Attacks on LWE

Martin R. Albrecht
Oxford Lattice School

Outline

BKZ Refresher

LWE

Dual Lattice Attack

Primal Lattice Attack (uSVP Version)

BKZ Refresher

BKZ

- Input basis is LLL reduced, the first block is $\mathbf{b}_1, \ldots, \mathbf{b}_\beta$.
- Call the SVP oracle to obtain a short vector, $\mathbf{b}_1'$, in the space spanned by these vectors.
- Now have $\beta + 1$ vectors spanning a β dimensional space, call LLL to obtain a set of β linearly independent vectors.
- The second block is made of vectors which are the projection of $\mathbf{b}_2, \ldots, \mathbf{b}_{\beta+1}$ onto the space which is orthognal to $\mathbf{b}_1$.
- Again, call SVP oracle to obtain a short vector in this space, $\mathbf{b}_2'$, which can be viewed as the projection of some $\mathbf{b}_2''$ in the lattice.
- Call LLL on $\mathbf{b}_2, \mathbf{b}_3, \ldots, \mathbf{b}_{\beta+1}, \mathbf{b}_2''$ to update the list of basis vectors.

… start again when reaching end, repeat until nothing changes

BKZ 2.0

early abort BKZ eventually terminates when there is nothing left to do. However, most work is done in the first few tours

recursive preprocessing use BKZ with smaller block size to preprocess blocks before calling the SVP oracle

(extreme) pruning choose pruning parameters which lead to low probability of success, rerandomise and repeat to boost probability

Gaussian heuristic use the Gaussian heuristic to set radius for enumeration search

Yuanmi Chen. Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe. PhD thesis. Paris 7, 2013, url: `https://www.nofile.io/f/PvRtI1VlkJ`, implementation: `https://github.com/fplll/fplll`

root-Hermite Factors

The shortest non-zero vector $\mathbf{b}_1$ in the output basis satisfies:

$$\|\mathbf{b}_1\| = \delta_0^d \cdot \mathrm{Vol}(\Lambda)^{1/d}.$$

- Hermite factor: δ_0^d
- root-Hermite factor: δ_0
- log root-Hermite factor: $\log_2 \delta_0$

Gaussian Heuristic

Let $\Lambda \subset \mathbb{Z}^d$ be a lattice and let $S \in \mathbb{R}^d$ be a measurable subset of the real space. Then

$$|S \cap \Lambda| \approx \mathrm{Vol}(S)/\,\mathrm{Vol}(\Lambda).$$

As a corollary, considering spheres, we get:

$$\lambda_1(\Lambda) \approx \sqrt{\frac{d}{2\pi e}}\,\mathrm{Vol}(\Lambda)^{1/d}.$$

Geometric Series Assumption

The norms of the Gram-Schmidt vectors after lattice reduction satisfy[1]

$$\|\mathbf{b}_i^*\| = \alpha^{i-1} \cdot \|\mathbf{b}_1\| \text{ for some } 0 < \alpha < 1.$$

Combining this with the root-Hermite factor $\|\mathbf{b}_1\| = \delta_0^d \cdot \mathrm{Vol}(\Lambda)^{1/d}$ and $\mathrm{Vol}(\Lambda) = \prod_{i=1}^{d} \|\mathbf{b}_i^*\|$, we get

$$\alpha = \delta^{-2d/(d-1)}.$$

[1] Claus-Peter Schnorr. Lattice Reduction by Random Sampling and Birthday Methods. In: *STACS 2003, 20th Annual Symposium on Theoretical Aspects of Computer Science, Berlin, Germany, February 27 - March 1, 2003, Proceedings*. Ed. by Helmut Alt and Michel Habib. Vol. 2607. Lecture Notes in Computer Science. Springer, 2003, pp. 145–156. DOI: 10.1007/3-540-36494-3_14. URL: http://dx.doi.org/10.1007/3-540-36494-3_14.

BKZ Quality

Assuming the Gaussian Heuristic (GH) and the Geometric Series Assumption (GSA), a limiting value of the root-Hermite factor δ_0 achievable by BKZ is[2]:

$$\lim_{n\to\infty} \delta_0 = \left(v_\beta^{\frac{-1}{\beta}} \right)^{\frac{1}{\beta-1}} \approx \left(\frac{\beta}{2\pi e} (\pi\beta)^{\frac{1}{\beta}} \right)^{\frac{1}{2(\beta-1)}}$$

where v_β is the volume of the unit ball in dimension β. Experimental evidence suggests that we may apply this as an estimate for δ_0 also in practice.

[2] Yuanmi Chen. Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe. PhD thesis. Paris 7, 2013.

BKZ Quality

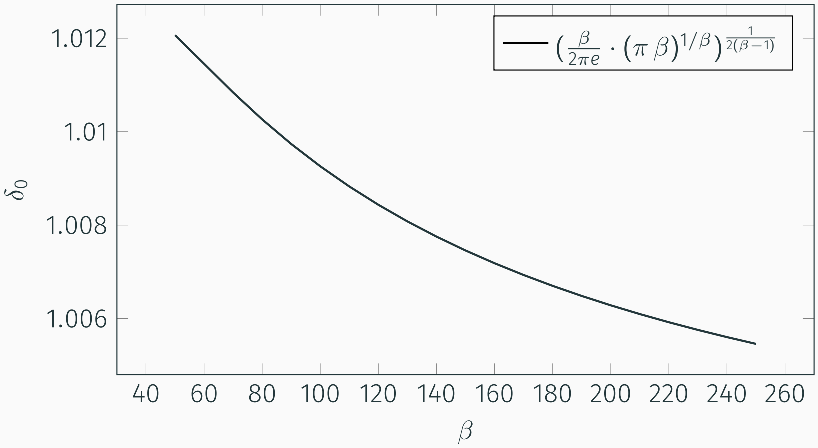
$\left(\frac{\beta}{2\pi e} \cdot (\pi\, \beta)^{1/\beta}\right)^{\frac{1}{2(\beta-1)}}$
1.012
1.01
1.008
1.006
δ_0
40
60
80
100
120
140
160
180
200
220
240
260
β

Running Time

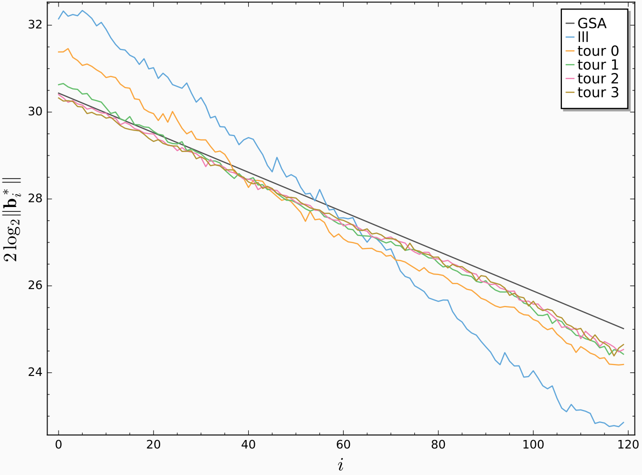

Most work is done in first 3-4 tours.

Running Time

Per tour, BKZ calls

$c_{\text{pre},\beta}$	prepare n SVP calls
$c_{\text{svp},\beta}$	n SVP oracle calls in block size $\leq \beta$
c_{lll}	n LLL calls to insert the vector into the basis

Total cost:

$$\approx 4\, n \cdot (c_{\text{pre},\beta} + c_{\text{svp},\beta} + c_{\text{lll}})$$

Running Time

We assume

- $c_{\text{pre},\beta} < c_{\text{svp},\beta}$[3] and
- $c_{\text{lll}} \ll c_{\text{svp},\beta}$

to obtain

$$\approx 4\, n\, c_{\text{svp},\beta}$$

Asymptotically, sieving is the most efficient heuristic SVP algorithm, with a cost[4] of

$$c_{\text{svp},\beta} = 2^{0.292\,\beta + o(1)}.$$

[3] For current code, this is a blatant lie.

[4] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In: *27th SODA*. ed. by Robert Krauthgamer. ACM-SIAM, Jan. 2016, pp. 10–24. DOI: 10.1137/1.9781611974331.ch2.

The log of the time complexity for running BKZ to achieve a root-Hermite factor δ_0 is:[5]

$$\Omega\left(\frac{-\log\left(\frac{-\log\log\delta_0}{\log\delta_0}\right)\log\log\delta_0}{\log\delta_0}\right)\text{for enumeration,}$$

$$\Omega\left(\frac{-\log\log\delta_0}{\log\delta_0}\right)\text{for sieving.}$$

[5]Martin R. Albrecht, Rachel Player, and Sam Scott. On The Concrete Hardness Of Learning With Errors. Cryptology ePrint Archive, Report 2015/046. http://eprint.iacr.org/2015/046. 2015.

Here's a Picture of a Kitten

LWE

Learning with Errors

Let n, q be positive integers, χ be a probability distribution on $\mathbb{Z}$ and $\mathbf{s}$ be a secret vector in $\mathbb{Z}_q^n$. We denote by $L_{n,q,\chi}$ the probability distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ obtained by choosing $\mathbf{a} \in \mathbb{Z}_q^n$ uniformly at random, choosing $e \in \mathbb{Z}$ according to χ and considering it in $\mathbb{Z}_q$, and returning $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.

Decision-LWE is the problem of deciding whether pairs $(\mathbf{a}, c) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ are sampled according to $L_{n,q,\chi}$ or the uniform distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$.

Search-LWE is the problem of recovering $\mathbf{s}$ from $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ sampled according to $L_{n,q,\chi}$.

Dual Lattice Attack

Short Integer Solutions

Consider the scaled (by q) dual lattice:

$$q\Lambda^* = \{\mathsf{x} \in \mathbb{Z}^m \mid \mathsf{x} \cdot \mathsf{A} \equiv 0 \bmod q\}.$$

A short vector of $q\Lambda^*$ is equivalent to solving SIS on A.

Short Integer Solutions (SIS)

Given $q \in \mathbb{Z}$, a matrix A, and $t < q$; find $\mathbf{y}$ with $0 < \|\mathbf{y}\| \leq t$ and

$$\mathsf{y} \cdot \mathsf{A} \equiv \mathbf{0} \pmod{q}.$$

Strategy

Given samples A, c:

1. Find a short y solving SIS on A.
2. Compute $\langle \mathsf{y}, \mathsf{c} \rangle$.

Either $\mathsf{c} = \mathsf{A}\mathsf{s} + \mathsf{e}$ or c uniformly random:

- If c is uniformly random, so is $\langle \mathsf{y}, \mathsf{c} \rangle$.
- If $\mathsf{c} = \mathsf{A} \cdot \mathsf{s} + \mathsf{e}$, then $\langle \mathsf{y}, \mathsf{c} \rangle = \langle \mathsf{y} \cdot \mathsf{A}, \mathsf{s} \rangle + \langle \mathsf{y}, \mathsf{e} \rangle \equiv \langle \mathsf{y}, \mathsf{e} \rangle \pmod{q}$. If y is sufficiently short, then $\langle \mathsf{y}, \mathsf{e} \rangle$ will also be short, since e is also small.

Required Quality

Given an LWE instance characterised by n, α, q and a vector $\mathbf{v}$ of length $\|\mathbf{v}\|$ in the scaled dual lattice

$$q\Lambda^* = \{\mathbf{x} \in \mathbb{Z}_q^m \mid \mathbf{x} \cdot \mathbf{A} \equiv 0 \bmod q\},$$

the advantage[6] of distinguishing $\langle \mathbf{v}, \mathbf{e} \rangle$ from random is close to

$$\exp\left(-\pi(\|\mathbf{v}\| \cdot \alpha)^2\right).$$

[6] Richard Lindner and Chris Peikert. Better Key Sizes (and Attacks) for LWE-Based Encryption. In: *CT-RSA 2011*. Ed. by Aggelos Kiayias. Vol. 6558. LNCS. Springer, Heidelberg, Feb. 2011, pp. 319–339.

A reduced lattice basis is made of short vectors, in particular the first vector has norm $\delta_0^m \cdot \mathsf{Vol}(q\Lambda^*)^{1/m}$

1. Construct bases of the dual for the instance.
2. Feed to a lattice reduction algorithm to obtain short vectors $\mathbf{v}_i$.
3. Check if $\mathbf{v}_i \cdot \mathbf{A}$ are small.

Constructing a Basis

- We seek a basis for the q-ary lattice

$$q\Lambda^* = \{\mathsf{x} \in \mathbb{Z}_q^m \mid \mathsf{x} \cdot \mathsf{A} \equiv 0 \bmod q\}$$

- Compute a row-echelon form Y of the basis for the left-kernel of A mod q using Gaussian elimination.
- With high probability it will have dimension $(m-n) \times m$
- Write $\mathsf{Y} = [\mathsf{I}_{(m-n)\times(m-n)} | \mathsf{Y}']$
- Extend to q-ary lattice by stacking on top of $[\mathsf{0}_{n\times(m-n)} \mid q \cdot \mathsf{I}_{n\times n}]$
- The basis is

$$\mathsf{L} = \begin{pmatrix} \mathsf{I}_{(m-n)\times(m-n)} & \mathsf{Y}' \\ 0 & q\,\mathsf{I}_{n\times n} \end{pmatrix}$$

Degrees of Freedom

- the dimension m, i.e. the number of samples we use, and
- the target advantage ε for distinguishing

CHOOSING m

Example: $q = 2^{17}, n = 1024, \delta_0 = 1.005$

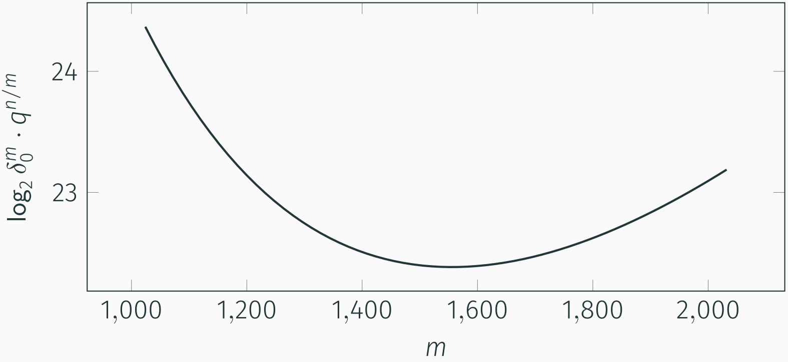

$$m = \sqrt{\frac{n \log q}{\log \delta_0}}$$

Choosing ε

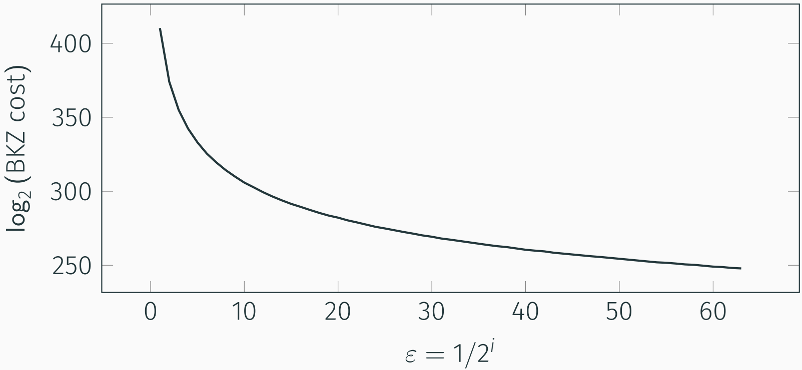

Choosing ε

Repeat experiment $\approx 1/\varepsilon^2$ times for majority vote to achieve constant advantage

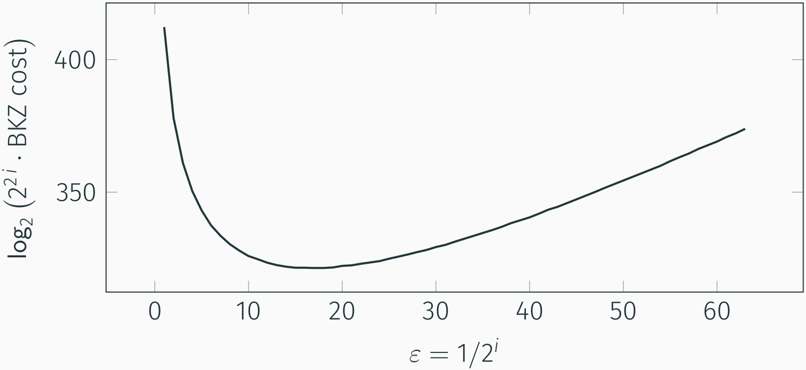

Amortising Costs

Producing $1/\varepsilon^2$ short vectors is cheaper than $1/\varepsilon^2$ calls to BKZ in block size β.

Two options:

- Use that sieving outputs $2^{0.2075 \cdot \beta}$ vectors.[7]
- Perform strong lattice reduction once, use light rerandomisation and cheaper lattice reduction for subsequent vectors.[8]

---

[7] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - a new hope. Cryptology ePrint Archive, Report 2015/1092. `http://eprint.iacr.org/2015/1092`. 2015.

[8] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in HElib and SEAL. Cryptology ePrint Archive, Report 2017/047. `http://eprint.iacr.org/2017/047`. 2017.

LWE Normal Form

Problem: most schemes give only n samples $\Rightarrow$ left kernel is trivial

But instances are in LWE normal form: $\mathbf{s}_i \leftarrow_{\$} \chi$

LWE Normal Form

Given samples $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ with $\mathbf{a} \leftarrow \mathcal{U}(\mathbb{Z}_q^n)$, $e \leftarrow \chi$ and $\mathbf{s} \in \mathbb{Z}_q^n$, we can construct samples

$$(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{e} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$$

with $\mathbf{a} \leftarrow \mathcal{U}(\mathbb{Z}_q^n)$, $e \leftarrow \chi$ and $\mathbf{e}$ such that all components

$$e_i \leftarrow \chi$$

in polynomial time.[9]

[9] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In: *CRYPTO 2009*. Ed. by Shai Halevi. Vol. 5677. LNCS. Springer, Heidelberg, Aug. 2009, pp. 595–618.

LWE Normal Form

- Construct basis for

$$\Lambda = \{(\mathbf{y}, \mathbf{x}) \in \mathbb{Z}^m \times \mathbb{Z}^n : \mathbf{y} \cdot \mathbf{A} \equiv \mathbf{x} \bmod q\}.$$

- Given a short vector in $(\mathbf{w}, \mathbf{v}) \in \Lambda$, we have

$$\mathbf{w} \cdot \mathbf{c} = \mathbf{w} \cdot (\mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = \langle \mathbf{v}, \mathbf{s} \rangle + \langle \mathbf{w}, \mathbf{e} \rangle .$$

- Analysis proceeds as before with $d \leq 2n$.

Small Secret

Assume $\|\mathbf{s}\| \ll \|\mathbf{e}\|$, e.g. $\mathbf{s}_i \leftarrow \{-1, 0, 1\}$.

- Aim is to balance $\| \langle \mathbf{v}, \mathbf{s} \rangle \| \approx \| \langle \mathbf{w}, \mathbf{e} \rangle \|$.
- Consider the scaled dual attack lattice

$$\Lambda(\mathbf{L}) = \{(\mathbf{x}, \mathbf{y}/c) \in \mathbb{Z}^m \times (1/c \cdot \mathbb{Z})^n : \mathbf{x} \cdot \mathbf{A} \equiv \mathbf{y} \bmod q\}$$

  for some constant c.
- Lattice reduction produces a vector $(\mathbf{v}', \mathbf{w}')$ with

$$\|(\mathbf{v}', \mathbf{w}')\| \approx \delta_0^{(m+n)} \cdot (q/c)^{n/(m+n)}.$$

- The final error we aim to distinguish from uniform is

$$e = \mathbf{v}' \cdot \mathbf{A} \cdot \mathbf{s} + \langle \mathbf{v}', \mathbf{e} \rangle = \langle c \cdot \mathbf{w}', \mathbf{s} \rangle + \langle \mathbf{v}', \mathbf{e} \rangle .$$

Honourable Mention: BKW

Assume $(\mathbf{a}_{21}, \mathbf{a}_{22}) = (0, 1)$, then:

$$\begin{pmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \mathbf{a}_{13} & \cdots & \mathbf{a}_{1n} & c_1 \\ \mathbf{a}_{21} & \mathbf{a}_{22} & \mathbf{a}_{23} & \cdots & \mathbf{a}_{2n} & c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \\ \mathbf{a}_{m1} & \mathbf{a}_{m2} & \mathbf{a}_{m3} & \cdots & \mathbf{a}_{mn} & c_m \end{pmatrix}$$

$$- \begin{bmatrix} 0 & 0 & \mathbf{t}_{13} & \cdots & \mathbf{t}_{1n} & c_{t,1} \\ 0 & 1 & \mathbf{t}_{23} & \cdots & \mathbf{t}_{2n} & c_{t,2} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \\ q-1 & q-1 & \mathbf{t}_{q^2 3} & \cdots & \mathbf{t}_{q^2 n} & c_{t,q^2} \end{bmatrix}$$

$$\Rightarrow \begin{pmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \mathbf{a}_{13} & \cdots & \mathbf{a}_{1n} & \tilde{c}_1 \\ 0 & 0 & \mathbf{a}_{23} & \cdots & \mathbf{a}_{2n} & \tilde{c}_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \\ \mathbf{a}_{m1} & \mathbf{a}_{m2} & \mathbf{a}_{m3} & \cdots & \mathbf{a}_{mn} & c_m \end{pmatrix}$$

Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In: *Journal of the ACM* 56.6 (Sept. 2009), 34:1–34:40. ISSN: 0004-5411 (print), 1557-735X (electronic). DOI: `http://doi.acm.org/10.1145/1568318.1568324`

Martin R. Albrecht et al. On the Complexity of the BKW Algorithm on LWE. Cryptology ePrint Archive, Report 2012/636. `http://eprint.iacr.org/2012/636`. 2012

Qian Guo, Thomas Johansson, and Paul Stankovski. Coded-BKW: Solving LWE Using Lattice Codes. Cryptology ePrint Archive, Report 2016/310. `http://eprint.iacr.org/2016/310`. 2016

Here's a Picture of a Kitten

Primal Lattice Attack (uSVP Version)

Bounded Distance Decoding and unique SVP

Given $\mathbf{A}, \mathbf{c}$ with $\mathbf{c} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$, we know that for some $\mathbf{w}$ we have that $\mathbf{A} \cdot \mathbf{w} - \mathbf{c} \pmod q$ is rather small.

In other words, we know there is an unusually short vector in the q-ary lattice

$$\mathbf{B} = \begin{pmatrix} \mathbf{A}^T & 0 \\ \mathbf{c}^T & t \end{pmatrix} \in \mathbb{Z}_q^{(n+1)\times(m+1)}$$

since

$$(\mathbf{s} \mid -1) \cdot \mathbf{B} = (\mathbf{e} \mid -t) \bmod q.$$

Let's find it.

Constructing a Basis

- Compute reduced row echelon form $[\mathsf{I}_{n\times n} \mid \mathsf{A}']$ of $\mathsf{A}^T \in \mathbb{Z}_q^{n\times m}$ with $m > n$.
- Stack on top of $[\mathsf{0}_{(m-n)\times n} \mid q\,\mathsf{I}_{(m-n)\times(m-n)}]$ to handle modular reductions
- Stack on top of $[\mathsf{c}^T \mid t]$
- Obtain

$$\mathsf{B} = \begin{pmatrix} \mathsf{I}_{n\times n} & \mathsf{A}' & 0 \\ \mathsf{0}_{(m-n)\times n} & q\,\mathsf{I}_{(m-n)\times(m-n)} & 0 \\ \mathsf{c}^T & & t \end{pmatrix} \in \mathbb{Z}^{(m+1)\times(m+1)}$$

- In practice, we always pick $t = 1$

HSVP vs uSVP

- Any algorithm which can solve κ-HSVP, such as a lattice reduction algorithm, can be used linearly many times to solve γ-uSVP with approximation factor $\gamma = \kappa^2$.[10]
- Whenever $\kappa > \sqrt{d}$ then any algorithm solving κ-HSVP can be used to solve γ-uSVP for $\gamma \approx \sqrt{d}\kappa$.[11]

[10] László Lovász. An algorithmic theory of numbers, graphs and convexity. CBMS-NSF regional conference series in applied mathematics. Philadelphia, Pa. Society for Industrial and Applied Mathematics, 1986. ISBN: 0-89871-203-3. URL: http://opac.inria.fr/record=b1086067.

[11] Cong Ling, Shuiyin Liu, Laura Luzzi, and Damien Stehlé. Decoding by embedding: Correct decoding radius and DMT optimality. In: *2011 IEEE International Symposium on Information Theory Proceedings, ISIT*. ed. by Alexander Kuleshov, Vladimir Blinovsky, and Anthony Ephremides. IEEE, 2011, pp. 1106–1110. DOI: 10.1109/ISIT.2011.6033703.

HSVP vs uSVP

- Any algorithm which can solve κ-HSVP, such as a lattice reduction algorithm, can be used linearly many times to solve γ-uSVP with approximation factor $\gamma = \kappa^2$.[10]
- Whenever $\kappa > \sqrt{d}$ then any algorithm solving κ-HSVP can be used to solve γ-uSVP for $\gamma \approx \sqrt{d}\kappa$.[11]

In practice

Algorithms behave better.

[10] László Lovász. An algorithmic theory of numbers, graphs and convexity. CBMS-NSF regional conference series in applied mathematics. Philadelphia, Pa. Society for Industrial and Applied Mathematics, 1986. ISBN: 0-89871-203-3. URL: http://opac.inria.fr/record=b1086067.

[11] Cong Ling, Shuiyin Liu, Laura Luzzi, and Damien Stehlé. Decoding by embedding: Correct decoding radius and DMT optimality. In: *2011 IEEE International Symposium on Information Theory Proceedings, ISIT*. ed. by Alexander Kuleshov, Vladimir Blinovsky, and Anthony Ephremides. IEEE, 2011, pp. 1106–1110. DOI: 10.1109/ISIT.2011.6033703.

Success Condition (2008)

Lattice reduction is expected/observed[12] to succeed if

$$\lambda_2/\lambda_1 \geq \tau \cdot \delta_0^d$$

where $\tau \approx 0.3$ is a constant that depends on the algorithm.

[12] Nicolas Gama and Phong Q. Nguyen. Predicting Lattice Reduction. In: *EUROCRYPT 2008*. Ed. by Nigel P. Smart. Vol. 4965. LNCS. Springer, Heidelberg, Apr. 2008, pp. 31–51.

Success Condition (2013, 2016)

- We can predict the length of the unusually short vector:

$$\lambda_1(\mathrm{B}) \approx \sqrt{m} \cdot \sigma.$$

- In general, we expect no other unusually short vectors, so we may assume[13]

$$\lambda_2(\mathrm{B}) \approx \sqrt{\frac{d}{2\,\pi, e}} \cdot \mathrm{Vol}(\mathrm{B})^{1/d}.$$

[13] Martin R. Albrecht, Robert Fitzpatrick, and Florian Gopfert. On the Efficacy of Solving LWE by Reduction to Unique-SVP. Cryptology ePrint Archive, Report 2013/602. http://eprint.iacr.org/2013/602. 2013; Florian Göpfert. Securely Instantiating Cryptographic Schemes Based on the Learning with Errors Assumption. http://tuprints.ulb.tu-darmstadt.de/5850/. PhD thesis. Technische Universität Darmstadt, 2016.

Success Condition (2015)

Lemma[14]

Given an LWE instance characterised by n, α, q. Any lattice reduction algorithm achieving log root-Hermite factor

$$\log \delta_0 = \frac{\log^2\left(\varepsilon' \tau \alpha \sqrt{2e}\right)}{4n \log q}$$

solves LWE with success probability greater than $\varepsilon_\tau \cdot \left(1 - \left(\varepsilon' \cdot \exp\left(\frac{1-\varepsilon'^2}{2}\right)\right)^m\right)$ for some $\varepsilon' > 1$ and some fixed $\tau \leq 1$, and $0 < \varepsilon_\tau < 1$ as a function of τ.

This lemma assumes $m = \sqrt{\frac{n \log q}{\log \delta_0}}$ which maximises the gap.

[14] Martin R. Albrecht, Rachel Player, and Sam Scott. On The Concrete Hardness Of Learning With Errors. Cryptology ePrint Archive, Report 2015/046. http://eprint.iacr.org/2015/046. 2015.

Success Condition (2016)

- Let $\mathbf{e}^*_{d-b}$ be the projection of $\mathbf{e}$ orthogonally onto the first $d-b$ vectors of the Gram-Schmidt basis $\mathbf{B}^*$
- BKZ-like algorithms will call an SVP oracle on th last block of dimension b.
- If $\mathbf{e}^*_{d-b}$ is a shortest vector in that block, it will be found
- If $\mathbf{e}^*_i$ is a shortest vector for all projections up to $d-b$ it will "travel to the front".

Success Condition (2016)

- Assume $\|\mathbf{e}^*_{d-b}\| \approx \sigma \cdot \sqrt{b}$.
- Applying the GSA, we expect the shortest vector to be found in the last block to have norm

$$\begin{aligned}\|\mathbf{b}^*_{d-b+1}\| &= \alpha^{d-b} \cdot \delta_0^d \cdot \mathrm{Vol}(\mathrm{B})^{1/d} \\ &= \delta_0^{-2(d-b)} \cdot \delta_0^d \cdot \mathrm{Vol}(\mathrm{B})^{1/d} \\ &= \delta_0^{2b-d} \cdot \mathrm{Vol}(\mathrm{B})^{1/d}.\end{aligned}$$

- Thus[15] we expect success if

$$\sigma \cdot \sqrt{b} \leq \delta_0^{2b-d} \cdot \mathrm{Vol}(\mathrm{B})^{1/d}$$

[15] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - a new hope. Cryptology ePrint Archive, Report 2015/1092. http://eprint.iacr.org/2015/1092. 2015.

Success Condition (2016)

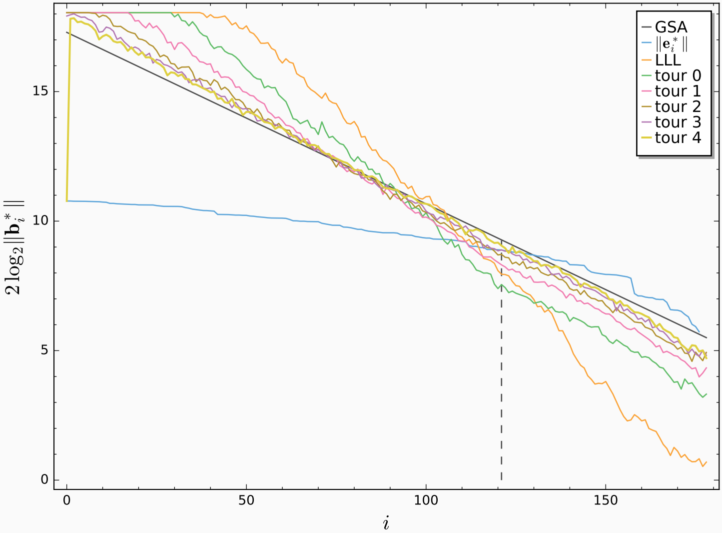

Comparison for $q = 2^{15}, \sigma = 3.2$

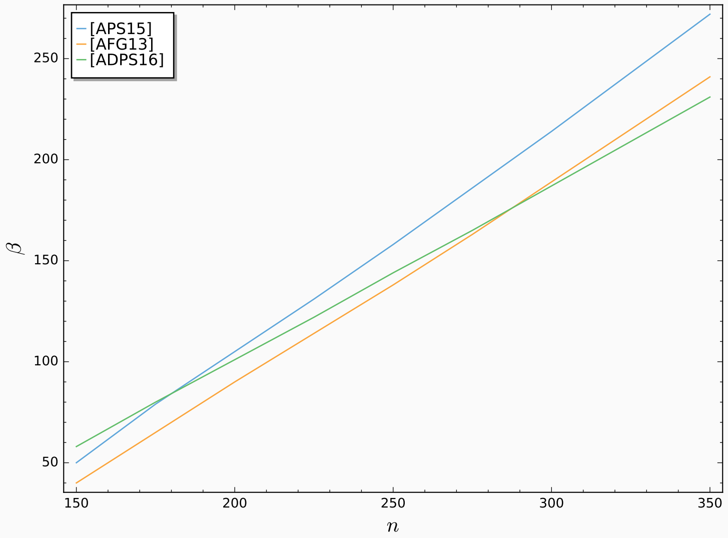

LWE Normal Form

- Consider the lattice

$$\Lambda = \{\mathbf{v} \in \mathbb{Z}^{n+m+1} | (\mathbf{A}|\mathbf{I}_m|\mathbf{c}) \cdot \mathbf{v} \equiv 0 \pmod q\}$$

- It contains an unusually short vector $(\mathbf{s}|\mathbf{e}|-1)$ since

$$(\mathbf{A}|\mathbf{I}_m|\mathbf{c}) \cdot (\mathbf{s}|\mathbf{e}|-1) \equiv \mathbf{A} \cdot \mathbf{s} + \mathbf{e} - \mathbf{c} \equiv 0 \pmod q$$

- Analysis proceeds as before with $d = n + m + 1$.

Small Secrets

- Let σ be the standard deviation of the components of $\mathbf{e}$.
- When $\|\mathbf{s}\| \ll \|\mathbf{e}\|$, the vector $(\mathbf{s}||\mathbf{e})$ is uneven in length.
- Rescale the first part to have the same norm as the second.[16]
  - When $\mathbf{s}_i \leftarrow_\$ \{-1, 0, 1\}$, the volume of the lattice is scaled by σ^n.
  - When $\mathbf{s} \leftarrow_\$ \{0, 1\}$ the volume of the lattice is scaled by $(2\sigma)^n$ because we can scale by 2σ and then rebalance.

[16] Shi Bai and Steven D. Galbraith. Lattice Decoding Attacks on Binary LWE. In: *ACISP 14*. Ed. by Willy Susilo and Yi Mu. Vol. 8544. LNCS. Springer, Heidelberg, July 2014, pp. 322–337. DOI: 10.1007/978-3-319-08344-5_21.

Thank You