Discrete gaussian sampling for BKZ-reduced basis

Amaury Pouly and Yixin Shen

Centre National de la Recherche Scientifique (CNRS), Paris, France
INRIA, Rennes, France

9 April 2025

Discrete Gaussian Sampling

$$\rho_s(\boldsymbol{x}) = \exp\left(-\pi \frac{\|\boldsymbol{x}\|^2}{s^2}\right), \qquad D_{\mathcal{L},s}(\boldsymbol{x}) = \frac{\rho_s(\boldsymbol{x})}{\rho_s(\mathcal{L})}, \qquad \boldsymbol{x} \in \mathbb{R}^n, s > 0.$$

Definition (Discrete Gaussian Distribution)

On lattice $\mathcal{L}$ with parameter s: probability of $\boldsymbol{x} \in \mathcal{L}$ is $D_{\mathcal{L},s}(\boldsymbol{x})$.

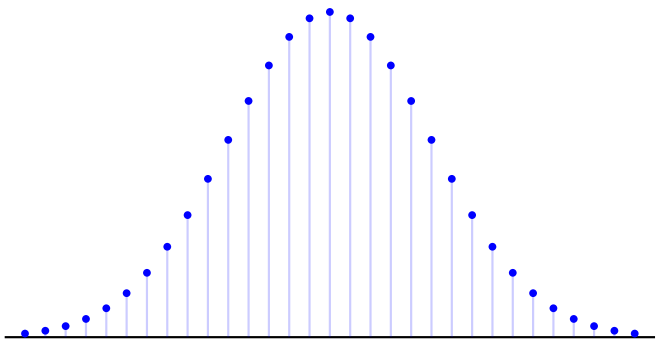

$\mathcal{L} = \mathbb{Z}, s = 7$

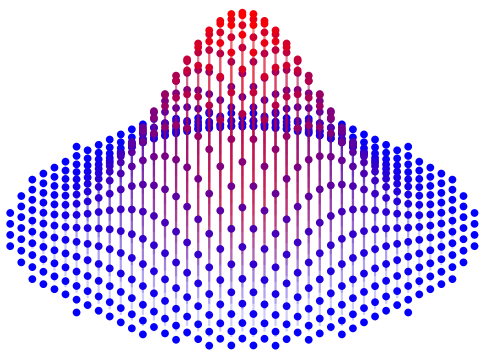

$\mathcal{L} = \mathbb{Z}^2, s = 7$

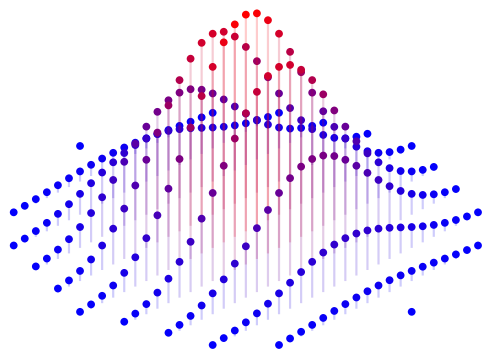

$\mathcal{L} = \mathbb{Z} \times 4\mathbb{Z}, s = 7$

Discrete Gaussian Sampling (DGS)

- **input:** $\mathcal{L}$ and s
- **output:** random $\boldsymbol{x} \in \mathcal{L}$ according to $D_{\mathcal{L},s}$.

Complexity of Discrete Gaussian Sampling

Sampling from the discrete Gaussian over $\mathcal{L}$ with parameter s:

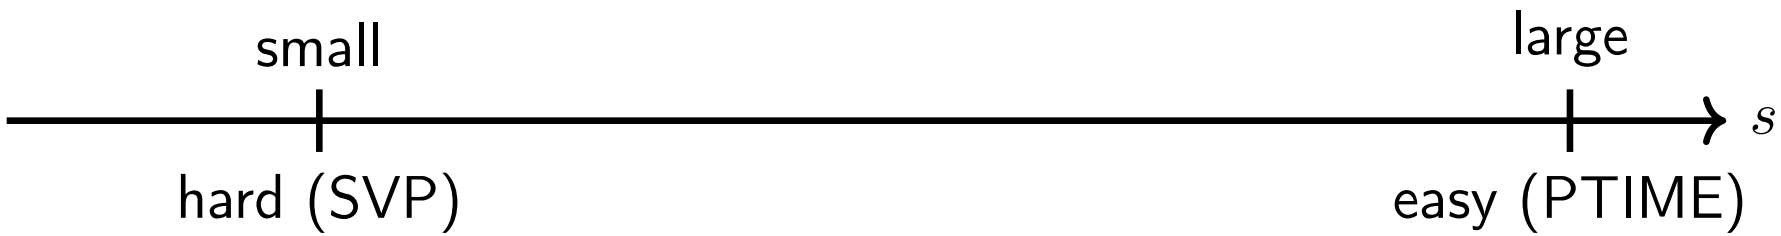

Complexity of Discrete Gaussian Sampling

Sampling from the discrete Gaussian over $\mathcal{L}$ with parameter s:

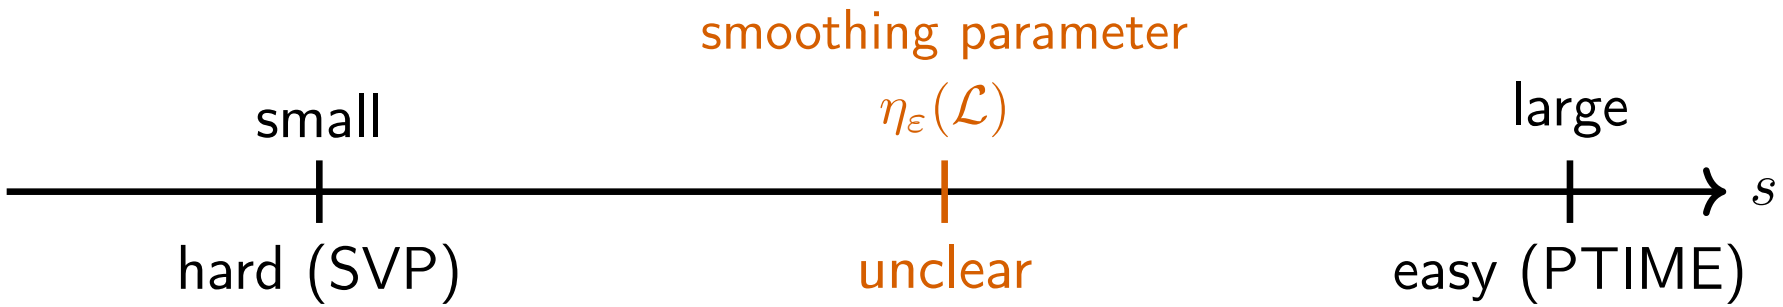

Smoothing parameter

For $\varepsilon > 0$, $\eta_\varepsilon(\mathcal{L}) = \inf\left\{s > 0 : \rho_{1/s}(\widehat{\mathcal{L}}) \leqslant 1 + \varepsilon\right\}$.

Complexity of Discrete Gaussian Sampling

Sampling from the discrete Gaussian over $\mathcal{L}$ with parameter s:

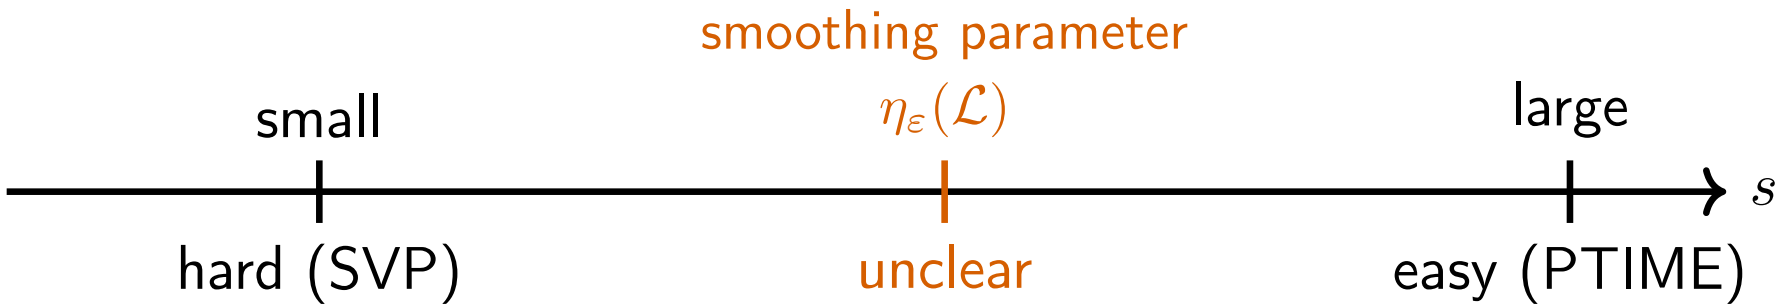

Smoothing parameter

For $\varepsilon > 0$, $\eta_\varepsilon(\mathcal{L}) = \inf\left\{s > 0 : \rho_{1/s}(\widehat{\mathcal{L}}) \leqslant 1 + \varepsilon\right\}$.

- Open problem: $2^{O(n)}$ time, $2^{o(n)}$ space algorithm for $s = \eta_\varepsilon(\mathcal{L})$
- $s \geqslant \eta_\varepsilon(\mathcal{L})$ is useful for hard problems: LWE, SVP, BDD, ...

Discrete Gaussian Samplers

Some samplers depends on the input basis $\mathbf{B}$, or its Gram-Schmidt orthogonalization $\tilde{\mathbf{B}}$.

Reference	s	Complexity
[GPV08]	$\|\tilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log n})$	polynomial
[ACKS21]	$\eta_{1/3}(\mathcal{L})$	$2^{n/2}$
[WL19]*	$\frac{1}{\sqrt{\pi}}\|\tilde{\mathbf{B}}\|$	1.0039^n
	arbitrary	depends on $\mathbf{B}$ and s

Limited choices for attacks:

- [GPV08]'s width is too large
- [ACKS21] is very expensive
- [WL19] so far unused?

*[BLP+13] seems to contain a similar result but it wasn't explicitly claimed.

Markov chain Monte Carlo (MCMC)-based sampling

Theorem ([WL19])

There is an algorithm that given a basis of $L \subset \mathbb{R}^n$ and any $s, \varepsilon > 0$, returns a sample according to some distribution ε-close to $\mathcal{D}_{L,s}$. It runs in time $\ln\left(\frac{1}{\varepsilon}\right) \cdot \frac{1}{\Delta} \cdot \mathrm{poly}(n)$ where $\frac{1}{\Delta} = \frac{1}{\rho_s(L)} \prod_{i=1}^{n} \rho_{s/\|\widetilde{\mathbf{b}}_i\|}(\mathbb{Z})$ and $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_d$ are the Gram-Schmidt vectors of the basis.

- no restriction on s
- complexity heavily depends on the shape of the basis
- complexity depends on $\rho_s(L)$ which is hard to estimate

Markov chain Monte Carlo (MCMC)-based sampling

Theorem ([WL19])

There is an algorithm that given a basis of $L \subset \mathbb{R}^n$ and any $s, \varepsilon > 0$, returns a sample according to some distribution ε-close to $\mathcal{D}_{L,s}$. It runs in time $\ln\left(\frac{1}{\varepsilon}\right) \cdot \frac{1}{\Delta} \cdot \mathrm{poly}(n)$ where $\frac{1}{\Delta} = \frac{1}{\rho_s(L)} \prod_{i=1}^{n} \rho_{s/\|\widetilde{\mathbf{b}}_i\|}(\mathbb{Z})$ and $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_d$ are the Gram-Schmidt vectors of the basis.

- no restriction on s
- complexity heavily depends on the shape of the basis
- complexity depends on $\rho_s(L)$ which is hard to estimate but

$$\frac{1}{\rho_s(L)} \prod_{i=1}^{n} \rho_{s/\|\widetilde{\mathbf{b}}_i\|}(\mathbb{Z}) = \frac{1}{\rho_{1/s}(\widehat{L})} \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$$

  is tight when $s \geqslant \eta_\varepsilon(L)$ and $\varepsilon = O(1)$.

MCMC complexity

Tight upper bound on complexity ($s \geqslant \eta_{O(1)}(L)$): $C \leqslant \prod_{i=1}^{n} \rho_{\|\tilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$

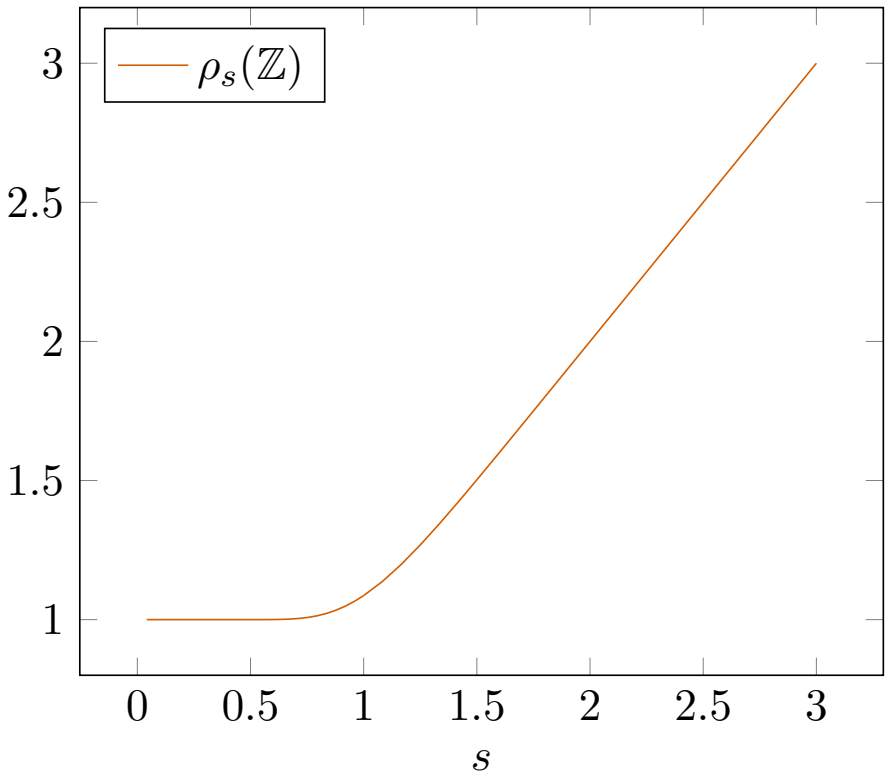

For a general basis:

$$C \leqslant (\rho_{\|\tilde{\mathbf{B}}\|/s}(\mathbb{Z}))^n \lessapprox \left(\frac{\|\tilde{\mathbf{B}}\|}{s}\right)^n$$

quite pessimistic (for $\|\tilde{\mathbf{B}}\| \geqslant s$)

MCMC complexity

Tight upper bound on complexity ($s \geqslant \eta_{O(1)}(L)$): $C \leqslant \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$

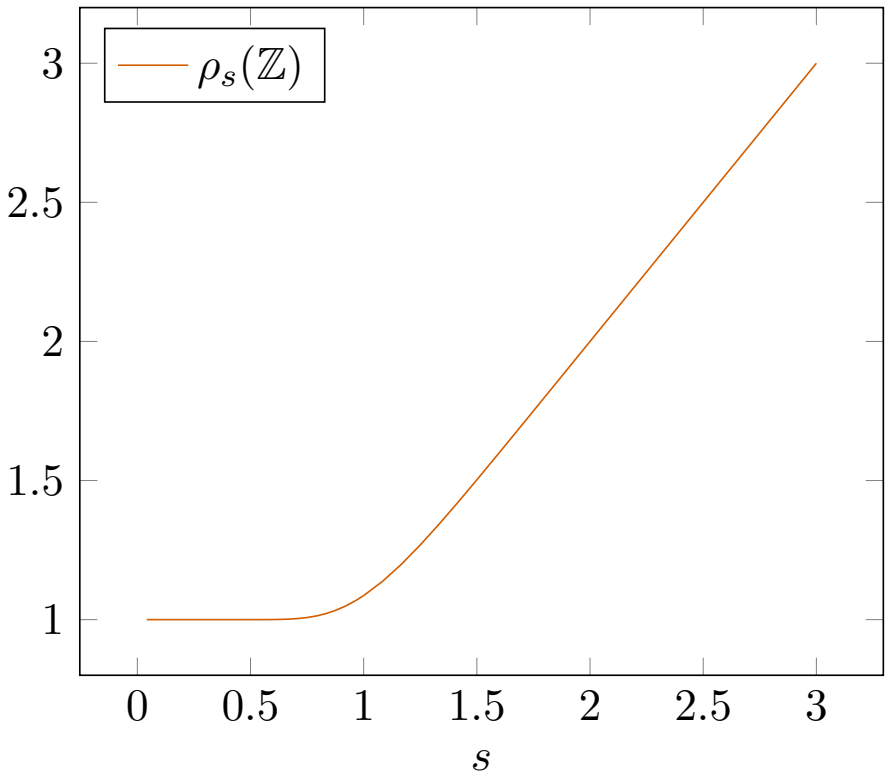

For a general basis:

$$C \leqslant (\rho_{\|\tilde{\mathbf{B}}\|/s}(\mathbb{Z}))^n \lessapprox \left(\frac{\|\tilde{\mathbf{B}}\|}{s}\right)^n$$

quite pessimistic (for $\|\tilde{\mathbf{B}}\| \geqslant s$)

Complexity is better if many $\|\widetilde{\mathbf{b}}_i\|$ are small.

$\rightsquigarrow$ BKZ-reduced basis have this feature

BKZ-reduced basis

Blockwise Korkine-Zolotare (BKZ): lattice basis reduction algorithm

- Parameter: blocksize $\beta \in [1, n]$
- Complexity: exponential in β
- $\|\widetilde{\mathbf{b}}_i\|$ decreases exponentially quickly

BKZ-reduced basis

Blockwise Korkine-Zolotare (BKZ): lattice basis reduction algorithm

- Parameter: blocksize $\beta \in [1, n]$
- Complexity: exponential in β
- $\|\widetilde{\mathbf{b}}_i\|$ decreases exponentially quickly

Geometric Series Assumption (GSA)

For a BKZ-β reduced basis,

$$\|\widetilde{\mathbf{b}}_i\| = \|\mathbf{b}_1\| H_\beta^{-2(i-1)}, \qquad \|\mathbf{b}_1\| = H_\beta^{n-1} \operatorname{vol}(L)^{1/n}$$

where $H_\beta := \left(\frac{\beta}{2\pi e}(\pi\beta)^{1/\beta}\right)^{1/2(\beta-1)}$.

- Reasonably accurate for $50 \leqslant \beta \ll n$

MCMC complexity for BKZ-reduced basis

- MCMC complexity:
$$C \leqslant \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$$
- GSA: $\|\widetilde{\mathbf{b}}_i\| = \|\mathbf{b}_1\| H_\beta^{-2(i-1)}$
- $\rho_s(\mathbb{Z}) \approx s$ for $s \geqslant 1$

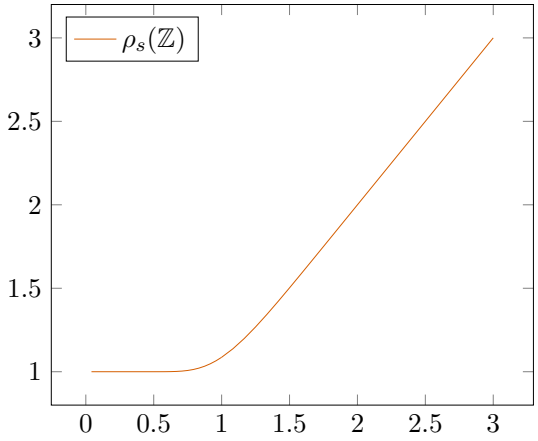

MCMC complexity for BKZ-reduced basis

- MCMC complexity:
$$C \leqslant \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$$
- GSA: $\|\widetilde{\mathbf{b}}_i\| = \|\mathbf{b}_1\| H_\beta^{-2(i-1)}$
- $\rho_s(\mathbb{Z}) \approx s$ for $s \geqslant 1$

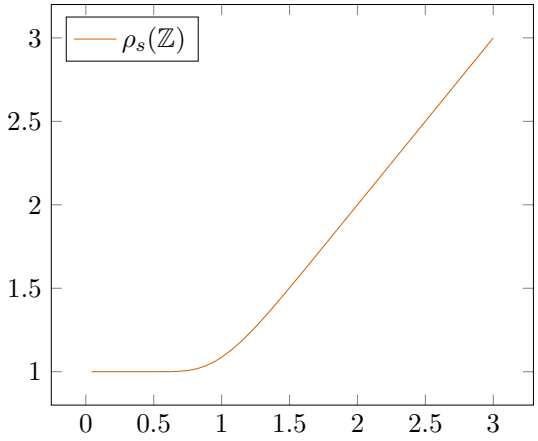

Very informally

For a BKZ-β reduced basis:

$$C \lessapprox \left(\frac{\sqrt{\|\tilde{\mathbf{B}}\|}}{s}\right)^{k_\beta}$$

for some $k_\beta \in [1, n]$ which depends on β

$\rightsquigarrow$ Much better than the generic bound $\left(\frac{\|\tilde{\mathbf{B}}\|}{s}\right)^{n}$.

What is in the paper?

Theorem

For a BKZ-β basis under the GSA,

$$\prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant \textit{(horrible formula)}$$

Formula is explicit and efficiently ($\approx$ constant time) computable

What is in the paper?

Theorem

For a BKZ-β basis under the GSA,

$$\prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant \textit{(horrible formula)}$$

Formula is explicit and efficiently ($\approx$ constant time) computable

- Comparison between BKZ and GSA for MCMC

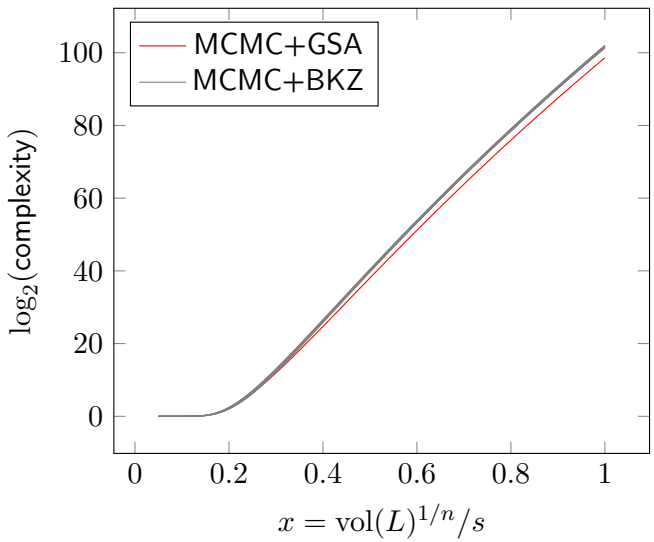

What is in the paper?

Theorem

For a BKZ-β basis under the GSA,

$$\prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant \textit{(horrible formula)}$$

Formula is explicit and efficiently ($\approx$ constant time) computable

- Comparison between BKZ and GSA for MCMC

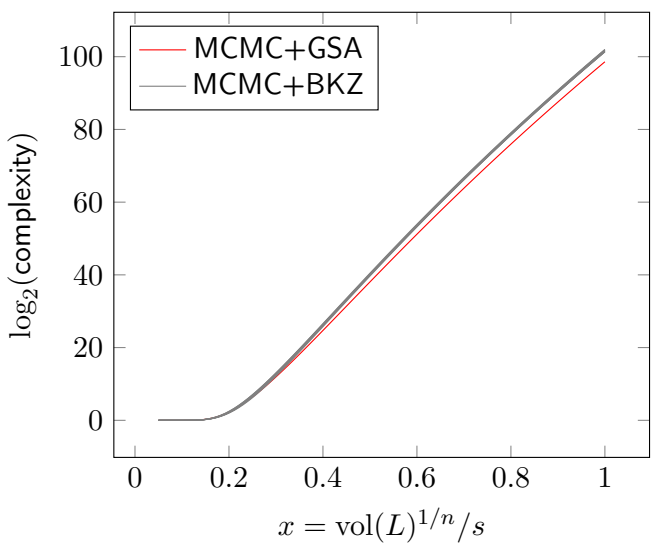

- Comparison between GSA and (horrible formula)

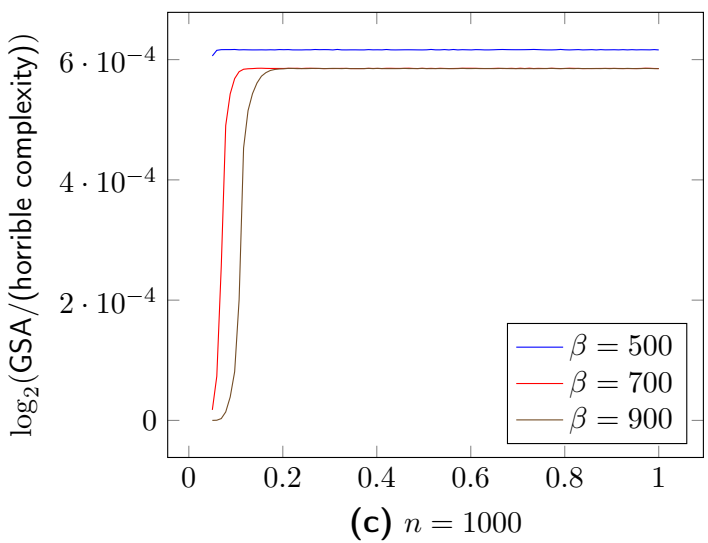

(c) $n = 1000$

Comparing BKZ and GSA for MCMC

GSA is relatively accurate but how does this affect the MCMC complexity?

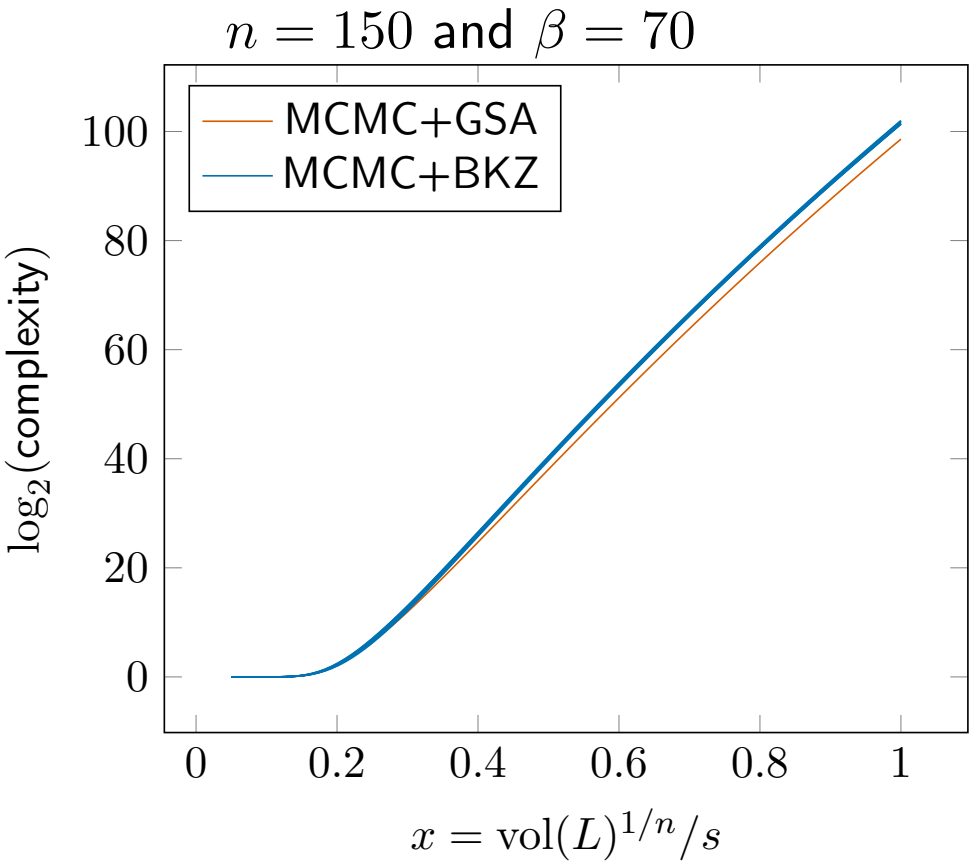

- Error increases with $1/s$
- Relatively small error, especially for $x \leqslant 1/4$
- Hard to extrapolate to $n = 1000$ without BKZ simulators

Tentative conclusion for MCMC complexity

GSA introduces only a minor relative error, especially for $s \gg \mathrm{vol}(L)^{1/n}$.

Approximation formula for MCMC+GSA

Theorem (Simplified)

For a BKZ-β basis under the GSA, $s > 0$ and any $p \in \{1, 3, \ldots\}$,

$$\log \prod_{i=1}^{n} \rho_{\|\tilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant \log A + \sum_{\ell=1}^{p} X_\ell$$

where

$$A = \left(\frac{\sqrt{\|\tilde{\mathbf{B}}\|}}{s}\right)^{k_\beta}, \qquad k_\beta \approx \frac{\ln \|\tilde{\mathbf{B}}\| - \ln s}{\ln H_\beta} \qquad X_\ell = \textit{corrective term}$$

- A is trivial to compute, good formula for intuitions
- X_ℓ ugly but cheap to compute and quickly get smaller as $\ell \to \infty$
- $p = 3$ gives almost perfect approximation (next slide)

Comparing GSA and the approximation formula

Theorem:

$$\log \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant k_\beta \log\left(\frac{\sqrt{\|\tilde{\mathbf{B}}\|}}{s}\right) + \sum_{\ell=1}^{p} X_\ell \qquad (\mathsf{X})$$

$\log_2(\text{MCMC+GSA complexity})$:

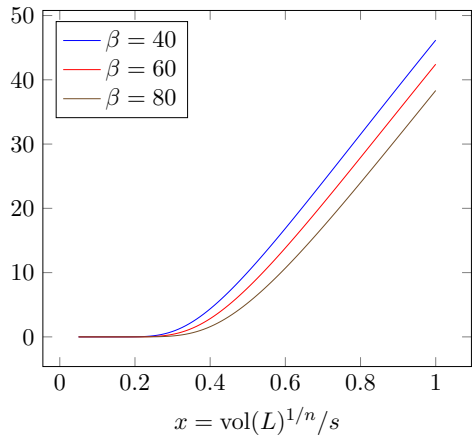

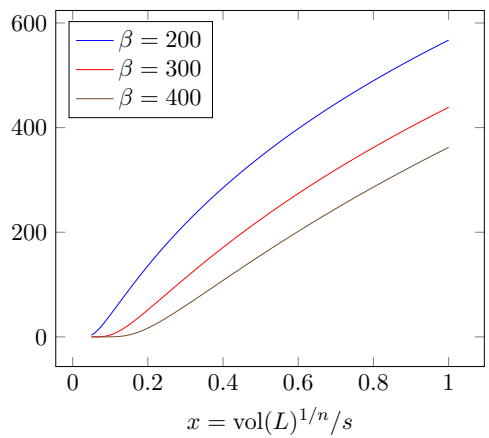

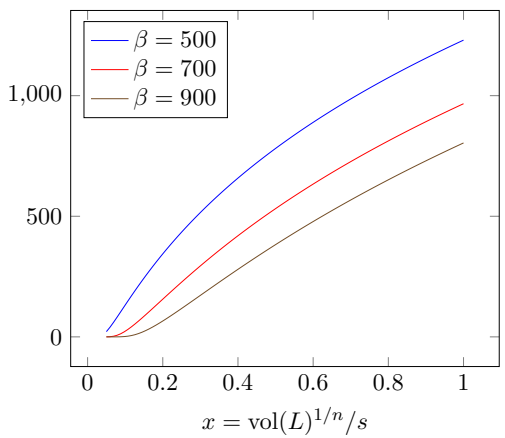

$\log_2((\mathsf{X})) - \log_2(\text{MCMC+GSA complexity})$:

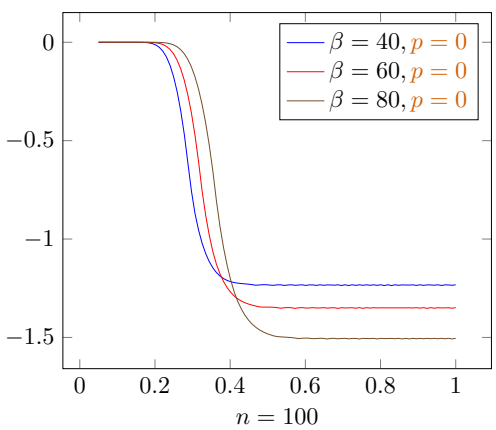

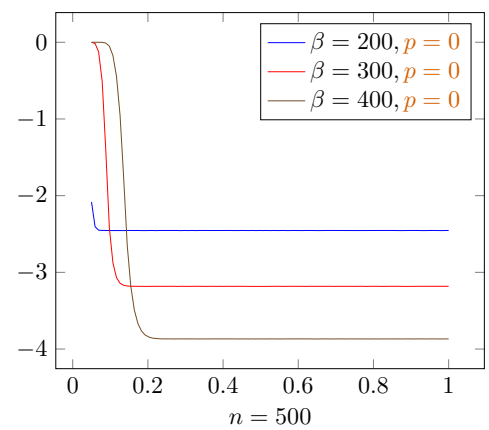

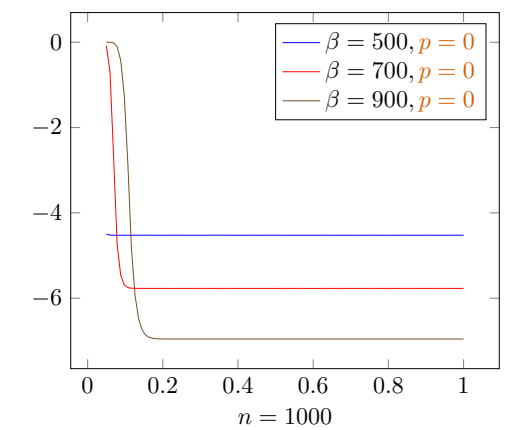

Comparing GSA and the approximation formula

Theorem:

$$\log \prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant k_\beta \log\left(\frac{\sqrt{\|\tilde{\mathbf{B}}\|}}{s}\right) + \sum_{\ell=1}^{p} X_\ell \qquad (\mathsf{X})$$

$\log_2(\text{MCMC+GSA complexity})$:

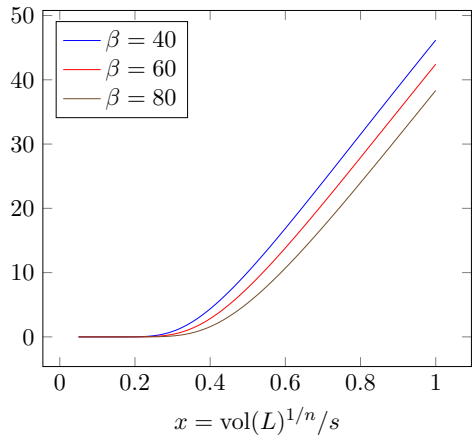

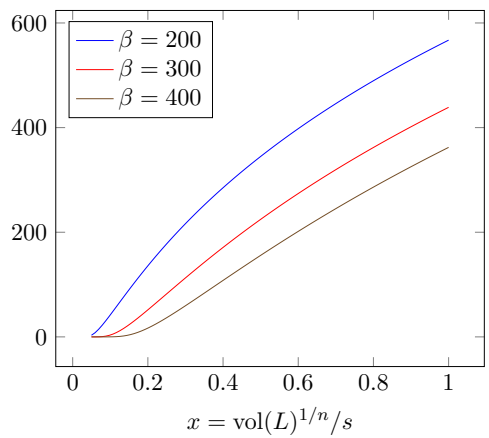

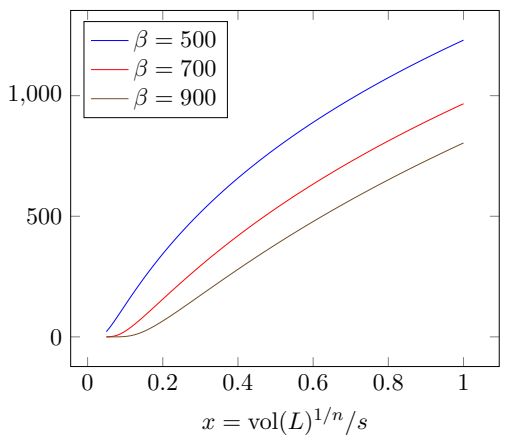

$\log_2((\mathsf{X})) - \log_2(\text{MCMC+GSA complexity})$:

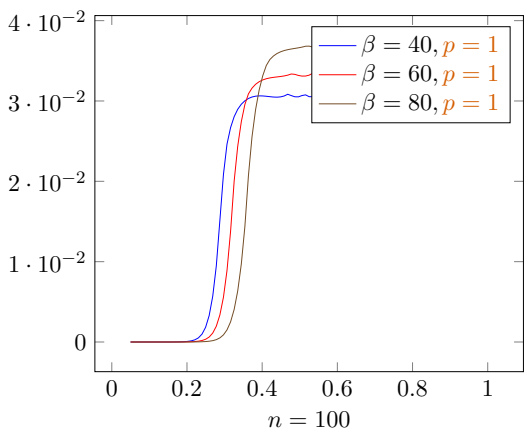

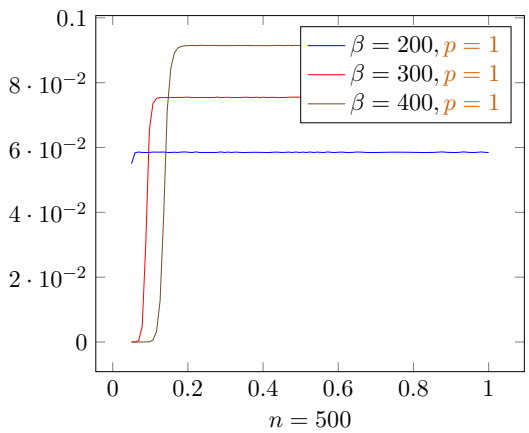

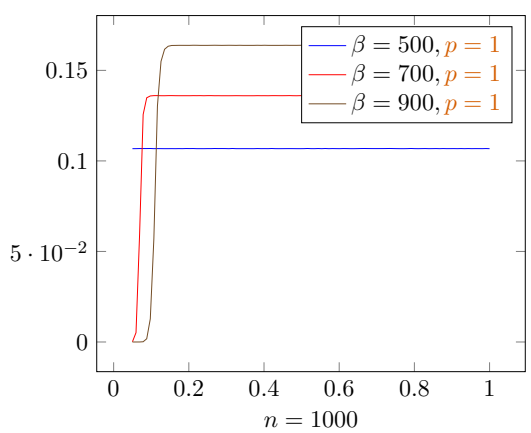

Comparing GSA and the approximation formula

Theorem:

$$\log \prod_{i=1}^{n} \rho_{\|\tilde{\mathbf{b}}_i\|/s}(\mathbb{Z}) \leqslant k_\beta \log\left(\frac{\sqrt{\|\tilde{\mathbf{B}}\|}}{s}\right) + \sum_{\ell=1}^{p} X_\ell \qquad (\mathsf{X})$$

$\log_2(\text{MCMC+GSA complexity})$:

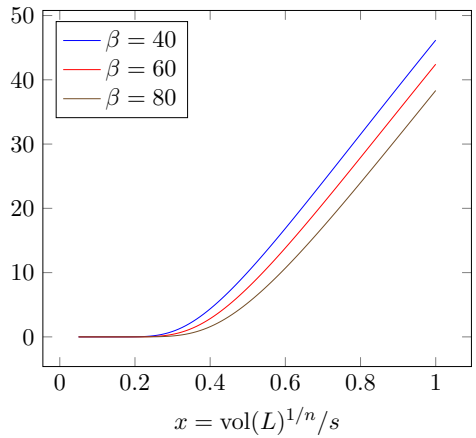

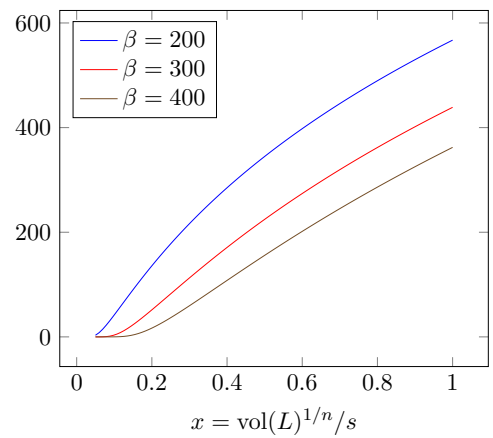

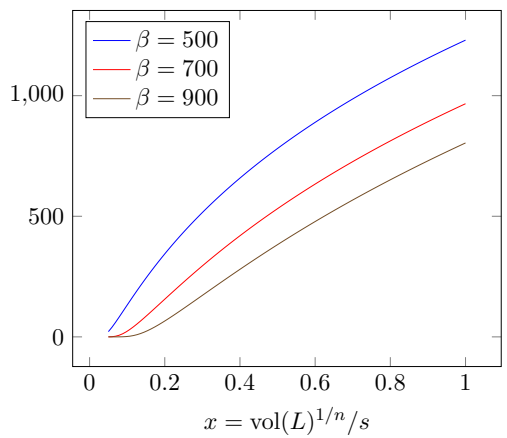

$\log_2((\mathsf{X})) - \log_2(\text{MCMC+GSA complexity})$:

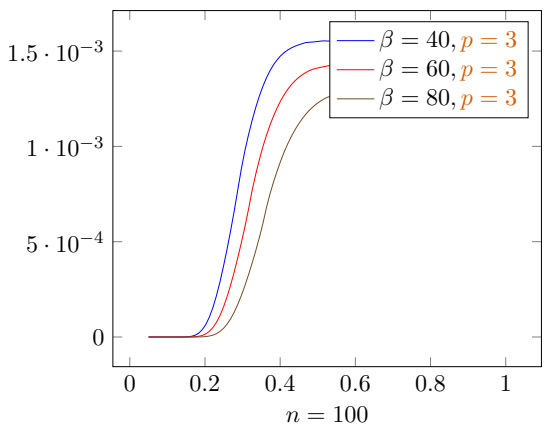

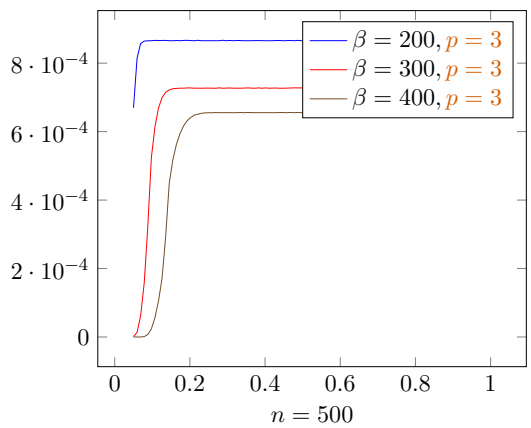

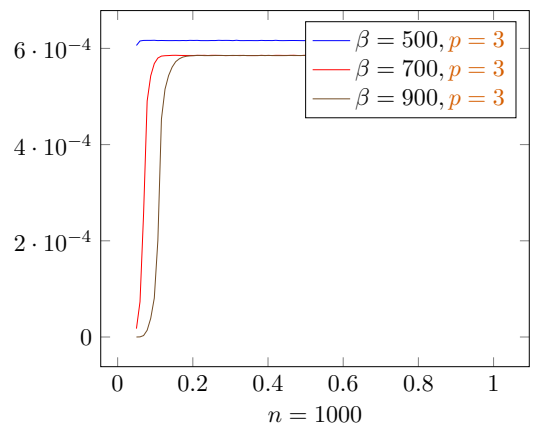

Application: dual attack on LWE

[PS24] describes a dual attack on LWE using [WL19] as Gaussian sampler.

Needs to optimize the choice of parameters (e.g. β, s):

- run an estimator for many parameters and choose the best

Application: dual attack on LWE

[PS24] describes a dual attack on LWE using [WL19] as Gaussian sampler.

Needs to optimize the choice of parameters (e.g. β, s):

- run an estimator for many parameters and choose the best
- estimating the complexity of MCMC using $\prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$ is slow
- instead uses a crude but fast approximation
- $\rightsquigarrow$ limits how low s can be chosen

Application: dual attack on LWE

[PS24] describes a dual attack on LWE using [WL19] as Gaussian sampler.

Needs to optimize the choice of parameters (e.g. β, s):

- run an estimator for many parameters and choose the best
- estimating the complexity of MCMC using $\prod_{i=1}^{n} \rho_{\|\widetilde{\mathbf{b}}_i\|/s}(\mathbb{Z})$ is slow
- instead uses a crude but fast approximation
- $\rightsquigarrow$ limits how low s can be chosen

Contribution: revisit [PS24] using our approximation formula

- better approximation $\rightsquigarrow$ less constraints on s
- unlocks better choices of parameters (next slide)

Dual attack on LWE: estimates

Scheme	No modulus switching		With modulus switching	
	This paper	[PS24]	This paper	[PS24]
Kyber512	182	185	141	141
Kyber768	267	273	201	202
Kyber1024	366	376	273	279

Take away

Nontrivial improvement in the complexity

In the paper: further improvements by doing a smarter parameter search

Dual attack on LWE: estimates

Scheme	No modulus switching		With modulus switching	
	This paper	[PS24]	This paper	[PS24]
Kyber512	182	185	141	141
Kyber768	267	273	201	202
Kyber1024	366	376	273	279

Take away

Nontrivial improvement in the complexity

In the paper: further improvements by doing a smarter parameter search

GSA/BKZ approximation error

Recall: error increases with $x = \mathrm{vol}(L)^{1/n}/s$.
Dual attack only uses $x \leqslant 0.1 \rightsquigarrow$ approx error expected to be small.

Conclusion and future work

Complexity analysis of the MCMC sampler [WL19] for BKZ-reduced basis:

- Comparison between actual BKZ and GSA
- Upper bound/approximation formula for MCMC+GSA
- Comparison between MCMC+GSA and formula

Application to dual attack on LWE

Also in the paper: analysis of $\eta_\varepsilon(L)$ for random q-ary lattices L


Divesh Aggarwal, Yanlin Chen, Rajendra Kumar, and Yixin Shen.
Improved (provable) algorithms for the shortest vector problem via bounded distance decoding.
In Markus Bläser and Benjamin Monmege, editors, *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPIcs*, pages 4:1–4:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.


Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé.
Classical hardness of learning with errors.
In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, page 575–584, New York, NY, USA, 2013. Association for Computing Machinery.


Shi Bai, Damien Stehlé, and Weiqiang Wen.
Measuring, simulating and exploiting the head concavity phenomenon in bkz.

In *Advances in Cryptology – ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part I*, page 369–404, Berlin, Heidelberg, 2018. Springer-Verlag.


Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan.
Trapdoors for hard lattices and new cryptographic constructions.
In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, page 197–206, New York, NY, USA, 2008. Association for Computing Machinery.


Amaury Pouly and Yixin Shen.
Provable dual attacks on learning with errors.
In *Advances in Cryptology – EUROCRYPT 2024: 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26–30, 2024, Proceedings, Part VII*, page 256–285, Berlin, Heidelberg, 2024. Springer-Verlag.


Zheng Wang and Cong Ling.

Lattice gaussian sampling by markov chain monte carlo: Bounded distance decoding and trapdoor sampling.
IEEE Transactions on Information Theory, 65(6):3630–3645, 2019.