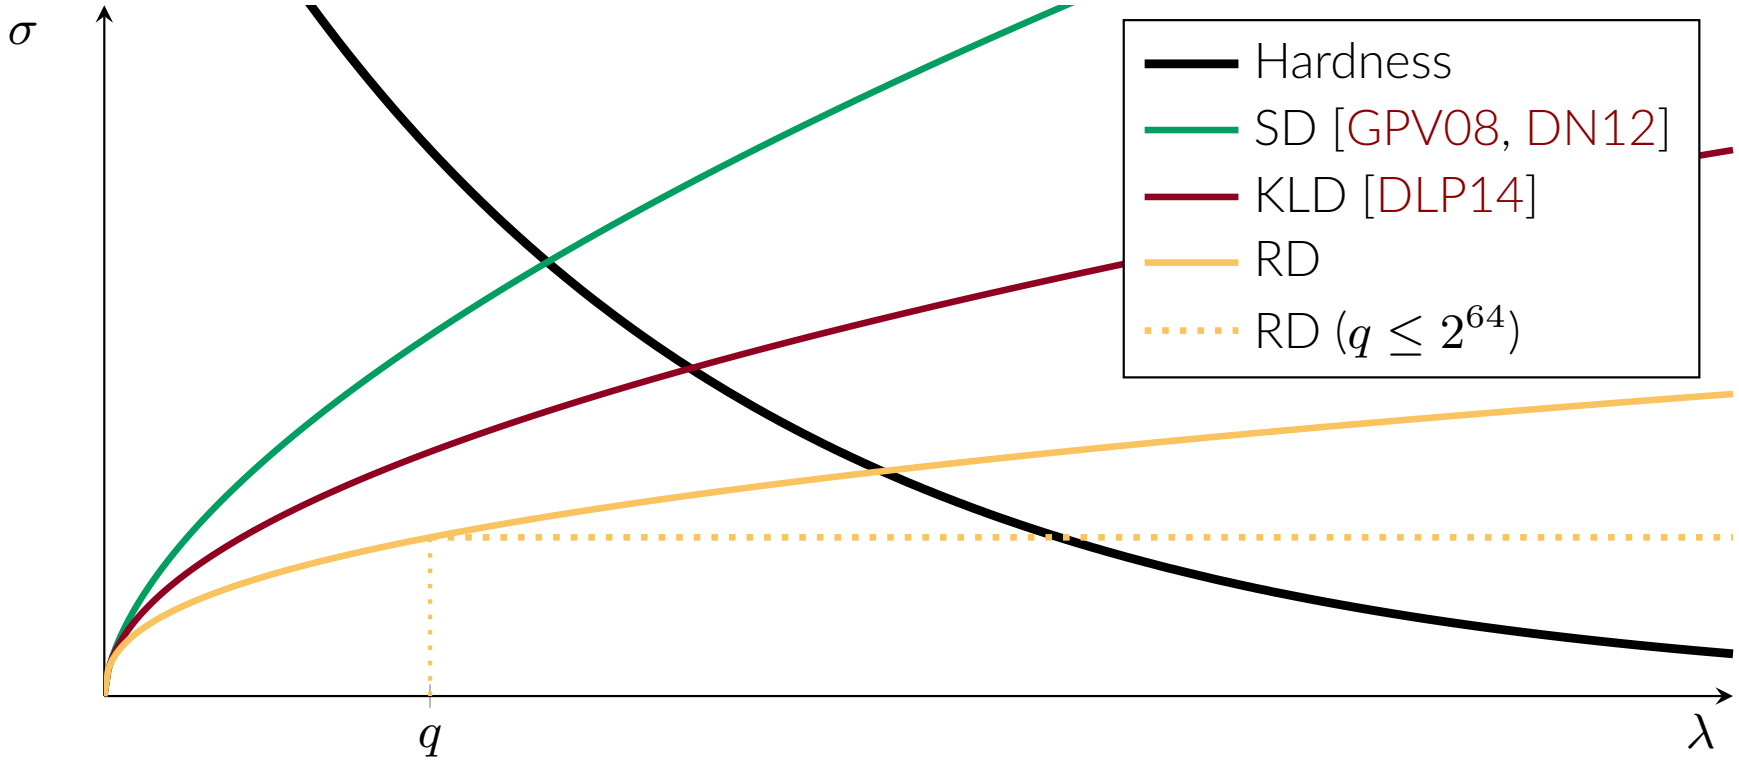

The adequate value for σ is at the intersection of the hardness curve (constraint 1) and the SD/KLD/RD curve (constraint 2).

- A Rényi divergence-based analysis proves to be much more efficient than an SD/KLD-based one.
- Interesting fact: in practice, σ is not conditioned by λ but by q.

In practice, we gain about 30 bits of security (compared to the SD).

What about the precision?

- **Previous works [LP15, Pre15].** a security of 256 bits requires about 150 bits of precision (completely impractical).
- **This work:** a security of 256 bits requires about 61 bits of precision.

Conclusion

Theory.

- We provide tools to optimize the use of the Rényi divergence for security proofs.
- Thanks to the reverse Pinsker's inequality, the fact that the Rényi divergence is not a distance is no longer a problem.
- These results are generic (not limited to lattice-based cryptography).

Practice.

- We get rid of high-precision arithmetic for a new sampler [MW17], rejection sampling, table-based sampling (with a new algorithm) and trapdoor sampling.
- We manage to squeeze the standard deviation of trapdoor samplers. This automatically increases security.

Open questions.

1. Implementation?
2. How can we compose the CoDF with other table-based approaches?
3. Other uses (in particular for reducing standard deviations)?
4. Other trapdoor samplers (probably easy but tedious)?
5. Can we get improvements outside of lattice-based cryptography?
6. Can we achieve a similar efficiency for decision problems?

Open questions.

1. Implementation?
2. How can we compose the CoDF with other table-based approaches?
3. Other uses (in particular for reducing standard deviations)?
4. Other trapdoor samplers (probably easy but tedious)?
5. Can we get improvements outside of lattice-based cryptography?
6. Can we achieve a similar efficiency for decision problems?

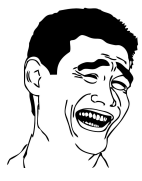

Thanks!


Shweta Agrawal, Dan Boneh, and Xavier Boyen.
Efficient lattice (H)IBE in the standard model.
In Gilbert [Gil10], pages 553–572.


Shweta Agrawal, Dan Boneh, and Xavier Boyen.
Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE.
In Rabin [Rab10], pages 98–115.


Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy.
Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits.
In Nguyen and Oswald [NO14], pages 533–556.


Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen.
On the hardness of learning with rounding over small modulus.
pages 209–224.


Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld.
Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance.
pages 3–24.


Xavier Boyen.
Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more.
In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 499–517. Springer, Heidelberg, May 2010.


Xavier Boyen.
Attribute-based functional encryption on lattices.
In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 122–142. Springer, Heidelberg, March 2013.


David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert.
Bonsai trees, or how to delegate a lattice basis.
In Gilbert [Gil10], pages 523–552.


Léo Ducas, Alain Durmus, Tancrède 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. Springer, Heidelberg, August 2013.


Nagarjun C. Dwarakanath and Steven D. Galbraith.
Sampling from discrete gaussians for lattice-based cryptography on a constrained device.
Appl. Algebra Eng. Commun. Comput., 25(3):159–180, 2014.


Léo Ducas, Vadim Lyubashevsky, and Thomas Prest.
Efficient identity-based encryption over NTRU lattices.
In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Heidelberg, December 2014.


Léo Ducas and Phong Q. Nguyen.
Faster Gaussian lattice sampling using lazy floating-point arithmetic.
In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 415–432. Springer, Heidelberg, December 2012.


Léo Ducas and Thomas Prest.
Fast fourier orthogonalization.
In Sergei A. Abramov, Eugene V. Zima, and Xiao-Shan Gao, editors, *Proceedings of the ACM on International Symposium on Symbolic and Algebraic Computation, ISSAC 2016, Waterloo, ON, Canada, July 19-22, 2016*, pages 191–198. ACM, 2016.


Henri Gilbert, editor.
EUROCRYPT 2010, volume 6110 of *LNCS*. Springer, Heidelberg, May 2010.


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. ACM Press, May 2008.


Philip N. Klein.
Finding the closest lattice vector when it's unusually close.
In *SODA*, 2000.


Vadim Lyubashevsky and Thomas Prest.
Quadratic time, linear space algorithms for Gram-Schmidt orthogonalization and Gaussian sampling in structured lattices.
In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015*, *Part I*, volume 9056 of *LNCS*, pages 789–815. Springer, Heidelberg, April 2015.


San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld.
Hardness of k-LWE and applications in traitor tracing.
In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014*, *Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Heidelberg, August 2014.


Adeline Langlois, Damien Stehlé, and Ron Steinfeld.
GGHLite: More efficient multilinear maps from ideal lattices.
In Nguyen and Oswald [NO14], pages 239–256.


Daniele Micciancio and Chris Peikert.
Trapdoors for lattices: Simpler, tighter, faster, smaller.
In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.


Daniele Micciancio and Michael Walter.
Gaussian sampling over the integers: Efficient, generic, constant-time.
CRYPTO, 2017.
http://eprint.iacr.org/2017/259.


Phong Q. Nguyen and Elisabeth Oswald, editors.
EUROCRYPT 2014, volume 8441 of *LNCS*. Springer, Heidelberg, May 2014.


Thomas Pöppelmann, Léo Ducas, and Tim Güneysu.
Enhanced lattice-based signatures on reconfigurable hardware.
In Lejla Batina and Matthew Robshaw, editors, *CHES 2014*, volume 8731 of *LNCS*, pages 353–370. Springer, Heidelberg, September 2014.


Chris Peikert.
An efficient and parallel Gaussian sampler for lattices.
In Rabin [Rab10], pages 80–97.


Thomas Prest.
Gaussian Sampling in Lattice-Based Cryptography.
Theses, École Normale Supérieure, December 2015.


Tal Rabin, editor.
CRYPTO 2010, volume 6223 of *LNCS*. Springer, Heidelberg, August 2010.