Introduction to lattice-based cryptography

Damien Stehlé

ENS de Lyon

Aussois, March 2019

Lattice-based cryptography

Maybe the most mature approach for post-quantum crypto.
Allows advanced cryptographic constructions
(homomorphic enc., some functional enc., some program obfuscation, etc)

Topics covered in this mini-course:

1. Hardness foundations: what are the assumptions?
2. Basic schemes: encrypting and signing
3. Fast(er) schemes using algebraic lattices

References:

- C. Peikert: a decade of lattice-based cryptography
  eprint 2015/939
- NewHope, Frodo, Kyber and Dilithium
  eprint 2015/1092, 2016/659, 2017/633 and 2017/634

Lattice-based cryptography

Maybe the most mature approach for post-quantum crypto.
Allows advanced cryptographic constructions
(homomorphic enc., some functional enc., some program obfuscation, etc)

Topics covered in this mini-course:

1. Hardness foundations: what are the assumptions?
2. Basic schemes: encrypting and signing
3. Fast(er) schemes using algebraic lattices

References:

- C. Peikert: a decade of lattice-based cryptography

  eprint 2015/939

- NewHope, Frodo, Kyber and Dilithium

  eprint 2015/1092, 2016/659, 2017/633 and 2017/634

Plan for this lecture

1. Background on Euclidean lattices.
2. The SIS and LWE problems.
3. Encrypting from LWE.

Euclidean lattices

Lattice $\equiv$ discrete subgroup of $\mathbb{R}^n$

$\equiv \quad \{\sum_{i \leq n} x_i \mathbf{b}_i : x_i \in \mathbb{Z}\}$

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:

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

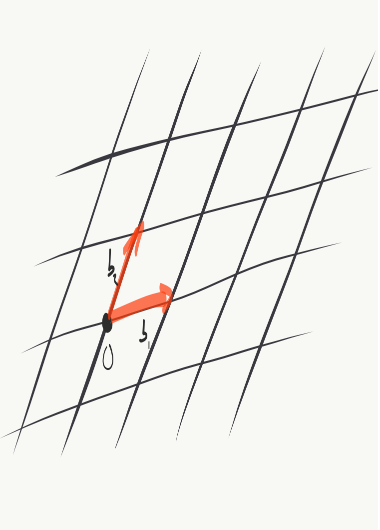

Euclidean lattices

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

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:

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

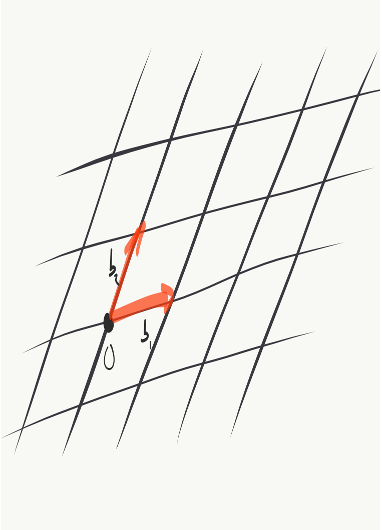

Euclidean lattices

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

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:

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

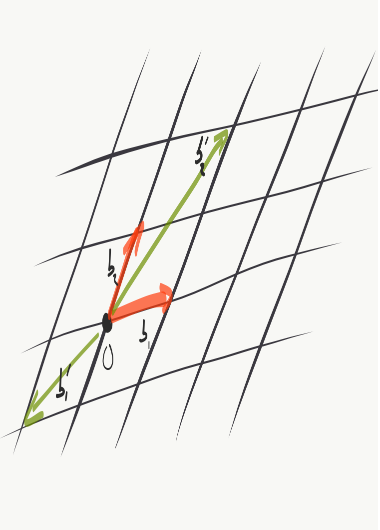

Lattice invariants

Dimension: n.

First minimum:
$\lambda_1 = \min(\|\mathbf{b}\| : \mathbf{b} \in L \setminus \mathbf{0})$.

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

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

Minkowski theorem:
$\lambda_1(L) \leq \sqrt{n} \cdot (\det\ L)^{1/n}$.

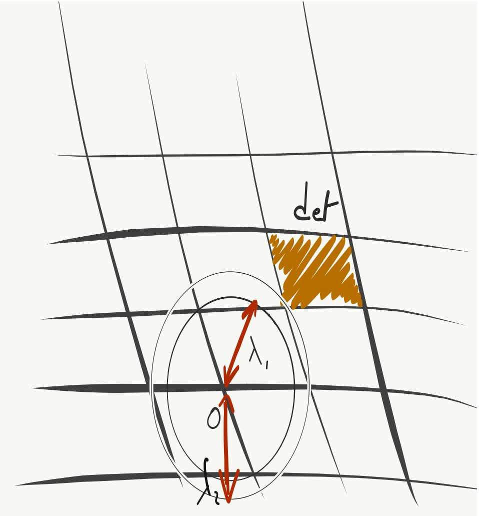

Lattice invariants

Dimension: n.

First minimum:
$\lambda_1 = \min(\|\mathbf{b}\| : \mathbf{b} \in L \setminus \mathbf{0})$.

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

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

Minkowski theorem:
$\lambda_1(L) \leq \sqrt{n} \cdot (\det L)^{1/n}$.

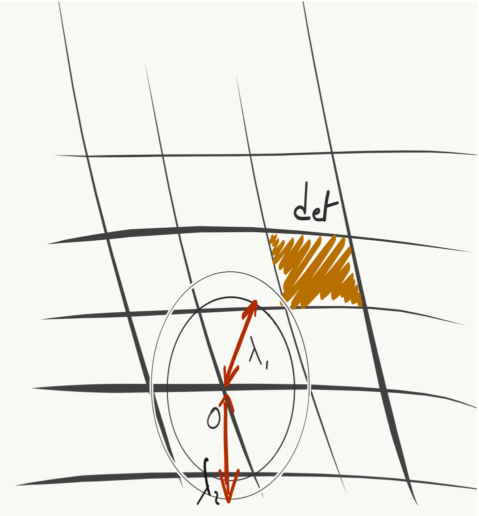

Lattice invariants

Dimension: n.

First minimum:
$\lambda_1 = \min(\|\mathbf{b}\| : \mathbf{b} \in L \setminus \mathbf{0})$.

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

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

Minkowski theorem:
$\lambda_1(L) \leq \sqrt{n} \cdot (\det\ L)^{1/n}$.

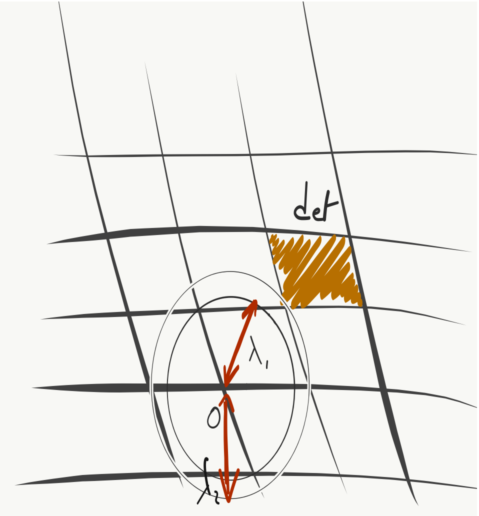

Lattice invariants

Dimension: n.

First minimum:
$\lambda_1 = \min(\|\mathbf{b}\| : \mathbf{b} \in L \setminus \mathbf{0})$.

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

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

Minkowski theorem:
$\lambda_1(L) \leq \sqrt{n} \cdot (\det L)^{1/n}$.

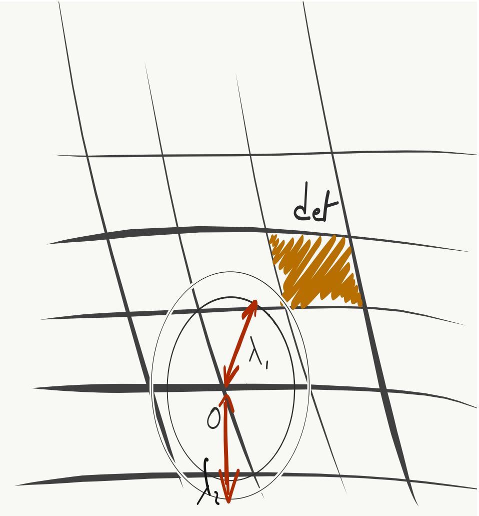

An example: construction A lattices

Construction A. Let $m \geq n \geq 1$ and $q \geq 2$ prime (for tranquility)

Let $A \in \mathbb{Z}_q^{m \times n}$. Then $L(A) := A \cdot \mathbb{Z}_q^n + q \cdot \mathbb{Z}^m$ is a lattice.

- Dimension: m
- Determinant, for full-rank A: q^{m-n}

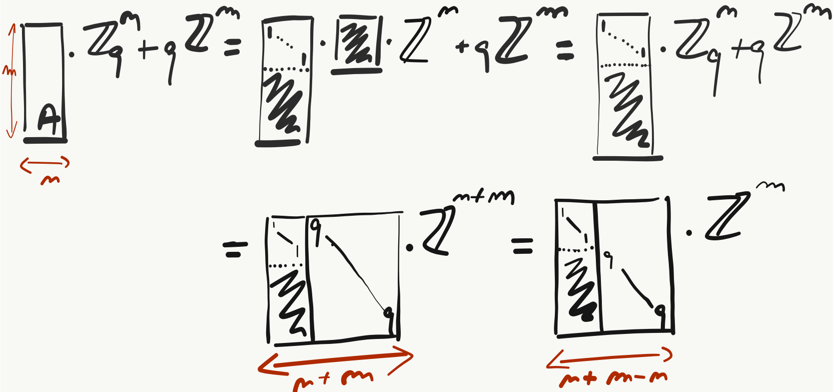

An example: construction A lattices

Construction A. Let $m \geq n \geq 1$ and $q \geq 2$ prime (for tranquility)

Let $A \in \mathbb{Z}_q^{m \times n}$. Then $L(A) := A \cdot \mathbb{Z}_q^n + q \cdot \mathbb{Z}^m$ is a lattice.

(For full-rank A.) Dim: m, det: q^{m-n} $\overset{\text{Minkowski}}{\Longrightarrow}$ $\lambda_1(L(A)) \leq \sqrt{m} \cdot q^{1-n/m}$.

$$\begin{aligned}
\Pr_A[\lambda_1 \leq B] &= \Pr_A\left[\exists \mathbf{s} \in \mathbb{Z}_q^n, \mathbf{b} \in \mathbb{Z}^m : 0 < \|\mathbf{b}\| < B \ \wedge \ \mathbf{b} = A \cdot \mathbf{s}\ [q]\right] \\
&\leq \sum_{\mathbf{s}} \sum_{\mathbf{b}} \max_{\mathbf{s},\mathbf{b}} \Pr_A[A \cdot \mathbf{s} = \mathbf{b}\ [q]] \\
&\lesssim q^n \cdot (B/\sqrt{m})^m \cdot \max_{\mathbf{s},\mathbf{b}} \Pr_A[A \cdot \mathbf{s} = \mathbf{b}\ [q]] \\
&\lesssim q^n \cdot (B/\sqrt{m})^m \cdot q^{-m}
\end{aligned}$$

(Third step requires $B \geq \sqrt{m}$, last step requires $B < q$)

Overall, if $q = \Omega(\sqrt{m})$, with probability ≈ 1 over a uniform A:

$$\lambda_1(L(A)) \geq \Omega\left(\min(q, \sqrt{m} \cdot q^{1-n/m})\right).$$

An example: construction A lattices

Construction A. Let $m \geq n \geq 1$ and $q \geq 2$ prime (for tranquility)

Let $A \in \mathbb{Z}_q^{m \times n}$. Then $L(A) := A \cdot \mathbb{Z}_q^n + q \cdot \mathbb{Z}^m$ is a lattice.

(For full-rank A.) Dim: m, det: q^{m-n} $\overset{\text{Minkowski}}{\Longrightarrow}$ $\lambda_1(L(A)) \leq \sqrt{m} \cdot q^{1-n/m}$.

$$
\begin{aligned}
\Pr_A[\lambda_1 \leq B] &= \Pr_A\left[\exists \mathbf{s} \in \mathbb{Z}_q^n, \mathbf{b} \in \mathbb{Z}^m : 0 < \|\mathbf{b}\| < B \ \wedge \ \mathbf{b} = A \cdot \mathbf{s} \ [q]\right] \\
&\leq \sum_{\mathbf{s}} \sum_{\mathbf{b}} \max_{\mathbf{s},\mathbf{b}} \Pr_A[A \cdot \mathbf{s} = \mathbf{b} \ [q]] \\
&\lesssim q^n \cdot (B/\sqrt{m})^m \cdot \max_{\mathbf{s},\mathbf{b}} \Pr_A[A \cdot \mathbf{s} = \mathbf{b} \ [q]] \\
&\lesssim q^n \cdot (B/\sqrt{m})^m \cdot q^{-m}
\end{aligned}
$$

(Third step requires $B \geq \sqrt{m}$, last step requires $B < q$)

Overall, if $q = \Omega(\sqrt{m})$, with probability ≈ 1 over a uniform A:

$$\lambda_1(L(A)) \geq \Omega\left(\min(q, \sqrt{m} \cdot q^{1-n/m})\right).$$

An example: construction A lattices

Construction A. Let $m \geq n \geq 1$ and $q \geq 2$ prime (for tranquility)

Let $A \in \mathbb{Z}_q^{m \times n}$. Then $L(A) := A \cdot \mathbb{Z}_q^n + q \cdot \mathbb{Z}^m$ is a lattice.

(For full-rank A.) Dim: m, det: q^{m-n} $\overset{\text{Minkowski}}{\Longrightarrow}$ $\lambda_1(L(A)) \leq \sqrt{m} \cdot q^{1-n/m}$.

$$
\begin{aligned}
\Pr_A[\lambda_1 \leq B] &= \Pr_A\left[\exists \mathbf{s} \in \mathbb{Z}_q^n, \mathbf{b} \in \mathbb{Z}^m : 0 < \|\mathbf{b}\| < B \ \wedge \ \mathbf{b} = A \cdot \mathbf{s}\ [q]\right] \\
&\leq \sum_{\mathbf{s}} \sum_{\mathbf{b}} \max_{\mathbf{s},\mathbf{b}} \Pr_A[A \cdot \mathbf{s} = \mathbf{b}\ [q]] \\
&\lesssim q^n \cdot (B/\sqrt{m})^m \cdot \max_{\mathbf{s},\mathbf{b}} \Pr_A[A \cdot \mathbf{s} = \mathbf{b}\ [q]] \\
&\lesssim q^n \cdot (B/\sqrt{m})^m \cdot q^{-m}
\end{aligned}
$$

(Third step requires $B \geq \sqrt{m}$, last step requires $B < q$)

Overall, if $q = \Omega(\sqrt{m})$, with probability ≈ 1 over a uniform A:

$$\lambda_1(L(A)) \geq \Omega\left(\min(q, \sqrt{m} \cdot q^{1-n/m})\right).$$

Another example

Let $m \geq n \geq 1$ and $q \geq 2$ prime.

Construction A for the orthogonal code

Let $A \in \mathbb{Z}_q^{m \times n}$. Then $A^{\perp} = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \cdot A = \mathbf{0}\ [q]\}$ is a lattice.

- Dimension: m
- Determinant: $q^{rk(A)}$.
- $\lambda_1 \approx \min(\sqrt{n \log q}, \sqrt{m} q^{n/m})$, with probability ≈ 1 for a uniform A.

Another example

Let $m \geq n \geq 1$ and $q \geq 2$ prime.

Construction A for the orthogonal code

Let $A \in \mathbb{Z}_q^{m \times n}$. Then $A^\perp = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \cdot A = \mathbf{0}\ [q]\}$ is a lattice.

- Dimension: m
- Determinant: $q^{rk(A)}$.
- $\lambda_1 \approx \min(\sqrt{n \log q}, \sqrt{m} q^{n/m})$, with probability ≈ 1 for a uniform A.

SVP and SIVP

The Shortest Vector Problem: SVP$_\gamma$

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

SVP and SIVP

The Shortest Vector Problem: SVP$_\gamma$

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$_\gamma$

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

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

SVP and SIVP

The Shortest Vector Problem: SVP$_\gamma$

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$_\gamma$

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 reductions for SVP).
- In lattice-based crypto: $\gamma = \mathcal{P}oly(n)$ (most often).
- Solvable in polynomial time when $\gamma = 2^{\widetilde{O}(n)}$.

CVP and BDD

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:

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

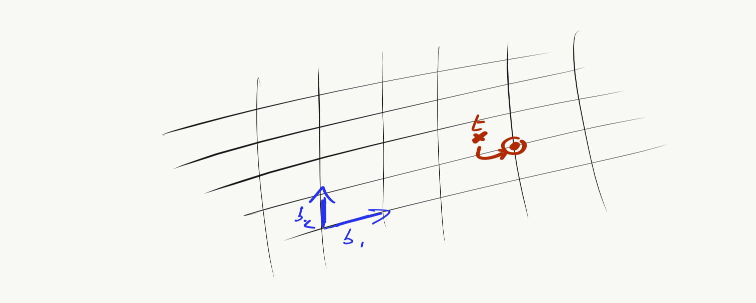

BDD_γ (Bounded Distance Decoding)

Find the closest $\mathbf{b} \in L$ to $\mathbf{t}$, under the promise that $\|\mathbf{b} - \mathbf{t}\| \leq \lambda_1(L)/\gamma$.

Hardness

- All known algorithms for SVP, SIVP, CVP, BDD with $\gamma = \mathcal{P}oly(n)$ cost $2^{\Omega(n)}$.
- Same landscape if we allow quantum algorithms.

Open problems

- How equivalent are these problems? See survey by Noah Stephens-Davidowitz
- Can we beat the $2^{\Omega(n)}$ cost barrier?

But these are worst-case problems, and worst-case hardness is not convenient for cryptographic purposes.

Hardness

- All known algorithms for SVP, SIVP, CVP, BDD with $\gamma = \mathcal{P}oly(n)$ cost $2^{\Omega(n)}$.
- Same landscape if we allow quantum algorithms.

Open problems

- How equivalent are these problems? See survey by Noah Stephens-Davidowitz
- Can we beat the $2^{\Omega(n)}$ cost barrier?

But these are worst-case problems, and worst-case hardness is not convenient for cryptographic purposes.

Plan for this lecture

1. Background on Euclidean lattices.
2. **The SIS and LWE problems.**
3. Encrypting from LWE.

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

The Short Integer Solution Problem

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

$$\|\mathbf{x}\| \leq \beta \quad \text{and} \quad \mathbf{x}^T \cdot A = \mathbf{0} \bmod q.$$

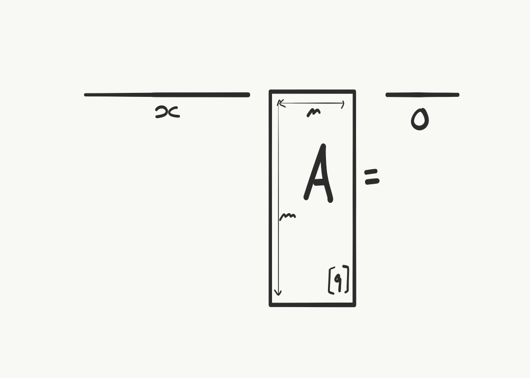

SIS as a lattice problem

Remember our lattice example:

$$A^{\perp} = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \cdot A = \mathbf{0} \ [q]\}.$$

SIS consists in finding a short non-zero vector in $A^{\perp}$, for a random A.

- If $\beta < \lambda_1 \approx \min(\sqrt{n \log q}, \sqrt{m} q^{n/m})$: trivially hard.
- If $\beta \geq q$: trivially easy.
- In between: interesting.

SIS is an average-case SVP/SIVP.

SIS as a lattice problem

Remember our lattice example:

$$A^{\perp} = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \cdot A = \mathbf{0} \ [q]\}.$$

SIS consists in finding a short non-zero vector in $A^{\perp}$, for a random A.

- If $\beta < \lambda_1 \approx \min(\sqrt{n \log q}, \sqrt{m} q^{n/m})$: trivially hard.
- If $\beta \geq q$: trivially easy.
- In between: interesting.

SIS is an average-case SVP/SIVP.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction $(\gamma \approx n\beta,\ q \geq \sqrt{n}\beta)$

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction ($\gamma \approx n\beta$, $q \geq \sqrt{n}\beta$)

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q} \sum x_i \mathbf{y}_i$ is a shorter vector in L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction ($\gamma \approx n\beta$, $q \geq \sqrt{n}\beta$)

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}$ $[q]$.
- $\frac{1}{q} \sum x_i \mathbf{y}_i$ is a shorter vector in L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction $(\gamma \approx n\beta,\ q \geq \sqrt{n}\beta)$

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q} \sum x_i \mathbf{y}_i$ is a shorter vector in L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction $(\gamma \approx n\beta,\ q \geq \sqrt{n}\beta)$

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q}\sum x_i \mathbf{y}_i$ is a shorter vector in L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction $(\gamma \approx n\beta,\ q \geq \sqrt{n}\beta)$

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.
(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q}\sum x_i \mathbf{y}_i$ is a shorter vector in L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction $(\gamma \approx n\beta,\ q \geq \sqrt{n}\beta)$

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q} \sum x_i \mathbf{y}_i$ is a shorter vector in L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction ($\gamma \approx n\beta$, $q \geq \sqrt{n}\beta$)

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.

(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q} \sum x_i \mathbf{y}_i$ is a shorter vector in L.

- Repeat to get a basis shorter than the initial one.
- Repeat to get shorter and shorter bases of L.

Hardness of SIS? [Ajtai96,...,GPV08]

Worst-case to average-case reduction ($\gamma \approx n\beta$, $q \geq \sqrt{n}\beta$)

Any efficient $\mathrm{SIS}_{\beta,q,m}$ algorithm succeeding with non-negligible probability leads to an efficient SIVP_γ algorithm.
(See [MP13] for smaller q)

Sketch:

- Start with a short basis B of the lattice $L \subseteq \mathbb{Z}^n$.
- Sample m short random lattice points $(\mathbf{y}_i)_{i \leq m}$.
- Look at their coordinates with respect to B, reduced modulo q. These form a SIS instance A.
- The SIS oracle gives $\mathbf{x} \in \mathbb{Z}^m$ short s.t. $\mathbf{x}^T \cdot A = \mathbf{0}\ [q]$.
- $\frac{1}{q} \sum x_i \mathbf{y}_i$ is a shorter vector in L.
- Repeat to get a basis shorter than the initial one.
- Repeat to get shorter and shorter bases of L.

$\mathrm{LWE}_{\alpha,q}$ [Regev'05]

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Let $D_{\mathbf{s},\alpha}$ be the distribution corresponding to:

$$(\mathbf{a}; \langle \mathbf{a}, \mathbf{s} \rangle + e\ [q]) \quad \text{with} \quad \mathbf{a} \hookleftarrow U(\mathbb{Z}_q^n),\ e \hookleftarrow \lfloor \nu_{\alpha q} \rceil,$$

where $\nu_{\alpha q}$ denotes the continuous Gaussian of st. dev. αq.

The Learning With Errors Problem — Search-LWE_α

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Given arbitrarily many samples from $D_{\mathbf{s},\alpha}$, find $\mathbf{s}$.

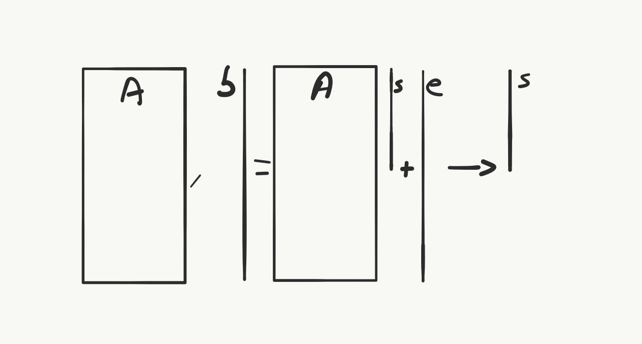

LWE as a lattice problem

Search-LWE$_\alpha$

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Given $(A; A\mathbf{s} + \mathbf{e}\ [q])$ with $A \hookleftarrow U(\mathbb{Z}_q^{m\times n})$ and $\mathbf{e} \hookleftarrow \lfloor \nu_{\alpha q}^m \rceil$ for and arbitrary m, find $\mathbf{s}$.

Remember our lattice example $L_A = A \cdot \mathbb{Z}_q^n + q \cdot \mathbb{Z}^m$.

- If $\alpha \approx 0$, then LWE is easy to solve.
- If $\alpha \gg 1$, then LWE is trivially hard.
- In between: interesting.

LWE is an average-case BDD.

LWE as a lattice problem

Search-LWE$_\alpha$

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Given $(A; A\mathbf{s} + \mathbf{e}\ [q])$ with $A \hookleftarrow U(\mathbb{Z}_q^{m \times n})$ and $\mathbf{e} \hookleftarrow \lfloor \nu_{\alpha q}^m \rceil$ for and arbitrary m, find $\mathbf{s}$.

Remember our lattice example $L_A = A \cdot \mathbb{Z}_q^n + q \cdot \mathbb{Z}^m$.

- If $\alpha \approx 0$, then LWE is easy to solve.
- If $\alpha \gg 1$, then LWE is trivially hard.
- In between: interesting.

LWE is an average-case BDD.

How hard is LWE? [Regev05]

Quantum worst-case to average-case reduction $(\gamma \approx n/\alpha,\ \alpha q \geq \sqrt{n})$

Assume that q is prime and $\mathcal{P}oly(n)$.
Any efficient $\mathrm{LWE}_{n,\alpha,q}$ algorithm succeeding with non-negligible probability leads to an efficient **quantum** SIVP_γ algorithm.

How hard is LWE? [Regev05]

Quantum worst-case to average-case reduction $(\gamma \approx n/\alpha,\ \alpha q \geq \sqrt{n})$

Assume that q is prime and $\mathcal{P}oly(n)$.
Any efficient $\mathrm{LWE}_{n,\alpha,q}$ algorithm succeeding with non-negligible probability leads to an efficient **quantum** SIVP_γ algorithm.

How hard is LWE? [Regev05]

Quantum worst-case to average-case reduction $(\gamma \approx n/\alpha,\ \alpha q \geq \sqrt{n})$

Assume that q is prime and $\mathcal{P}oly(n)$.
Any efficient $\mathrm{LWE}_{n,\alpha,q}$ algorithm succeeding with non-negligible probability leads to an efficient **quantum** SIVP_γ algorithm.

- [Peikert09]: classical reduction, for $q \approx 2^n$, from BDD.
- [SSTX09]: simpler (but weaker) quantum reduction, from SIS.
- [BLPRS13]: de-quantized reduction, for any q that is at least some $\mathcal{P}oly(n)$, from a weaker worst-case lattice problem.
- [BKSW18]: yet another quantum reduction, from BDD.

Decision LWE

$$D_{\mathbf{s},\alpha}: \quad (\mathbf{a}; \langle \mathbf{a}, \mathbf{s} \rangle + e\ [q]) \quad \text{with} \quad \mathbf{a} \hookleftarrow U(\mathbb{Z}_q^n),\ e \hookleftarrow \lfloor \nu_{\alpha q} \rceil.$$

Search-LWE$_\alpha$

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Given arbitrarily many samples from $D_{\mathbf{s},\alpha}$, find $\mathbf{s}$.

Dec-LWE$_\alpha$

Let $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$. With non-negligible probability over $\mathbf{s}$, distinguish between an oracle access to $D_{\mathbf{s},\alpha}$ or an oracle access to $U(\mathbb{Z}_q^{n+1})$.

Decision LWE

$$D_{\mathbf{s},\alpha}: \quad (\mathbf{a}; \langle \mathbf{a}, \mathbf{s} \rangle + e \ [q]) \quad \text{with} \quad \mathbf{a} \hookleftarrow U(\mathbb{Z}_q^n), \ e \hookleftarrow \lfloor \nu_{\alpha q} \rceil.$$

Search-LWE$_\alpha$

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Given arbitrarily many samples from $D_{\mathbf{s},\alpha}$, find $\mathbf{s}$.

Dec-LWE$_\alpha$

Let $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$. With non-negligible probability over $\mathbf{s}$, distinguish between an oracle access to $D_{\mathbf{s},\alpha}$ or an oracle access to $U(\mathbb{Z}_q^{n+1})$.

Decision LWE

$$D_{\mathbf{s},\alpha}: \quad (\mathbf{a}; \langle \mathbf{a}, \mathbf{s}\rangle + e\ [q]) \quad \text{with} \quad \mathbf{a} \hookleftarrow U(\mathbb{Z}_q^n),\ e \hookleftarrow \lfloor \nu_{\alpha q} \rceil.$$

Search-LWE$_\alpha$

Let $\mathbf{s} \in \mathbb{Z}_q^n$. Given arbitrarily many samples from $D_{\mathbf{s},\alpha}$, find $\mathbf{s}$.

Dec-LWE$_\alpha$

Let $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$. With non-negligible probability over $\mathbf{s}$, distinguish between an oracle access to $D_{\mathbf{s},\alpha}$ or an oracle access to $U(\mathbb{Z}_q^{n+1})$.

Dec-LWE and Search-LWE efficiently reduce to one another.

Decision LWE and SIS

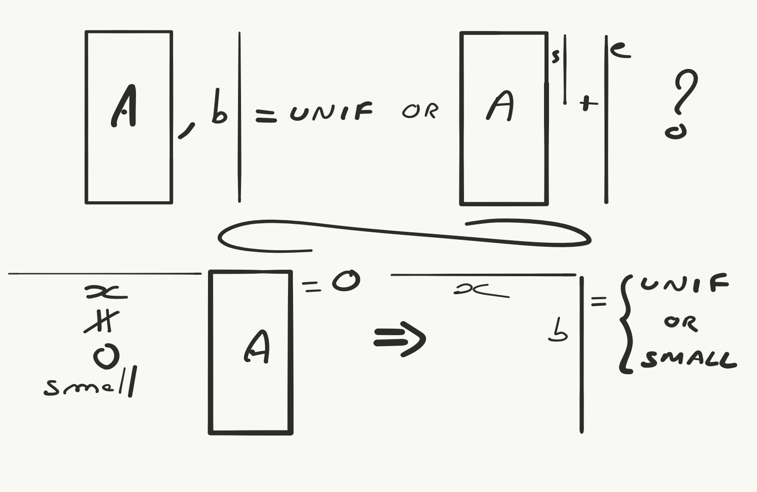

Nice properties of LWE

1. Arbitrary number of samples
   $\Rightarrow$ can amplify success probability and distinguishing advantage.
2. Random self-reducibility
   $\Rightarrow$ solving for a non-negligible fraction of $\mathbf{s}$'s suffices.

$$(A, A \cdot \mathbf{s} + \mathbf{e}) + (0, A \cdot \mathbf{t}) = (A, A \cdot (\mathbf{s} + \mathbf{t}) + \mathbf{e})$$

3. A distinguishing oracle allows to check a guess for a coordinate of $\mathbf{s}$.

$\Rightarrow$ These lead to a search-to-decision reduction.

4. Can take different types of noises:
   - Discrete Gaussian
   - Uniform integer in an interval
   - Deterministic, using rounding

Nice properties of LWE

1. Arbitrary number of samples
   $\Rightarrow$ can amplify success probability and distinguishing advantage.
2. Random self-reducibility
   $\Rightarrow$ solving for a non-negligible fraction of $\mathbf{s}$'s suffices.

$$(A, A \cdot \mathbf{s} + \mathbf{e}) + (0, A \cdot \mathbf{t}) = (A, A \cdot (\mathbf{s} + \mathbf{t}) + \mathbf{e})$$

3. A distinguishing oracle allows to check a guess for a coordinate of $\mathbf{s}$.

$\Rightarrow$ These lead to a search-to-decision reduction.

4. Can take different types of noises:
   - Discrete Gaussian
   - Uniform integer in an interval
   - Deterministic, using rounding

Nice properties of LWE

1. Arbitrary number of samples
   $\Rightarrow$ can amplify success probability and distinguishing advantage.
2. Random self-reducibility
   $\Rightarrow$ solving for a non-negligible fraction of $\mathbf{s}$'s suffices.

$$(A, A \cdot \mathbf{s} + \mathbf{e}) + (0, A \cdot \mathbf{t}) = (A, A \cdot (\mathbf{s} + \mathbf{t}) + \mathbf{e})$$

3. A distinguishing oracle allows to check a guess for a coordinate of $\mathbf{s}$.

$\Rightarrow$ These lead to a search-to-decision reduction.

4. Can take different types of noises:
   - Discrete Gaussian
   - Uniform integer in an interval
   - Deterministic, using rounding

Nice properties of LWE

1. Arbitrary number of samples
   $\Rightarrow$ can amplify success probability and distinguishing advantage.
2. Random self-reducibility
   $\Rightarrow$ solving for a non-negligible fraction of $\mathbf{s}$'s suffices.

$$(A, A \cdot \mathbf{s} + \mathbf{e}) + (0, A \cdot \mathbf{t}) = (A, A \cdot (\mathbf{s} + \mathbf{t}) + \mathbf{e})$$

3. A distinguishing oracle allows to check a guess for a coordinate of $\mathbf{s}$.

$\Rightarrow$ These lead to a search-to-decision reduction.

4. Can take different types of noises:
   - Discrete Gaussian
   - Uniform integer in an interval
   - Deterministic, using rounding

Open problems

Selected problems on SIS/LWE

- Can we get hardness of SIS/LWE based on SIVP with approximation factor less than n?
- Can we reduce SVP_γ to SIS/LWE?
- Can we get a classical reduction from SIVP to LWE with parameters equivalent to those of Regev's quantum reduction?
- Or is this discrepancy intrinsic and there is a quantum acceleration for solving LWE and SIVP?

Open problems

Selected problems on SIS/LWE

- Can we get hardness of SIS/LWE based on SIVP with approximation factor less than n?
- Can we reduce SVP_γ to SIS/LWE?
- Can we get a classical reduction from SIVP to LWE with parameters equivalent to those of Regev's quantum reduction?
- Or is this discrepancy intrinsic and there is a quantum acceleration for solving LWE and SIVP?

Open problems

Selected problems on SIS/LWE

- Can we get hardness of SIS/LWE based on SIVP with approximation factor less than n?
- Can we reduce SVP_γ to SIS/LWE?
- Can we get a classical reduction from SIVP to LWE with parameters equivalent to those of Regev's quantum reduction?
- Or is this discrepancy intrinsic and there is a quantum acceleration for solving LWE and SIVP?

Plan for this lecture

1. Background on Euclidean lattices.
2. The SIS and LWE problems.
3. **Encrypting from LWE.**

SVP/SIVP/CVP/BDD are here only implicitly:
(almost) no need to know lattices for designing lattice-based schemes!

LWE with small secret

Small-secret-LWE$_\alpha$

Let $\mathbf{s} \hookleftarrow \lfloor \nu_{\alpha q} \rceil^n$. With non-negligible probability over $\mathbf{s}$, distinguish between (arbitrarily many) samples from $D_{\mathbf{s},\alpha}$ or from $U(\mathbb{Z}_q^{n+1})$.

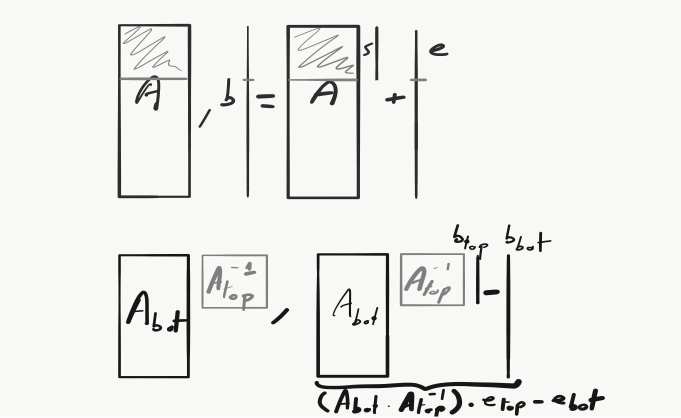

LWE-based encryption

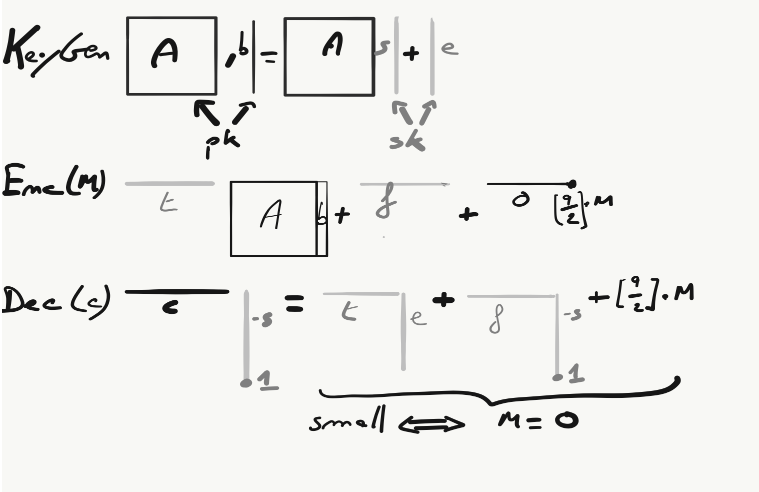

Decryption correctness

To ensure correctness, it suffices that

$$\left|\mathbf{t}^T\mathbf{e} + \mathbf{f}^T(-\mathbf{s}|1)\right| \; < \; q/4,$$

with probability very close to 1.

Up to the roundings of Gaussians:

- Gaussian tail bound $\Rightarrow$ $\|\mathbf{t}\|, \|\mathbf{e}\|, \|\mathbf{f}\|, \|\mathbf{s}\| \lesssim \sqrt{n}\alpha q$ with probability $1 - 2^{-\Omega(n)}$.
- It suffices that $(\sqrt{n}\alpha q)^2 \lesssim q/4$, i.e., $\alpha \lesssim 1/(n\sqrt{q})$.

Better:

- $\mathbf{t}^T\mathbf{e}$ is a 1-dim Gaussian of parameter $\alpha q\|\mathbf{e}\|$.
- Gaussian tail bound $\Rightarrow$ $|\mathbf{t}^T\mathbf{e}| \leq \sqrt{n \log n} \cdot (\alpha q)^2$ with probability $\geq 1 - 1/\mathcal{P}oly(n)$.
- It suffices that $\alpha \lesssim 1/\sqrt{qn \log n}$.

Decryption correctness

To ensure correctness, it suffices that

$$\left|\mathbf{t}^T \mathbf{e} + \mathbf{f}^T(-\mathbf{s}|1)\right| \;<\; q/4,$$

with probability very close to 1.

Up to the roundings of Gaussians:

- Gaussian tail bound $\Rightarrow$ $\|\mathbf{t}\|, \|\mathbf{e}\|, \|\mathbf{f}\|, \|\mathbf{s}\| \lesssim \sqrt{n}\alpha q$ with probability $1 - 2^{-\Omega(n)}$.
- It suffices that $(\sqrt{n}\alpha q)^2 \lesssim q/4$, i.e., $\alpha \lesssim 1/(n\sqrt{q})$.

Better:

- $\mathbf{t}^T\mathbf{e}$ is a 1-dim Gaussian of parameter $\alpha q \|\mathbf{e}\|$.
- Gaussian tail bound $\Rightarrow$ $|\mathbf{t}^T\mathbf{e}| \leq \sqrt{n \log n} \cdot (\alpha q)^2$ with probability $\geq 1 - 1/\mathcal{P}oly(n)$.
- It suffices that $\alpha \lesssim 1/\sqrt{qn \log n}$.

Decryption correctness

To ensure correctness, it suffices that

$$\left|\mathbf{t}^T\mathbf{e} + \mathbf{f}^T(-\mathbf{s}|1)\right| \ < \ q/4,$$

with probability very close to 1.

Up to the roundings of Gaussians:

- Gaussian tail bound $\Rightarrow$ $\|\mathbf{t}\|, \|\mathbf{e}\|, \|\mathbf{f}\|, \|\mathbf{s}\| \lesssim \sqrt{n}\alpha q$ with probability $1 - 2^{-\Omega(n)}$.
- It suffices that $(\sqrt{n}\alpha q)^2 \lesssim q/4$, i.e., $\alpha \lesssim 1/(n\sqrt{q})$.

Better:

- $\mathbf{t}^T\mathbf{e}$ is a 1-dim Gaussian of parameter $\alpha q\|\mathbf{e}\|$.
- Gaussian tail bound $\Rightarrow$ $|\mathbf{t}^T\mathbf{e}| \leq \sqrt{n \log n} \cdot (\alpha q)^2$ with probability $\geq 1 - 1/\mathcal{P}oly(n)$.
- It suffices that $\alpha \lesssim 1/\sqrt{qn \log n}$.

Do decryption errors matter?

- We can cut Gaussian tails and use the first error bound to guarantee perfect correctness.
- The probability is quite close to 1, so it does not matter much.
- We can use an error correcting code to boost the correct decryption probability.

For security against chosen ciphertext attacks, it does matter ([HHK17,AGJNVV19], 2017/604, 2018/1089, 2019/043). $\Rightarrow$ tune parameters to make it very small.

Do decryption errors matter?

- We can cut Gaussian tails and use the first error bound to guarantee perfect correctness.
- The probability is quite close to 1, so it does not matter much.
- We can use an error correcting code to boost the correct decryption probability.

For security against chosen ciphertext attacks, it does matter ([HHK17,AGJNVV19], 2017/604, 2018/1089, 2019/043). $\Rightarrow$ tune parameters to make it very small.

Passive security (IND-CPA)

KeyGen A, $b = \cancel{A s + e}$ | b UNIFORM

pk sk

Enc(M) t $\begin{bmatrix} A & b \end{bmatrix}$ + f + $0 \ \lceil \frac{q}{2} \rfloor \cdot M$

Passive security (IND-CPA)

KeyGen $\boxed{A}$, b UNIFORM

pk

Enc(M) $\underbrace{t \; \boxed{A} \; b + f}_{\text{UNIFORM}} + \underbrace{0 \quad \lfloor \frac{q}{2} \rceil \cdot M}_{\text{statistically hidden}}$

Setting parameters (asymptotically)

How do we choose n, α and q?

Minimize bandwidth/key-size/run-times under the conditions that:

- Correctness holds
- Some security is guaranteed

Take $\sqrt{n}/q \approx 1/\sqrt{qn \log n}$, i.e., $q \approx n^2 \log n$.

Take $\alpha \approx \sqrt{n}/q \approx 1/(n^{3/2} \log n)$.

SIVP_γ in dimension n quantumly reduces to $\mathrm{LWE}_{n,\alpha,q}$ for $\gamma \approx n/\alpha$.

Setting parameters (asymptotically)

How do we choose n, α and q?

Minimize bandwidth/key-size/run-times under the conditions that:

- Correctness holds
- Some security is guaranteed

Take $\sqrt{n}/q \approx 1/\sqrt{qn \log n}$, i.e., $q \approx n^2 \log n$.

Take $\alpha \approx \sqrt{n}/q \approx 1/(n^{3/2} \log n)$.

SIVP_γ in dimension n quantumly reduces to $\mathrm{LWE}_{n,\alpha,q}$ for $\gamma \approx n/\alpha$.

Setting parameters (asymptotically)

How do we choose n, α and q?

Minimize bandwidth/key-size/run-times under the conditions that:

- Correctness holds $\alpha \lesssim 1/\sqrt{qn \log n}$
- Some security is guaranteed $\alpha q \geq \sqrt{n}$

Take $\sqrt{n}/q \approx 1/\sqrt{qn \log n}$, i.e., $q \approx n^2 \log n$.

Take $\alpha \approx \sqrt{n}/q \approx 1/(n^{3/2} \log n)$.

SIVP_γ in dimension n quantumly reduces to $\mathrm{LWE}_{n,\alpha,q}$ for $\gamma \approx n/\alpha$.

Setting parameters (asymptotically)

How do we choose n, α and q?

Minimize bandwidth/key-size/run-times under the conditions that:

- Correctness holds $\alpha \lesssim 1/\sqrt{qn \log n}$
- Some security is guaranteed $\alpha q \geq \sqrt{n}$

Take $\sqrt{n}/q \approx 1/\sqrt{qn \log n}$, i.e., $q \approx n^2 \log n$.

Take $\alpha \approx \sqrt{n}/q \approx 1/(n^{3/2} \log n)$.

SIVP_γ in dimension n quantumly reduces to $\mathrm{LWE}_{n,\alpha,q}$ for $\gamma \approx n/\alpha$.

From passive to active security

IND-CCA security

The encryptions of two plaintexts chosen by the adversary $\mathcal{A}$ should remain indistinguishable in $\mathcal{A}$'s view, **even if $\mathcal{A}$ can request decryptions of ciphertexts of its choice** (except the challenge ciphertexts).

How do we upgrade IND-CPA security to IND-CCA security?

- **OAEP**: requires decryption to recover the encryption randomness
  This is not our case: we recover $\mathbf{t}^T\mathbf{e} + \mathbf{f}^T(-\mathbf{s}|1)$.
- **Fujisaki-Okamoto**: upgraded decryption uses initial encryption and decryption algorithms.
- FO is secure in the **random oracle model**, if decryption errors occur with exponentially small probability.
- FO is also secure in the **quantum random oracle model**, but with a big advantage loss.

D. Hofheinz, K. Hövelmanns, E. Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. Eprint 2017/604.

Open problems

Selected problems on LWE encryption

- Do the diverse noise distributions have an impact?
- What is the exact impact of decryption failures to CCA security of the FO upgrade?
- Can we get efficient CCA security without the random oracle heuristic?

Open problems

Selected problems on LWE encryption

- Do the diverse noise distributions have an impact?
- What is the exact impact of decryption failures to CCA security of the FO upgrade?
- Can we get efficient CCA security without the random oracle heuristic?

Open problems

Selected problems on LWE encryption

- Do the diverse noise distributions have an impact?
- What is the exact impact of decryption failures to CCA security of the FO upgrade?
- Can we get efficient CCA security without the random oracle heuristic?

Plan for this lecture

1. Background on Euclidean lattices.
2. The SIS and LWE problems.
3. Encrypting from LWE.

Wrapping up

Lattices ae conjectured to provide exponentially hard worst-case problems, even for quantum algorithms.

SIS and LWE are average-case variants that are proved to be no easier than some such hard lattice problems.

- There is no fundamental weakness in SIS/LWE, compared to worst-case lattices.
- The reductions are not meant for setting parameters, but for ensuring that there is no fundamental weakness.
- Average-case problems are better suited for cryptographic design.

SIS and LWE are linear algebra problems.

- Leads to simple cryptographic design.
- Allows advanced cryptographic constructions.

Next time

- Signing from SIS
- Efficient variants of SIS/LWE
- NTRU