Introduction to modern lattice-based cryptography (Part I)

Damien Stehlé

LIP – CNRS/ENSL/INRIA/UCBL/U. Lyon

Singapore, June 2010

Modern lattice-based cryptography

- Cryptography: the science of information hiding.
- "Lattice-based": the schemes involve Euclidean lattices.
- Standard lattice problems provably reduce to attacks against those schemes.
- Modern: we won't be interested in GGH and NTRU.
  More recent schemes offer similar performance with rigorous security guarantees.

Why lattice-based cryptography?

(why not business as usual, with factoring and discrete log?)

- LBC provides unmatched security properties:
  its security stems from worst-case hardness assumptions.
- LBC seems to remain secure even against quantum computers.
- LBC is asymptotically extremely efficient.
- LBC is simple and flexible: this leads to easier design of complicated cryptographic functions.
- Diversity fosters cross-pollination.

Goal of this course

To give an overview of recent developments in LBC, and a flavour of the techniques/results.

Disclaimer: This is not a practical crypto course.

Contents: Complexity theory, distributions, quantum computing, cryptography, structured matrices, lattices.

Highlights: Worst-case to average-case reductions, encryption with quasi-optimal complexity, fully homomorphic encryption.

Bibliography

- The LLL Algorithm. Survey and Applications.
  P. Nguyen and B. Vallée (Eds.), Springer.
- The Learning with Errors Problem. Survey by O. Regev.
- Lattice-based Cryptography. Survey by D. Micciancio and O. Regev.
- Webpage of C. Peikert (including slides of several talks).

Plan

1- Background on Euclidean lattices.

2- The SIS problem, or how to hash.

3- The LWE problem, or how to encrypt.

4- Cryptanalysis.

5- Advanced topics: IBE and FHE.

Plan

1- **Background on Euclidean lattices**.

2- The SIS problem, or how to hash.

3- The LWE problem, or how to encrypt.

4- Cryptanalysis.

5- Advanced topics: IBE and FHE.

Background on Euclidean lattices

a- **Arbitrary lattices.**

b- Ideal lattices.

c- Lattice Gaussians.

(Arbitrary) lattices

Lattice $\equiv$ discrete subgroup of $\mathbb{R}^n$
$\equiv \{\sum_{i \le n} x_i \mathbf{b}_i : x_i \in \mathbb{Z}\}$

If the $\mathbf{b}_i$'s are linearly independent, they are called a **basis**.

Bases are not unique, but they can be obtained from each other by integer transforms of determinant ± 1:

$$\begin{bmatrix} -2 & 1 \\ 10 & 6 \end{bmatrix} = \begin{bmatrix} 4 & -3 \\ 2 & 4 \end{bmatrix} \cdot \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix}.$$

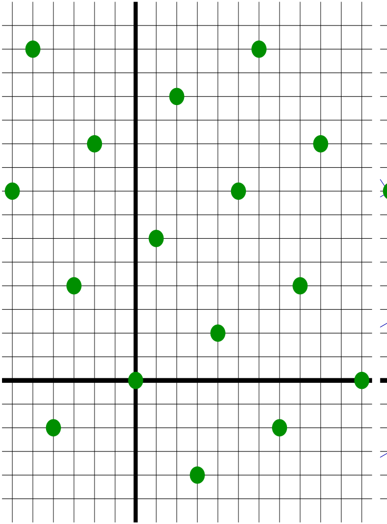

Lattice invariants

First minimum:
$\lambda = \min(\|\mathbf{b}\| : \mathbf{b} \in L \setminus \mathbf{0})$.

Successive minima:
$\lambda_k = \min(r : \dim \operatorname{span}(L \cap \mathcal{B}(r)) \geq k)$.

Lattice volume:
$\mathrm{vol}\ L = |\det(\mathbf{b}_i)_i|$, for any basis.

Minkowski theorem (1889):
$\lambda(L) \leq \sqrt{n} \cdot (\mathrm{vol}\ L)^{1/n}$.

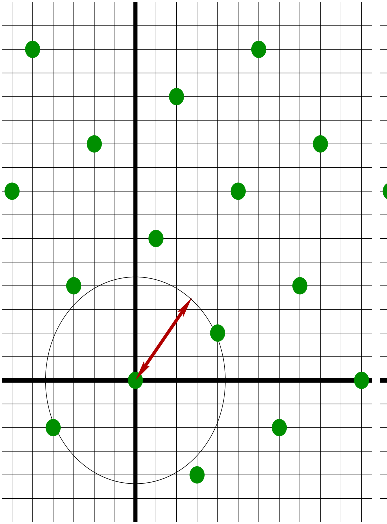

SVP and SIVP

The Shortest Vector Problem: SVP_γ

Given a basis of L, find $\mathbf{b} \in L \setminus \mathbf{0}$ such that: $\|\mathbf{b}\| \leq \gamma \cdot \lambda(L)$.

The Shortest Independent Vectors Problem: SIVP_γ

Given a basis of L, find $\mathbf{b}_1, \ldots, \mathbf{b}_n \in L$ lin. indep. such that:

$$\max \|\mathbf{b}_i\| \leq \gamma \cdot \lambda_n(L).$$

- NP-hard when $\gamma = O(1)$ (under randomized red.).
- In lattice-based crypto: $\gamma = \mathcal{P}oly(n)$ (most often).
- Solvable in polynomial time when $\gamma = 2^{\widetilde{O}(n)}$.

CVP

The Closest Vector Problem: CVP_γ

Given a basis of L and a target $\mathbf{t} \in \mathbb{Q}^n$, find $\mathbf{b} \in L$ such that:

$$\|\mathbf{b} - \mathbf{t}\| \leq \gamma \cdot \min(\|\mathbf{c} - \mathbf{t}\| : \mathbf{c} \in L).$$

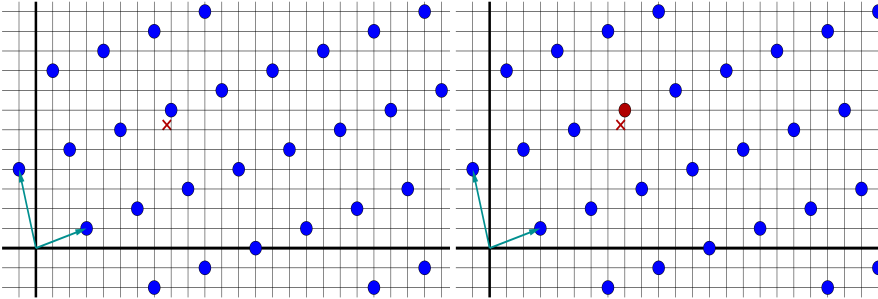

- NP-hard when $\gamma = O(1)$.

Gram-Schmidt Orthogonalisation

- A lattice may have infinitely many bases.
- Quality of a basis: measured by the GSO.

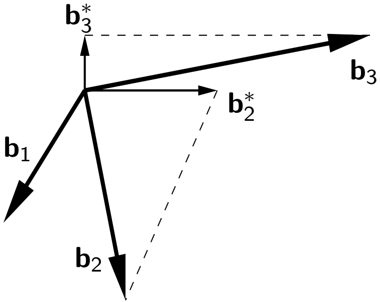

- $\mathbf{b}_i^* = \operatorname{argmin} \|\mathbf{b}_i + \sum_{j<i} \mathbb{R}\mathbf{b}_j\|$
- Quality measure: $\max_i \|\mathbf{b}_i^*\|$.

Properties of the GSO

GSO and basis vectors:

- For any i, $\|\mathbf{b}_i^*\| \leq \|\mathbf{b}_i\|$.
- Size-reduction: any basis can be efficiently transformed so that: $\max \|\mathbf{b}_i\| \leq \sqrt{n} \cdot \max \|\mathbf{b}_i^*\|$.

GSO and lattice invariants:

- $\mathrm{vol}\ L = \prod \|\mathbf{b}_i^*\|$, for any basis $(\mathbf{b}_i)$.
- Also, $\lambda(L) \geq \min \|\mathbf{b}_i^*\|$.

From short vectors to a short basis

- Let $(\mathbf{b}_i)_i$ be a basis of a lattice L.
- Let $(\mathbf{s}_i)_i$ in L be linearly independent with small GSO.
- Can we compute a basis of L with small GSO?

- Write $(\mathbf{s}_i)_i = (\mathbf{b}_i)_i \cdot T$, with $T \in \mathbb{Z}^{n \times n}$.
- Triangularize T, i.e., $T = U \cdot T'$ with $|\det U| = 1$ and $T' \in \mathbb{Z}^{n \times n}$ upper triangular. Let $(\mathbf{c}_i)_i = (\mathbf{b}_i)_i \cdot U$.
- $(\mathbf{c}_i)_i$ is a basis of L and $(\mathbf{s}_i)_i = (\mathbf{c}_i)_i \cdot T'$.
- Since T' is upper triangular: $\forall i, \|\mathbf{c}_i^*\| \leq \|\mathbf{s}_i^*\|$.

With a size-reduction, we get: $\max \|\mathbf{c}_i\| \leq \sqrt{n} \cdot \max \|\mathbf{s}_i\|$.

The dual lattice

- The dual of L is

$$\hat{L} = \left\{ \hat{\mathbf{b}} : \forall \mathbf{b} \in L, \langle \hat{\mathbf{b}}, \mathbf{b} \rangle \in \mathbb{Z} \right\}.$$

- B basis matrix of $L \implies B^{-T}$ basis matrix of $\hat{L}$.
- Let $B' = \text{reverse}(B^{-T})$. Then $\frac{1}{\|\mathbf{b'}^*_{n-i+1}\|} = \|\mathbf{b}_i^*\|$. Therefore:

$$\frac{1}{\min \|\mathbf{b'}_i^*\|} = \max \|\mathbf{b}_i^*\|.$$

Background on lattices

- **a-** Arbitrary lattices.
- **b-** **Ideal lattices.**
- **c-** Lattice Gaussians.

Ideal lattices

A lattice L is **ideal** if membership is preserved under negacyclic shifts of the coordinates:

$$\begin{array}{cccccccccc} & (& b_0 & b_1 & b_2 & b_3 & \ldots & b_{n-2} & b_{n-1} &) & \in L \\ \Rightarrow & (& -b_{n-1} & b_0 & b_1 & b_2 & \ldots & b_{n-3} & b_{n-2} &) & \in L \\ \Rightarrow & (& -b_{n-2} & -b_{n-1} & b_0 & b_1 & \ldots & b_{n-4} & b_{n-3} &) & \in L \\ \Rightarrow & (& -b_{n-3} & -b_{n-2} & -b_{n-1} & b_0 & \ldots & b_{n-5} & b_{n-4} &) & \in L \end{array}$$

A lattice L is **ideal** if it is an ideal of $\mathbb{Z}[x]/(x^n+1)$.

Easy property: all minima of an ideal lattice are equal.

$$\lambda_1(L) = \lambda_2(L) = \ldots = \lambda_n(L).$$

How special are ideal lattices?

Advantages

- The negacyclic structure allows one to save space.
  Warning: an ideal lattice may have no negacyclic basis.
- We can multiply vectors together.
- Fast polynomial arithmetic.

Drawbacks

- NP-hardness results not valid anymore.
- Decisional SVP becomes easier.

But no known computational advantage for Id-SVP/Id-SIVP.

Decisional SVP becomes easier

- Decisional SVP$_\gamma$ consists in approximating $\lambda(L)$.
- Minkowski: $\lambda \leq \sqrt{n} \cdot (\mathrm{vol}\ L)^{1/n}$.
- Let $(\mathbf{b}_i)_i$ be a basis, and $(\mathbf{s}_i)_i$ be lin. indep. vectors reaching the λ_i's: $\|\mathbf{s}_i\| = \lambda_i = \lambda$.
- Since $(\mathbf{s}_i)_i = (\mathbf{b}_i)_i \cdot T$, with $T \in \mathbb{Z}^{n \times n}$:

$$\mathrm{vol}\ L = |\det(\mathbf{b}_i)| \leq |\det(\mathbf{s}_i)| \leq \prod \|\mathbf{s}_i\| = \lambda^n.$$

- Overall: $1 \leq \dfrac{\lambda}{(\mathrm{vol}\ L)^{1/n}} \leq \sqrt{n}$.

Ideal lattices are famous objects

- A lattice L is **ideal** if it is an ideal of $\mathbb{Z}[x]/(x^n+1)$.
- We choose $n = 2^k$, making $x^n + 1$ irreducible.
- We play with the $2n$-th cyclotomic number field.
- We could use other number fields.
- These ideals have been studied for decades in the field of **algebraic number theory**.

Background on lattices

a- Arbitrary lattices.

b- Ideal lattices.

c- **Lattice Gaussians.**

A handy distribution: the discrete Gaussian

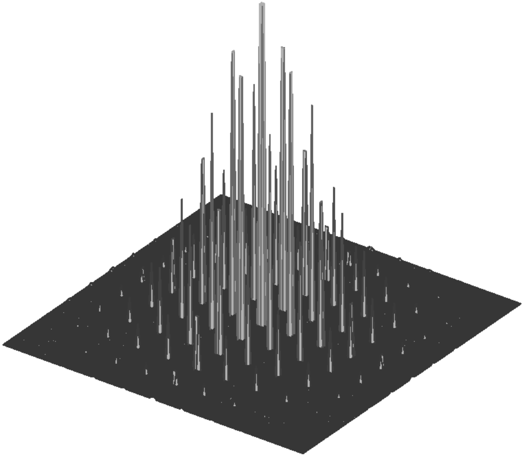

A handy distribution: the discrete Gaussian

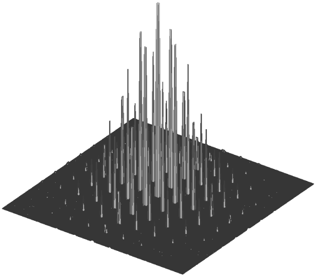

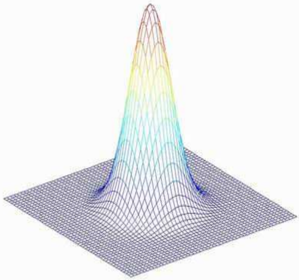

A discrete Gaussian is a discretization of a continuous Gaussian, with support being a lattice.

A handy distribution: the discrete Gaussian

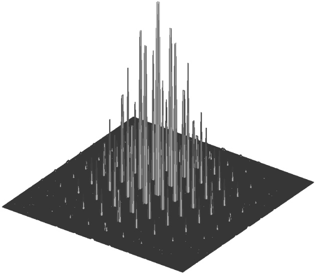

For $\mathbf{b} \in \mathbb{R}^n$ and $\mathbf{c} \in \mathbb{R}^n$:

$$\rho_{\sigma,\mathbf{c}}(\mathbf{b}) := e^{-\pi \frac{\|\mathbf{b}-\mathbf{c}\|^2}{\sigma^2}}.$$

σ is the standard deviation.

For $L \subseteq \mathbb{R}^n$ and $\mathbf{c} \in \mathbb{R}^n$: $\rho_{\sigma,\mathbf{c}}(L) = \sum_{\mathbf{b}\in L} \rho_{\sigma,\mathbf{c}}(\mathbf{b})$ is finite.
Discrete n-dimensional Gaussian:

$$\forall \mathbf{b} \in L : D_{L,\sigma,\mathbf{c}}(\mathbf{b}) = \frac{\rho_{\sigma,\mathbf{c}}(\mathbf{b})}{\rho_{\sigma,\mathbf{c}}(L)}.$$

Why are discrete Gaussians interesting?

- This is a lattice invariant.
- We can do **Fourier analysis** for lattice distributions, and (discrete) Gaussians interact nicely with (discrete) Fourier transforms.
- Many properties carry over from continuous Gaussians to discrete Gaussians. E.g.:

$$\forall \sigma \geq 1 : \rho_\sigma(L \setminus \mathcal{B}(\mathbf{0}, 2\sigma\sqrt{n})) \leq 2^{-n-1} \cdot \rho_\sigma(L).$$

  (i.e., the probability of getting a large vector is tiny)
- Discrete Gaussians can be sampled from efficiently.

The smoothing parameter

- Define $\eta(L)$ as the smallest σ such that $\rho_{1/\sigma}(\hat{L} \setminus \mathbf{0}) \leq 2^{-n}$.
- Intuition: If the standard deviation is larger than η, then discrete Gaussians behave like continuous ones.
- If $\sigma \geq \eta$, then $\rho_{\sigma,\mathbf{c}}(L)$ is quasi-constant:

$$\forall \mathbf{c} \in \mathbb{R}^n : \; \rho_{\sigma,\mathbf{c}}(L) \in \sigma^n \cdot (\mathrm{vol}\ \hat{L}) \cdot \left[1 \pm 2^{-n}\right].$$

- If $(\mathbf{b}_i)_i$ is a basis of L:

$$\eta(L) \leq \sqrt{n} \cdot \max \|\mathbf{b}_i^*\|.$$

- Consequence: $\eta \leq n \cdot \lambda_n$.

Proof that $\eta(L) \leq \sqrt{n} \cdot \max \|\mathbf{b}_i^*\|$

- First: $\eta(L) \leq \sqrt{n}/\lambda(\hat{L})$.

$$\begin{aligned} \rho_{1/\sigma}(\hat{L} \setminus \mathbf{0}) &= \rho(\sigma\hat{L} \setminus \mathcal{B}(\mathbf{0}, 2\sqrt{n})) \\ &\leq 2^{-n-1}\rho(\sigma\hat{L}) \\ &= 2^{-n-1}\rho_{1/\sigma}(\hat{L}) \\ &\leq 2^{-n}. \end{aligned}$$

- Second: $1/\lambda(\hat{L}) \leq \max \|\mathbf{b}_i^*\|$. With $B' = \text{reverse}(B^{-T})$:

$$\lambda(\hat{L}) \geq \min \|\mathbf{b}'^*_i\| \;\text{ and }\; \frac{1}{\min \|\mathbf{b}'^*_i\|} = \max \|\mathbf{b}_i^*\|.$$

Sampling from $D_{L,\sigma}$ (Gentry et al.'08)

There exists an efficient algorithm s.t. given as inputs a basis $(\mathbf{b}_i)_i$ of a lattice L, $\mathbf{c} \in \mathbb{Q}^n$ and $\sigma \geq \sqrt{n} \max \|\mathbf{b}_i^*\|$, produces vectors of L with distribution within statistical distance 2^{-n} of $D_{L,\sigma,\mathbf{c}}$:

$$\sum_{\mathbf{b} \in L} |\Pr[\mathbf{b}] - D_{L,\sigma,\mathbf{c}}(\mathbf{b})| \leq 2^{-n}.$$

- This may not exactly produce $D_{L,\sigma,\mathbf{c}}$, but no algorithm can see the difference with advantage $\geq \frac{1}{2} + 2^{-n}$.
- Being able to sample from $D_{L,\sigma,\mathbf{c}}$ with small σ is (almost) equivalent to having a small basis.
- But samples from $D_{L,\sigma,\mathbf{c}}$ do not provide information on the utilized small basis.

Plan

1- Background on Euclidean lattices.

2- **The SIS problem, or how to hash.**

3- The LWE problem, or how to encrypt.

4- Cryptanalysis.

5- Advanced topics: IBE and FHE.

The SIS problem

a- **Non structured SIS.**

b- Structured SIS.

c- A trapdoor for SIS.

$\mathrm{SIS}_{\beta,q,m}$ [Ajtai'96]

The Small Integer Solution Problem

Given a uniform $A \in \mathbb{Z}_q^{m \times n}$, find $\mathbf{s} \in \mathbb{Z}^m \setminus \mathbf{0}$ such that:

$$\|\mathbf{s}\| \leq \beta \quad \text{and} \quad \mathbf{s}A = \mathbf{0} \mod q.$$

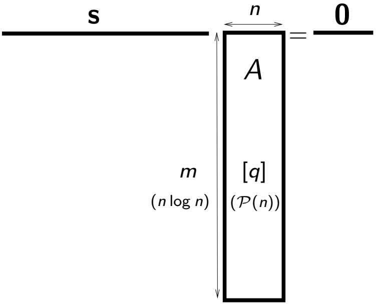

Many interpretations:

- Small codeword problem.
- Short lattice vector problem: $A^{\perp} = \{\mathbf{s} \in \mathbb{Z}^m : \mathbf{s}A = \mathbf{0}\ [q]\}$.

Cryptographic application of SIS

- Hash: an efficiently computable function $H : \mathcal{D} \mapsto \mathcal{R}$ with $|\mathcal{R}| \ll |\mathcal{D}|$ is collision resistant if finding $x \neq x'$ in $\mathcal{D}$ such that $H(x) = H(x')$ is computationally hard.
- Applications: message integrity, password verification, file identification, digital signature, etc.

- SIS-based hash: $\mathbf{s} \in \{0,1\}^m \mapsto \mathbf{s} \cdot A \ [q]$.
- By linearity, SIS reduces to finding a collision.
- Compression ratio: $\dfrac{m}{n \log q}$.

How hard is SIS? A unique level of security.

Worst-case to average-case reduction $(\gamma \approx n\beta)$

Any efficient SIS algorithm succeeding with non-negligible probability leads to an efficient SIVP algorithm.

Intuition:

- Start with a short basis of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points.
- Look at their coordinates wrt the basis, modulo q.
- A SIS solution provides a shorter vector of L.
- Repeat to get a basis shorter than the initial one.
- Repeat to get shorter and shorter bases of L.

The $D_{L,\sigma}$ sampler provides valid SIS inputs

- We start with a basis $(\mathbf{b}_i)$ with $B = \|\mathbf{b}_i\|$.
- Use the sampler with $\sigma = \sqrt{n}B$. Let $(\mathbf{c}_j)_{j \leq m}$ be the samples.
- With high probability: $\forall j, \|\mathbf{c}_j\| \leq \sqrt{n}\sigma = nB$.
- Are their coordinates wrt the $\mathbf{b}_i$'s uniform mod q?
- Yes, because $D_{L,\sigma}$ mod qL is (quasi)-uniform:
  $D_{qL,\sigma,\mathbf{c}}$ is (quasi)-independent of $\mathbf{c} \in L$,
  when $\sigma \geq \eta(qL) = q \cdot \eta(L)$.

Sufficient condition: $B \geq q\sqrt{n}\lambda_n$.

Shortness of the output vectors

- The $\mathbf{c}_j$'s satisfy $\|\mathbf{c}_j\| \leq nB$.
  Let $\mathbf{x}_j$ be their coordinates vectors, reduced mod q.
- The SIS oracle finds $\mathbf{s} \in \mathbb{Z}^m$ with $\sum s_j \mathbf{x}_j = \mathbf{0}\ [q]$ and $0 < \|\mathbf{s}\| \leq \beta$.
- Take $\mathbf{c} = \frac{1}{q} \sum s_j \mathbf{c}_j$: $\mathbf{c} \in L$ and $\|\mathbf{c}\| \leq \frac{\beta n^2 B}{q}$.
- If q is large enough, we obtain a shorter lattice vector.
- By analyzing $D_{L,\sigma}$ further, one can prove that by iterating, w.h.p. we can find a full rank set of short lattice vectors.
- We can convert the latter into a short basis.

Sufficient condition: $\dfrac{\beta n^2 B}{q} \leq \dfrac{B}{2}$.

The SIS problem

a- Non structured SIS.

b- Structured SIS.

c- A trapdoor for SIS.

Id-SIS, graphically

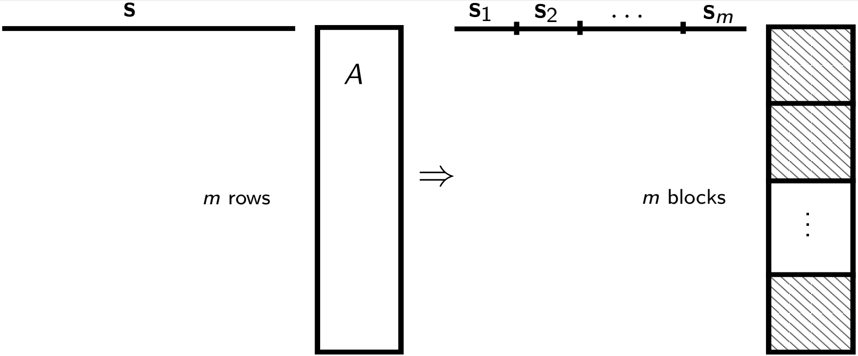

- Each block is negacyclic.
- The ith row is: $x^i \cdot a(x) \mod x^n + 1$.
- The structure allows us to decrease m by a factor n.
- Structured matrices $\equiv$ polynomials $\equiv$ fast algorithms.

Ideal SIS, algebraically

SIS

Given a uniform $A \in \mathbb{Z}_q^{m \times n}$, find $\mathbf{s} \in \mathbb{Z}^m \setminus \mathbf{0}$ such that:

$$\|\mathbf{s}\| \leq \beta \quad \text{and} \quad \mathbf{s}A = \mathbf{0} \mod q.$$

Let $R = \frac{\mathbb{Z}[x]}{x^n+1}$ and $R_q = \frac{\mathbb{Z}_q[x]}{x^n+1}$, with $n = 2^k$ and q prime.

Id-SIS

Given $a_1, \ldots, a_m \hookleftarrow U(R_q)$, find $s_1, \ldots, s_m \in R$ not all 0 s.t.:

$$\|\mathbf{s}\| \leq \beta \quad \text{and} \quad \sum s_i a_i = 0 \mod (q, x^n + 1).$$

Worst-case to average-case reduction

Any efficient **Id-SIS** algorithm succeeding with non-negligible probability leads to an efficient **Id-SIVP** algorithm.

Efficient hashing

- SIS hash: $\mathbf{s} \in \{0,1\}^m \mapsto \mathbf{s} \cdot A \; [q]$.
- Id-SIS hash: $s_1, \ldots, s_m \in \{0,1\}[x]$ of degrees $< n$ are mapped to $\sum s_i(x) a_i(x) \; [q, x^n + 1]$.
- If $2n | q - 1$, then $x^n + 1$ splits completely mod q.
  $\Rightarrow$ Fast Discrete Fourier Transform mod q.
- Storage: $\widetilde{O}(n^2) \to \widetilde{O}(n)$; complexity: $\widetilde{O}(n^2) \to \widetilde{O}(n)$.

This is SWIFFT and it was proposed to the SHA-3 contest.
With $n = 2^6, m = 2^4, q \approx 2^8$: A can be stored on $\approx 2^{13}$ bits.

The SIS problem

- **a-** Non structured SIS.
- **b-** Structured SIS.
- **c-** **A trapdoor for SIS.**

A uniform A with a good basis for $A^\perp$

If $m = \widetilde{\Omega}(n)$, we can efficiently sample $A \in \mathbb{Z}_q^{m \times n}$ and T_A s.t.

- The statistical distance from A to uniform is $2^{-\Omega(n)}$.
- The rows of T_A are small: $\max \|\mathbf{t}_i^*\| = \widetilde{O}(\sqrt{n})$.
- $T_A \in \mathbb{Z}^{m \times m}$ is a basis of $A^\perp$.

T_A (small)

$$T_A \cdot A = 0$$

Regularity principle:

- Assume $(\mathbf{a}_i)_{i \leq k}$ are iid uniform.
- Take $(x_i)_i$ iid uniform in $\{-1, 0, 1\}$.
- Then $\mathbf{a}_{k+1} = \sum_{i \leq k} x_i \mathbf{a}_i$ is close to uniform.

A trapdoor for (inhomogeneous) SIS

- Suppose we have $\mathbf{u} \in \mathbb{Z}_q^n$, A and T_A.
  How do we find a small $\mathbf{s} \in \mathbb{Z}^m$ such that $\mathbf{s}A = \mathbf{u}\ [q]$?
- With linear algebra, find $\mathbf{c} \in \mathbb{Z}^m$ such that $\mathbf{c}A = \mathbf{u}\ [q]$.
- It suffices to find a vector $\mathbf{b}$ of $A^{\perp}$ that is close to $\mathbf{c}$:
  $\|\mathbf{c} - \mathbf{b}\|$ is small and $(\mathbf{c} - \mathbf{b})A = \mathbf{u}\ [q]$.
- Use the sampler from $D_{L,\sigma,\mathbf{c}}$ with:

$$\sigma = \sqrt{n} \cdot \max \|\mathbf{t}_i^*\| = \widetilde{O}(n).$$

- We have $\|\mathbf{c} - \mathbf{b}\| \leq \sigma\sqrt{n} = \widetilde{O}(n^{1.5})$ w.h.p.
- And we do not leak any information about the trapdoor!

Cryptographic application: hash-and-sign

- Signature: to ensure the authenticity of a document.
- Signer's public key: A; private key: T_A.
- To sign M, use the trapdoor to find $\mathbf{s}$ short with $\mathbf{s}A = \mathcal{H}(M \| r)$, where $\mathcal{H}$ is a public random oracle.
- To verify $(M, \mathbf{s}, r)$, see whether $\mathbf{s}A = \mathcal{H}(M \| r)$ and $\|\mathbf{s}\|$ small.
- Can be made at least as hard to break as to solve SIS, in the random oracle model.

A trapdoor for Id-SIS

The trapdoor can be adapted to the structured case:

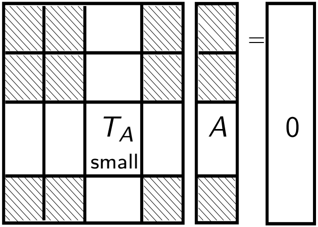

- Compact trapdoor: $\widetilde{O}(n^2) \rightarrow \widetilde{O}(n)$ bits.
- Verifying the signature boils down to FFT mod q.
- [Peikert'10] can be used to sign in quasi-linear time.