The art of cryptography, summer 2013

Lattices and cryptography

Prof. Dr. Joachim von zur Gathen

α	$\log \alpha$	class
$2^{n \log^2 \log n / \log n}$	$n \log^2 \log n / \log n$	P
$2^{n \log \log n / \log n}$	$n \log \log n / \log n$	BPP
$\vdots$	$\vdots$	$\vdots$
$\sqrt{n}$	$\frac{1}{2} \log n$	NP $\cap$ coNP not NP-hard
$\sqrt{\frac{n}{\log n}}$	$\frac{1}{2}(\log n - \log \log n)$	NP $\cap$ coAM not NP-hard
$\vdots$	$\vdots$	$\vdots$
$n^{1/\log \log n}$	$\log n / \log \log n$	hard
1	0	NP-hard (random)

Table : Complexity of α-approximations to SVP.

We define below a problem called *learning with errors* (LWE). The idea is that we are given positive integers q and n, several (a, b') with uniformly and independently chosen $a \xleftarrow{\text{\$}} \mathbb{Z}_q^n$ and $b' \in \mathbb{Z}_q$, and want to find $u \in \mathbb{Z}_q^n$ under the guarantee that the errors

$$v = b' - a \star u \in \mathbb{Z}_q$$

follow a Gaussian distribution.

For a positive integer n and positive real r, the Gaussian function $\gamma_r^{(n)}$ is

$$\begin{aligned} \gamma_r^{(n)} \colon \mathbb{R}^n &\longrightarrow \mathbb{R}, \\ x &\longmapsto e^{-\pi(\|x\|/r)^2}. \end{aligned}$$

The total volume of $\mathbb{R}^n$ under $\gamma_r^{(n)}$ is

$$\int_{\mathbb{R}^n} \gamma_r^{(n)}(x)\mathsf{d}x = r^n.$$

Thus we can define the continuous Gaussian distribution $\mathcal{G}_r^{(n)}$ on $\mathbb{R}^n$ by its density $\rho_r^{(n)}(x) = r^{-n} \cdot \gamma_r^{(n)}(x)$. Then $\mathcal{G}_r^{(n)}(A) = r^{-n} \int_A \rho_r^{(n)}(x)dx$ for a measurable set $A \subseteq \mathbb{R}^n$ is the probability that some $x \in A$ is chosen if $x \xleftarrow{\$} \mathcal{G}_r^{(n)}$. We abbreviate $\mathcal{D}_{s,\mathcal{G}_r^{(1)}}$ as $\mathcal{D}_{s,r}$.

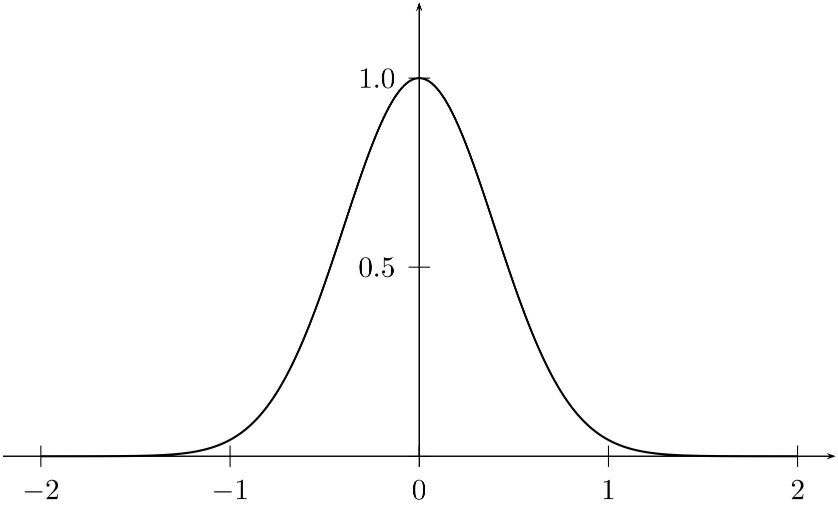
1.0
0.5
−2
−1
0
1
2

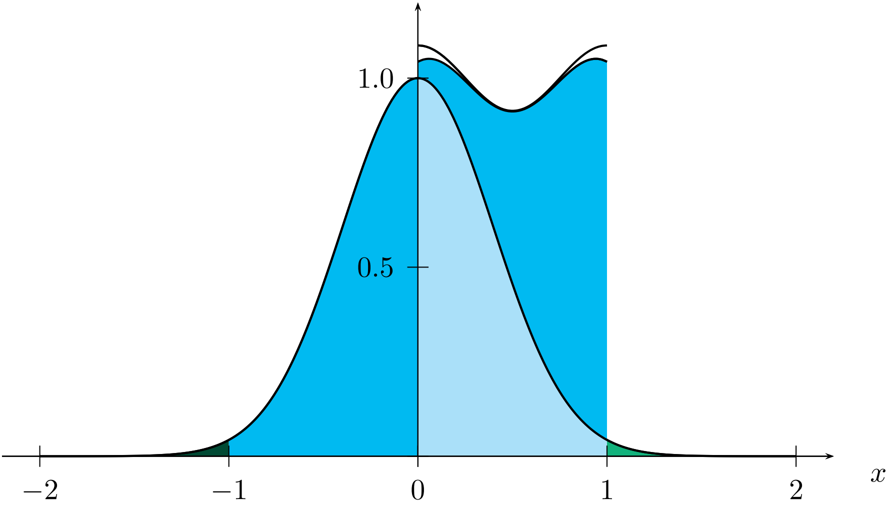
1.0
0.5
−2
−1
0
1
2
x

Definition 1. *Let* $q, r : \mathbb{N} \longrightarrow \mathbb{R}$ *with integral* $q(n) \geq 2$ *and* $r(n) > 0$ *for all* n*. An algorithm solves the learning with errors problem LWE*$_{s,r}$ *if it determines* $s \in \mathbb{Z}^n_{q(n)}$ *with overwhelming probability, given access to any number, polynomial in* n*, of samples* $(a, b) \in \mathbb{Z}^n_{q(n)} \times \mathbb{T}$ *according to* $\mathcal{D}_{s,r}$*.*

Stage 1: reduction (n/r)-GapSVP $\leq_p$ LWE,

Stage 2: reduction LWE $\leq_p$ DLWE,

Stage 3: LWE-based cryptosystem.

Definition 2. *For a function* $\alpha\colon \mathbb{N} \longrightarrow \mathbb{R}$ *with* $\alpha(n) \geq 1$ *for all* n*, we define the* α*-gap shortest vector problem* α*-GapSVP as follows. Input is a basis* A *of an* n*-dimensional lattice* L *and a positive real number* d*. The answer is*

$$\begin{cases} \textit{yes} & \textit{if } \lambda_1(L) \leq d, \\ \textit{no} & \textit{if } \lambda_1(L) \geq \alpha(n) \cdot d. \end{cases}$$

When $d < \lambda_1(L) < \alpha(n) \cdot d$*, any answer is permitted.*

Definition 3. *For functions* $\alpha, \beta\colon \mathbb{N} \longrightarrow \mathbb{R}$ *with* $\beta(n) \geq \alpha(n) \geq 1$ *for all* n*, we define the* β*-to-*α*-gap shortest vector problem* α*-to-*β*-GapSVP as follows. Input is a basis* A *of an* n*-dimensional lattice* L *in* $\mathbb{R}^n$ *with GSO* $(a_1^*, \ldots, a_n^*)$ *and a positive integer* d *so that*

i. $\lambda_1(L) \leq \beta(n)$,

ii. $\|a_i^*\| \geq 1$ *for* $1 \leq i \leq n$,

iii. $1 \leq d \leq \beta(n)/\alpha(n)$.

The answer is, as in Definition 2,

$$\begin{cases} \textit{yes} & \textit{if } \lambda_1(L) \leq d, \\ \textit{no} & \textit{if } \lambda_1(L) \geq \alpha(n) \cdot d. \end{cases}$$

DEFINITION 4. *For functions* $\alpha, \beta \colon \mathbb{N} \longrightarrow \mathbb{R}$ *with* $\beta(n) \geq \alpha(n) \geq 1$ *for all* n*, we define the* β*-to-*α*-gap shortest vector problem* α*-to-*β*-GapSVP as follows. Input is a basis* A *of an* n*-dimensional lattice* L *in* $\mathbb{R}^n$ *with GSO* $(a_1^*, \ldots, a_n^*)$ *and a positive integer* d *so that*

i. $\lambda_1(L) \leq \beta(n)$,

ii. $\|a_i^*\| \geq 1$ *for* $1 \leq i \leq n$,

iii. $1 \leq d \leq \beta(n)/\alpha(n)$.

The answer is, as in Definition 2,

$$\begin{cases} \textit{yes} & \textit{if } \lambda_1(L) \leq d, \\ \textit{no} & \textit{if } \lambda_1(L) \geq \alpha(n) \cdot d. \end{cases}$$

LEMMA 5. *For any* $c, d > 0$ *and* $z \in \mathbb{R}^n$ *with* $\|z\| \leq d$, *and* $d' = d\sqrt{cn/\log n}$, *we have*

$$\Delta(\mathcal{U}_{d'\mathcal{B}_n}, \mathcal{U}_{z+d'\mathcal{B}_n}) \leq 1 - \frac{1}{\mathsf{poly}(n)}.$$

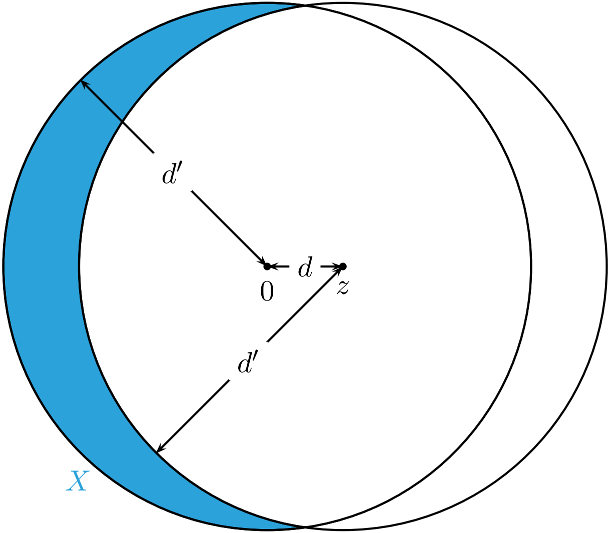

Figure : Δ of two shifted balls.

LEMMA 6. *There is a probabilistic polynomial-time algorithm that takes as input a basis A of an n-dimensional lattice L and some $r > \max\{\|a_i^*\| \colon 1 \leq i \leq n\} \cdot \omega(\sqrt{\log n})$. As output it produces samples from a distribution whose statistical distance to $\mathcal{G}_{L,r}$ is negligible in n.*

Definition 7. *Let L be an n-dimensional lattice and $\epsilon > 0$. The smoothing parameter $\eta_\epsilon(L)$ is the smallest s so that*

$$\rho_{1/s}^{(n)}(L^* \setminus \{0\}) = \sum_{x \in L^* \setminus \{0\}} \rho_{1/s}^{(n)}(x) \leq \epsilon.$$

LEMMA 8. *Let L be an n-dimensional lattice and $\epsilon, c > 0$.*

i. *If $s' > \eta_\epsilon(L)$, then $\rho^{(n)}_{1/s'}(L^* \setminus \{0\}) \leq \epsilon$.*

ii. $\eta_\epsilon(c \cdot L) = c \cdot \eta_\epsilon(L)$.

iii. $\eta_{2^{-n}}(L) \leq \frac{\sqrt{n}}{\lambda_1(L^*)}$.

iv. *For any function f with $f(n) = \omega(\sqrt{\log n})$, there exists a negligible function ϵ so that $\eta_{\epsilon(n)}(\mathbb{Z}) \leq f(n)$.*

v. *If $0 < \epsilon < 1$, $r \geq \eta_\epsilon(L)$ and $d \in \mathbb{R}^n$, then*

$$\frac{1-\epsilon}{1+\epsilon} \leq \frac{\rho^{(n)}_r(L+d)}{\rho^{(n)}_r(L)} \leq 1.$$

PROPOSITION 9. *Let* $\gamma, \epsilon, q \colon \mathbb{N} \longrightarrow \mathbb{R}_{>0}$ *be functions with* $\gamma(n) < 1, \epsilon$ *negligible, and* $q(n) \geq 2$ *an integer for all* n*. There exists a reduction* $\mathcal{R}$ *that takes as input a basis* A *of a lattice* $L \subseteq \mathbb{R}^n$*, real* $r \geq \sqrt{2}q(n) \cdot \eta_{\epsilon(n)}(L^*)$ *and* $z \in \mathbb{R}^n$ *with* $d(z, L) \leq \gamma(n)q(n)/\sqrt{2}r < \lambda_1(L)/2$*. It makes use of two subroutines* W *and* D*, where* W *solves* $\mathsf{LWE}_{q(n),\gamma(n)}$ *using polynomially in* n *many samples, and* D *generates samples from* $\mathcal{G}_{L^*,r}$*. The output is with overwhelming probability (the unique)* $x \in L$ *closest to* z*.*

Algorithm 10. Reduction from β-to-α-GapSVP to LWE.

Input: A basis A of an n-dimensional lattice L, and $d \geq 1$.
Output: "yes" or "no".

1. Choose a large N, polynomial in n.
2. Do step 3 through 7 N times.
3. $d' \longleftarrow d \cdot \sqrt{n/(4 \log n)}$.
4. Choose w uniformly at random in the ball $d' \cdot \mathcal{B}_n = \{u \in \mathbb{R}^n \colon \|u\| \leq d'\}$.
5. $x \longleftarrow w \operatorname{srem} L$.
6. Call the reduction $\mathcal{R}$ from Proposition 9 with input A, x and
$$r = \frac{q\sqrt{2n}}{\alpha d}.$$
The sampler for $\mathcal{G}_{L^*,r}$ is implemented by the algorithm from Lemma 6 on the reversed dual basis D of L^*. Let v be the output of $\mathcal{R}$.
7. If $v \neq x - w$, then return "yes".
8. Return "no".

Theorem 11. *Let* α, β, γ, $q\colon \mathbb{N} \longrightarrow \mathbb{R}_{>0}$ *be such that* $\gamma(n) < 1$, $\alpha(n) \geq n/(\gamma(n)\sqrt{\log n})$, $\beta(n) \geq \alpha(n)$, $q(n) \in \mathbb{Z}$, *and* $q(n) \geq \beta(n) \cdot \omega(\sqrt{n^{-1} \log n})$ *for all* n. *Then Algorithm 10 provides a probabilistic polynomial time reduction from solving worst-case* β*-to-*α*-GapSVP with overwhelming probability to solving* $\mathsf{LWE}_{q(n),\gamma(n)}$ *with polynomially in* n *many samples.*

LEMMA 12. *Let* $q, \alpha \colon \mathbb{N} \longrightarrow \mathbb{R}$ *with* $0 < \alpha(n) < 1$ *and all prime factors* p *of the squarefree* n*-bit integer* $q(n)$ *satisfying* $\omega(\sqrt{\log n})/\alpha(n) \leq p \leq \mathsf{poly}(n)$. *Then there is a probabilistic polynomial-time reduction from solving* $\mathsf{LWE}_{q(n),\alpha}$ *with overwhelming probability to distinguishing between* $\mathcal{D}_{s,\alpha}$ *and* $\mathcal{U}(\mathbb{Z}_{q(n)}^n \times \mathbb{T})$ *for unknown* $s \in \mathbb{Z}_{q(n)}^n$ *with overwhelming advantage.*

LEMMA 13. *Let* $q\colon \mathbb{N} \longrightarrow \mathbb{N}_{\geq 2}$, *let* $\mathcal{C}$ *be a distribution on* $\mathbb{T}$, *and* $\mathcal{U}_n = \mathcal{U}_{\mathbb{Z}^n_{q(n)} \times \mathbb{T}}$. *There is a probabilistic polynomial time reduction from distinguishing between* $\mathcal{D}_{s,\mathcal{C}}$ *and* $\mathcal{U}_n$ *for an arbitrary* $s \in \mathbb{Z}^n_{q(n)}$ *with overwhelming advantage to distinguishing between* $\mathcal{D}_{t,\mathcal{C}}$ *and* $\mathcal{U}_n$ *for uniformly random* $t \xleftarrow{\text{⚄⚄}} \mathbb{Z}^n_{q(n)}$ *with nonnegligible advantage.*

For simplicity we write q instead of $q(n)$. We now construct a trapdoor function based on lattices. For starters, we consider matrices $A \in \mathbb{Z}_q^{n\times\ell}$ and their (left) kernel

$$\text{lker}\, A = \{x \in \mathbb{Z}_q^n \colon xA = 0 \text{ in } \mathbb{Z}_q^\ell\}.$$

We always have $0 = (0, \ldots, 0) \in \ker A$. Notions like kernel and rank are well understood when q is prime, so that $\mathbb{Z}_q$ is a field. For general q, we have following bound.

LEMMA 14. *Let* $\ell \geq n \geq 1$, $q \geq 2$, $\delta > 0$, *and*

$$p = \mathsf{prob}\{\mathsf{Iker}\, A \neq \{0\} \colon A \overset{\text{⚂⚄}}{\longleftarrow} \mathcal{U}_{\mathbb{Z}_q^{n \times \ell}}\}.$$

Then $p < q^n \cdot 2^{-\ell}$.

Given q and $A \in \mathbb{Z}_q^{n\times\ell}$, we define two lattices:

$$\begin{aligned}\Lambda(A) &= \{x \in \mathbb{Q}^\ell \colon q\cdot x \in \mathbb{Z}^\ell,\ \exists s \in \mathbb{Z}_q^n \quad q\cdot x = sA \text{ in } \mathbb{Z}_q^\ell\},\\ \Lambda^\perp(A) &= \{y \in \mathbb{Z}^\ell \colon Ay = 0 \text{ in } \mathbb{Z}_q^n\}.\end{aligned}$$

Then $\mathbb{Z}^\ell \subseteq \Lambda(A)$ and $q\mathbb{Z}^\ell \subseteq \Lambda^\perp(A)$, and the two lattices are duals of each other.

We use an algorithm that generates an almost uniform A together with a "trapdoor" basis T of $\Lambda^{\perp}(A)$, whose vectors are fairly short.

FACT 15. *There is a probability polynomial-time algorithm which on input n in unary, odd $q \geq 3$, and $\ell \geq 6n \log_2 q$ with $\ell \in \mathsf{poly}(n)$, outputs a pair (A, T) of matrices with the following properties.*

i. $A \in \mathbb{Z}_q^{n \times \ell}$ *is distributed within negligible (in n) statistical distance from uniform,*
ii. $T \in \mathbb{Z}^{\ell \times \ell}$ *is a basis of $\Lambda^{\perp}(A)$,*
iii. *there is some $C \in O(\sqrt{n \log_2 q})$ so that each row of the GSO basis T^* has norm at most C.*

We now have the following trapdoor function, including the family $\{g_A \colon \mathbb{Z}_q^n \longrightarrow \mathbb{T}_{q'}^\ell\}_{n\in\mathbb{N}}$, where we leave out the argument n in most places. The integers $q, q' \geq 2$ and real $r > 0$ are further parameters.

- gen: Run the algorithm from Fact 15 to generate a function index $A \in \mathbb{Z}_q^{n\times\ell}$ and a trapdoor basis $T \in \mathbb{Z}^{\ell\times\ell}$.
- eval(A, s): Obtain $x \xleftarrow{\text{⚄}} \mathcal{G}_r^{(\ell)}$ and output

$$b = g_A(s, x) = \lfloor (sA)/q + x \rceil_{q'} \in \mathbb{T}_{q'}^\ell. \tag{16}$$

- inv(T, z): Run the nearest hyperplane algorithm with input z to find some $y \in \Lambda(A)$ with $\|z - y\| \leq 2^{n-1} d(z, \Lambda(A))$. Compute $s \in \mathbb{Z}_q^n$ with $(sA)/q = y$ in $\mathbb{T}$.

THEOREM 17. *Let* $A \in \mathcal{A}_q^{n\times\ell}$, $q' \geq 2C\sqrt{\ell}$, *and* $r^{-1} \geq C \cdot \omega(\sqrt{\log n})$. *For any* $s \in \mathbb{Z}_q^n$, *the algorithm inv, on input* (T, b) *with* $b = \lfloor (sA)/q + x \rceil_{q'} \in \mathbb{T}_{q'}^{\ell}$, *outputs* s *with overwhelming probability over the choice of* $x \xleftarrow{\$} \mathcal{G}_r^{(\ell)}$.

- *Correctness.* For every $s \in D_n$ and $b \xleftarrow{\text{⚄⚅}} g_a(s)$, $\mathsf{ver}(a, s, b)$ accepts with overwhelming probability over the random parameter $x \in X_n$.
- *Unique preimage.* For every $b \in R_n$ there is at most one $s \in D_n$ so that $\mathsf{ver}(a, s, b)$ accepts.
- *Findable preimage.* For every $s \in D_n$ and $b \in R_n$ with $\mathsf{ver}(a, s, b)$ accepting, we have $\mathsf{inv}(t, b) = s$.

Peikert cryptosystem key generation 18.

Input: n in unary.

Output: Public key pk and secret key sk.

1. $U \xleftarrow{\$} \mathbb{Z}_q^{n\times\ell}$.
2. For $1 \leq i \leq k$ and $b \in \{0,1\}$ do
3. $(A_{i,b}, T_{i,b}) \xleftarrow{\$} T.\mathsf{gen}(n)$.
4. Output $\mathsf{pk} = (\{A_{i,b} \colon 1 \leq i \leq k, b \in \{0,1\}\}, U)$ and $\mathsf{sk} = (T_{1,0}, T_{1,1})$.

Peikert cryptosystem encapsulation 19.

Input: pk.

Output: encap(pk).

1. $(S.\mathsf{pk}, S.\mathsf{sk}) \xleftarrow{🎲} S.\mathsf{gen}(n)$.
2. $y \xleftarrow{🎲} \{0,1\}^j$, $s \xleftarrow{🎲} \mathbb{Z}_q^n$ uniformly, $x_0 \xleftarrow{🎲} \mathcal{G}_r^{(j)}$.
3. $b_0 \longleftarrow \lfloor (sU)/q + x_0 + y/2 \rceil_{q'} \in \mathbb{T}_{q'}^{\ell}$.
4. For $1 \leq i \leq k$ do.
5. indent $b_i \xleftarrow{🎲} T.\mathsf{eval}(A_i, (s.\mathsf{pk})_i, s) \in \mathbb{T}_{q'}^{\ell}$.
6. $b \longleftarrow (b_0, b_1, \ldots, b_k) \in \mathbb{T}_{q'}^{k\ell+j}$.
7. $\sigma \longleftarrow S.\mathsf{sign}(S.\mathsf{sk}, b)$.
8. Output $\tau = (S.\mathsf{pk}, b, \sigma)$.

Peikert cryptosystem decapsulation 20.

Input: sk, τ.

Output: an element of $\{0,1\}^\ell$ or "failure".

1. Write $b = (b_0, b_1, \ldots, b_k)$ with $b_0 \in \mathbb{T}^j_{q'}$ and $b_i \in \mathbb{T}^\ell_{q'}$ for $1 \le i \le k$. If b cannot be parsed in this way, then return "failure".
2. Verify the signature by running $S.\mathsf{ver}$ on τ. If this is rejected, then return "failure".
3. $s \longleftarrow T.\mathsf{inv}(T_{1,(S.\mathsf{sk})_1}, b_1) \in \mathbb{Z}_q^n$.
4. For $1 \le i \le k$ do
5. Run $T.\mathsf{ver}$ on $(A_{i,S.\mathsf{pk}}, s, b_i)$. If $T.\mathsf{ver}$ rejects, then return "failure".
6. $h \longleftarrow b_0 - (sU)/q \in \mathbb{T}^j = [0,1)^j$.
7. For $1 \le i \le j$ do 8–9
8. $y_i \longleftarrow 1$.
9. If $h_i \in [0,1/4) \cup [3/4,1)$ then $y_i \longleftarrow 0$.
10. Return $y = (y_1, \ldots, y_j) \in \{0,1\}^j$.

LEMMA 21. *The decapsulation procedure works correctly with overwhelming probability.*

THEOREM 22. *Assume that the signature scheme* S *is strongly unforgeable under one-time chosen message attacks, and that for* $s \xleftarrow{\text{🎲}} \mathcal{U}_{\mathbb{Z}_q^n}$, $G_{s,r}$ *is pseudorandom. Then the above key encapsulation mechanism is indistinguishable under chosen message attacks.*