Short Integer Solutions
A Worst-case to Average-case Reduction

Agnese Gini

University of Luxembourg
Interdisciplinary Centre for Security, Reliability and Trust

June 14, 2019

Short Integer Solutions ($\mathsf{SIS}_{q,m,\beta}$)

Given a positive integer q, a matrix $A \in \mathbb{Z}_q^{n \times m}$ and $\beta \in \mathbb{R}$, find $e \in \mathbb{Z}^m \setminus \{0\}$ such that

$$Ae = 0 \mod q \quad \text{and} \quad \|e\|_2 \leq \beta$$

Short Integer Solutions ($\mathsf{SIS}_{q,m,\beta}$)

Given a positive integer q, a matrix $A \in \mathbb{Z}_q^{n \times m}$ and $\beta \in \mathbb{R}$, find $e \in \mathbb{Z}^m \setminus \{0\}$ such that

$$Ae = 0 \mod q \quad \text{and} \quad \|e\|_2 \leq \beta$$

For functions $q(n), m(n)$ and $\beta(n)$, $\mathsf{SIS}_{q,m,\beta}$ is the probability ensemble over instances $(q(n), A, \beta(n))$ with A chosen uniformly at random among $\mathbb{Z}_q^{n \times m(n)}$.

Overview

Solutions for SIS

A Worst-case to Average-case Reduction: SIVP to SIS

- Smoothing parameter
- IncIVD to SIS
- SIVP to IncIVD

Solutions for SIS

Proposition

For any $A \in \mathbb{Z}_q^{n \times m}$ and $\beta \geq \sqrt{m} q^{n/m}$, then the instance (q, A, β) of SIS admits a solution.

Proof:
The set of all vectors $\{0, \ldots, q^{n/m}\}^m$ has cardinality greater than q^n. Hence there exist $z_1 \neq z_2$ such that $Az_1 = Az_2 \mod q$ and $z = z_1 - z_2 \neq 0$. Then $Az = 0 \mod q$ and $\|z\| \leq \sqrt{m} q^{n/m}$, since it has coordinates in $[-q^{n/m}, q^{n/m}] \cap \mathbb{Z}$.

Recall:

$$\Lambda_q(A) \coloneqq \{y \in \mathbb{Z}^m : y = Ax \mod q \text{ for some } x \in \mathbb{Z}^n\}$$

$$\Lambda_q^{\perp}(A) \coloneqq \{e \in \mathbb{Z}^m : Ae = 0 \mod q\}$$

Recall:

$$\Lambda_q(A) \coloneqq \{y \in \mathbb{Z}^m : y = Ax \mod q \text{ for some } x \in \mathbb{Z}^n\}$$

$$\Lambda_q^{\perp}(A) \coloneqq \{e \in \mathbb{Z}^m : Ae = 0 \mod q\}$$

$e \in \mathbb{Z}^m$ is a solution if and only if $e \in \Lambda_q^{\perp}(A)$ and $\|e\| \leq \beta$

Recall:

$$\Lambda_q(A) := \{y \in \mathbb{Z}^m : y = Ax \mod q \text{ for some } x \in \mathbb{Z}^n\}$$

$$\Lambda_q^\perp(A) := \{e \in \mathbb{Z}^m : Ae = 0 \mod q\}$$

$e \in \mathbb{Z}^m$ is a solution if and only if $e \in \Lambda_q^\perp(A)$ and $\|e\| \leq \beta$

Proposition

- $\Lambda_q^\perp(A) = q(\Lambda_q(A)^\vee)$ and $\Lambda_q(A) = q((\Lambda_q^\perp(A))^\vee)$,
- $\det(\Lambda_q^\perp(A)) \leq q^n$,
- $\det(\Lambda_q(A)) \geq q^{m-n}$,
- If q is prime, the above equalities hold if and only if A is full-rank.

Corollary

Let A be in $\mathbb{Z}_q^{n \times m}$. Then $\lambda_1(\Lambda_q^{\perp}(A)) \leq \sqrt{m} q^{m-n}$.

Corollary

Let A be in $\mathbb{Z}_q^{n \times m}$. Then $\lambda_1(\Lambda_q^{\perp}(A)) \leq \sqrt{m} q^{m-n}$.

Suppose to have (q, A, β) where $\beta = \gamma q^{m-n}$ and $\gamma = \exp(n)$, we can solve $\mathsf{SIS}_{q,m,\beta}$ in polynomial time by LLL.

Corollary

Let A be in $\mathbb{Z}_q^{n \times m}$. Then $\lambda_1(\Lambda_q^{\perp}(A)) \leq \sqrt{m} q^{m-n}$.

Suppose to have (q, A, β) where $\beta = \gamma q^{m-n}$ and $\gamma = \exp(n)$, we can solve $\mathsf{SIS}_{q,m,\beta}$ in polynomial time by LLL.

Usually in crypto: $\beta = \text{subexp}(n) q^{m-n}$.

A Worst-case to Average-case Reduction: SIVP to SIS

Shortest Independent Vectors (SIVP_γ)

Given a n-dimensional full-rank lattice basis B and $\gamma \in \mathbb{R}$, find a set S of linearly independent vectors in $\mathcal{L}(B)$ such that $\|S\| \leq \gamma \lambda_n(B)$.

A Worst-case to Average-case Reduction: SIVP to SIS

Shortest Independent Vectors (SIVP_γ)

Given a n-dimensional full-rank lattice basis B and $\gamma \in \mathbb{R}$, find a set S of linearly independent vectors in $\mathcal{L}(B)$ such that $\|S\| \leq \gamma \lambda_n(B)$.

Theorem [GPV08]

For any $m(n), \beta(n)$ polynomial in n and $\gamma = \beta \cdot \widetilde{O}(\sqrt{n})$ there exists a polynomial time reduction from solving SIVP_γ in the worst case to solving $\mathsf{SIS}_{q,m,\beta}$ on the average, for any prime $q \geq \beta(n) \cdot \omega(\sqrt{\log n})$.

Proof Sketch:

1. IncIVD to SIS (A-W)
2. SIVP to IncIVD (Lattice Preserving Reduction)

Proof Sketch:

1. IncIVD to SIS (A-W)
2. SIVP to IncIVD (Lattice Preserving Reduction)

Incremental Independent Vectors Decoding ($\mathsf{IncIVD}^{\phi}_{\gamma,g}$)

Given a n-dimensional full-rank lattice basis B, $\gamma \in \mathbb{R}$, a set S of linearly independent vectors in $\mathcal{L}(B)$, such that $\|S\|_2 \geq \gamma\phi(B)$, and a target vector $t \in \mathbb{R}^n$, find $v \in \mathcal{L}(B)$ such that $\|v - t\|_2 \leq \|S\|_2/g$.

Proof Sketch:

1. IncIVD to SIS (A-W)
2. SIVP to IncIVD (Lattice Preserving Reduction)

Incremental Independent Vectors Decoding ($\mathsf{IncIVD}^{\phi}_{\gamma,g}$)

Given a n-dimensional full-rank lattice basis B, $\gamma \in \mathbb{R}$, a set S of linearly independent vectors in $\mathcal{L}(B)$, such that $\|S\|_2 \geq \gamma\phi(B)$, and a target vector $t \in \mathbb{R}^n$, find $v \in \mathcal{L}(B)$ such that $\|v - t\|_2 \leq \|S\|_2/g$.

Question: who is ϕ?

Smoothing parameter $\eta_\varepsilon(\mathcal{L})$

Naive characterization: *"If one picks a noise vector from a Gaussian distribution with radius at least as large as the smoothing parameter, and reduces the noise vector modulo the fundamental parallelepiped of the lattice, then the resulting distribution is very close to uniform.[MR07]"*

For any vector c, x and any $s > 0$, let

$$\rho_{s,c}(x) = \exp(-\pi \|(x - c)/s\|^2)$$

be a *Gaussian* functions centered in c and scaled by a factor s.

The associated (continuos) *Gaussian distribution* can be defined by the probability density function

$$D_{s,c}(x) = \frac{\rho_{s,c}(x)}{s^n}$$

for each $x \in \mathbb{R}^n$.

For any vector c, x and any $s > 0$, let

$$\rho_{s,c}(x) = \exp(-\pi \|(x-c)/s\|^2)$$

be a *Gaussian* functions centered in c and scaled by a factor s.

The associated (continuos) *Gaussian distribution* can be defined by the probability density function

$$D_{s,c}(x) = \frac{\rho_{s,c}(x)}{s^n}$$

for each $x \in \mathbb{R}^n$.

The *discrete Gaussian distribution* of a given $\mathcal{L}$ a lattice is

$$D_{\mathcal{L},s,c}(x) := \frac{D_{s,c}(x)}{D_{s,c}(\mathcal{L})} = \frac{\rho_{s,c}(x)}{\rho_{s,c}(\mathcal{L})} \quad \forall x \in \mathcal{L}$$

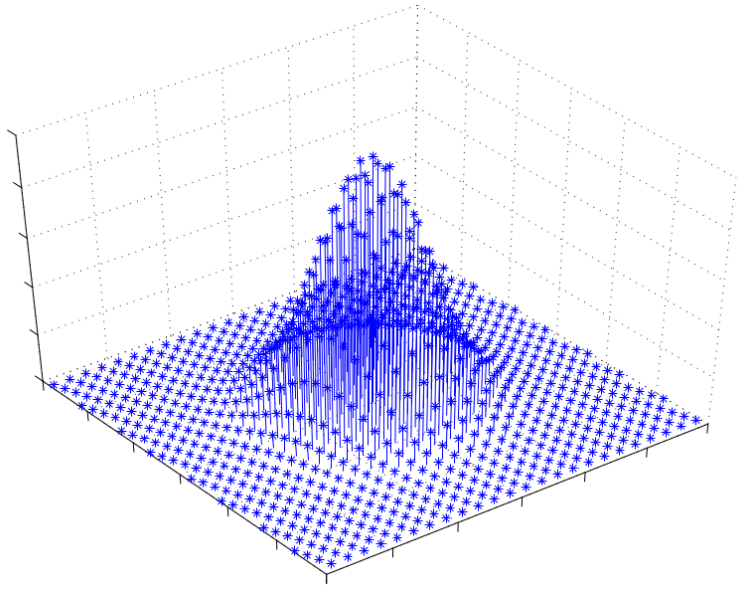

Figure: [MR07]

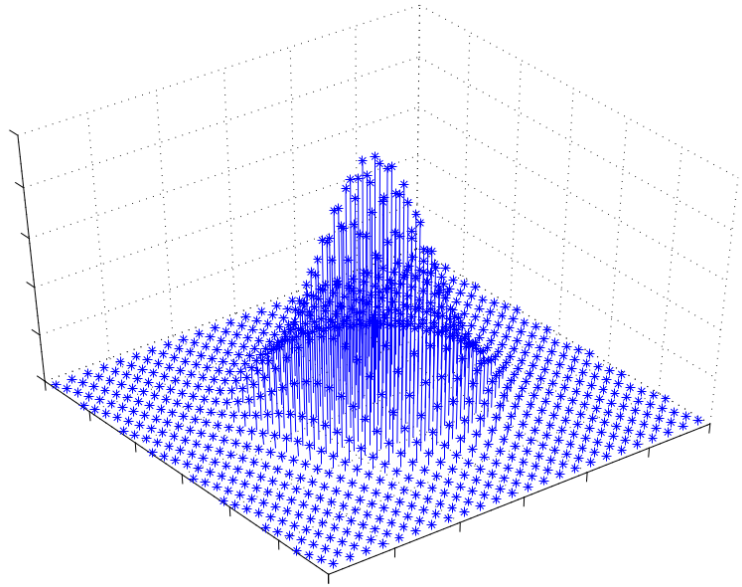

Figure: [MR07]

- If x is distributed according to $D_{s,c}$ and we condition on $x \in \mathcal{L}$, the conditional distribution of x is $D_{\mathcal{L},s,c}$.

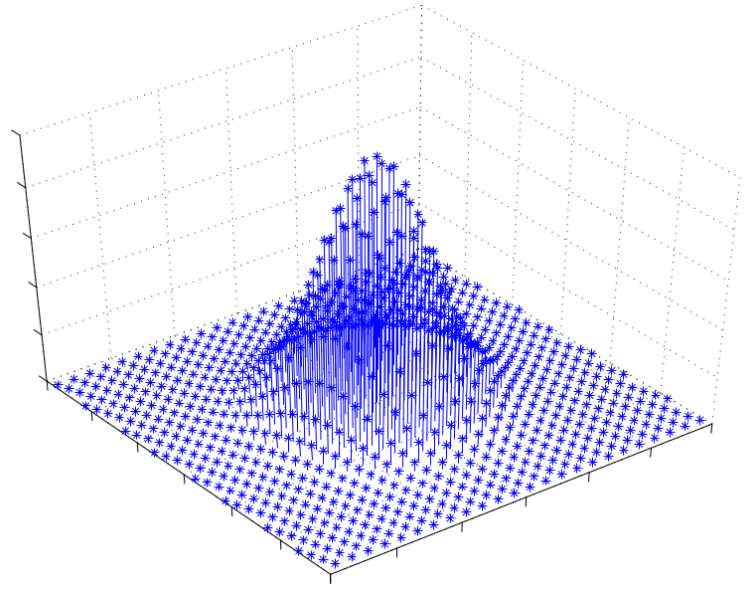

Figure: [MR07]

- If x is distributed according to $D_{s,c}$ and we condition on $x \in \mathcal{L}$, the conditional distribution of x is $D_{\mathcal{L},s,c}$.
- (*) The *smoothing parameter* is the minimal s such that the vectors distributed $D_{\mathcal{L},s,c}$ have an average value very close to c and expected squared distance from c very close to $s^2 n/2\pi$.

Smoothing Parameter

For a n-dimensional lattice $\mathcal{L}$ and $\varepsilon > 0$, the *smoothing parameter* $\eta_\varepsilon(\mathcal{L})$ is the smallest s such that $\rho_{1/s}(\mathcal{L}^\vee \setminus \{0\}) \leq \varepsilon$.

- [MR07] For any $\varepsilon > 0$, $s \geq \eta_\varepsilon(\mathcal{L})$, $c \in \mathbb{R}^n$ and lattice $\mathcal{L}(B)$, the statistical distance between $D_{s,c} \mod \mathcal{P}(B)$ and the uniform distribution over $\mathcal{P}(B)$ is at most $\varepsilon/2$.

Smoothing Parameter

For a n-dimensional lattice $\mathcal{L}$ and $\varepsilon > 0$, the *smoothing parameter* $\eta_\varepsilon(\mathcal{L})$ is the smallest s such that $\rho_{1/s}(\mathcal{L}^\vee \setminus \{0\}) \leq \varepsilon$.

- [MR07] For any $\varepsilon > 0$, $s \geq \eta_\varepsilon(\mathcal{L})$, $c \in \mathbb{R}^n$ and lattice $\mathcal{L}(B)$, the statistical distance between $D_{s,c} \mod \mathcal{P}(B)$ and the uniform distribution over $\mathcal{P}(B)$ is at most $\varepsilon/2$.
- [GPV08] Let $\mathcal{L}, \mathcal{L}'$ n-dimensional lattices such that $\mathcal{L} \supseteq \mathcal{L}'$. Then for any $\varepsilon \in (0, 1/2)$, any $s \geq \eta_\varepsilon(\mathcal{L}')$ and any $c \in \mathbb{R}^n$,the statistical distance between $D_{s,c} \mod \mathcal{L}'$ and the uniform distribution over $\mathcal{L} \mod \mathcal{L}'$ is at most 2ε

IncIVD to SIS

Theorem [GPV08]

For any $g(n) > 1$ and negligible $\varepsilon(n)$, there exists a probabilistic polynomial time reduction from solving $\mathsf{IncIVD}^{\eta_\varepsilon}_{\gamma,g}$ in the worst case for $\gamma(n) = g(n) \cdot \beta(n) \cdot \sqrt{n}$ to solving $\mathsf{SIS}_{q,m,\beta}$ on the average with non-negligible probability, for any $q(n) \geq \gamma(n) \cdot \omega(\sqrt{\log n})$ and $m(n), \beta(n)$ polynomial in n.

Proof:
Suppose to have an oracle $\mathcal{O}$ that solves $\mathsf{SIS}_{q,m,\beta}$ on average with non-negligible probability. The input of the reduction is an instance (B, S, t) of $\mathsf{InclVD}^{\eta_\varepsilon}_{\gamma,g}$.

1. $j \xleftarrow{\$} \{1, \ldots, m\}$
2. $\alpha \xleftarrow{\$} \{-\beta, \ldots, \beta\} \setminus \{0\}$
3. $c_j \leftarrow: \frac{q}{\alpha} t$
4. For each $i \in [m] \setminus \{j\}$: $c_i \leftarrow 0$
5. $s \leftarrow \frac{q}{\gamma} \|S\|$
6. For each $i \in [m]$: $y_i \leftarrow: D_{\mathcal{L}(B), s, c_i}$
7. $Y \leftarrow: [y_1, \ldots, y_m] \in \mathbb{R}^{n \times m}$
8. $A \leftarrow: B^{-1} Y \mod q$
9. $e \leftarrow \mathcal{O}(q, A, \beta)$
10. $v \leftarrow: \frac{1}{q} Y e$
11. if $\|v - t\| \leq \|S\|/g$: return v else goto 1.

δ-Correctness:

- For any j, α, the distribution of A in statistically close to uniform over $\mathbb{Z}_q^{m\times n}$ and $\mathcal{O}$ outputs a nonzero solution e such that $e_j = \alpha$ with non-negligible probability.
  Proof:
  $\|S\| \geq \gamma\eta_\varepsilon(\mathcal{L}(B)) \Rightarrow s \geq q\eta_\varepsilon(\mathcal{L}(B)) = \eta_\varepsilon(q\mathcal{L}(B))$.
  For $\varepsilon \in (0, 1/2)$ and any $c \in \mathbb{R}^n$,the statistical distance between $D_{\mathcal{L}(B),s,c} \mod q\mathcal{L}(B)$ and the uniform distribution over $\mathcal{L}(B) \mod q\mathcal{L}(B)$ is at most 2ε. Hence each $y_i \mod q\mathcal{L}(B)$ is statistically close to uniform in $\mathcal{L}(B)/q\mathcal{L}(B)$.
  $\exists z \in \mathbb{Z}_q^n$ uniform : $Bz = y_i \mod q \Rightarrow z = B^{-1}y_i \mod q$ is uniform.
  We may assume $e \neq 0$ and $\|e\| \leq \beta$, moreover the outputs of $\mathcal{O}$ is statistically close to uniform for each j, α. Then the probability that $e_j = \alpha$ is $1/(2\beta m) = 1/poly(n)$. ▲

- If e is a valid solution and $e_j = \alpha$, then $v \in \mathcal{L}(B)$.
  Proof:
  $Ae = 0 \mod q \Rightarrow \exists z \in \mathbb{Z}^n : Ae = qz \Rightarrow B^{-1}Ye = qz.$
  Then $v = Ye/q = Bz \in \mathcal{L}(B)$.▲

- If e is a valid solution and $e_j = \alpha$, then $v \in \mathcal{L}(B)$.
  Proof:
  $Ae = 0 \mod q \Rightarrow \exists z \in \mathbb{Z}^n : Ae = qz \Rightarrow B^{-1}Ye = qz$.
  Then $v = Ye/q = Bz \in \mathcal{L}(B)$.▲
- If e is a valid solution and $e_j = \alpha$, then $\|v - t\| \leq \|S\|/g$.
  Proof:
  $e_j = \alpha \Rightarrow t = \frac{Ce}{q}$.
  For each i exists $w_i \in \mathcal{L}(B)$ $w_i + c_i = y_i$. Then
  $v - t = \frac{1}{q}(W + C)e - t = \frac{1}{q}(W + C) - \frac{1}{q}Ce = \frac{1}{q}We$.
  It is a combination of vectors sampled by $D_{\mathcal{L}(B),s,0}$. Since
  $s \geq \eta_\varepsilon$, $\|W\| \leq s\sqrt{n}$.
  $\|v - t\| \leq \frac{\|e\|\|W\|}{q} \leq \frac{\beta s\sqrt{n}}{q} = \frac{\gamma s\sqrt{n}}{qg\sqrt{n}} = \frac{\|S\|}{g}$.▲

- If e is a valid solution and $e_j = \alpha$, then $v \in \mathcal{L}(B)$.
  Proof:
  $Ae = 0 \mod q \Rightarrow \exists z \in \mathbb{Z}^n : Ae = qz \Rightarrow B^{-1}Ye = qz$.
  Then $v = Ye/q = Bz \in \mathcal{L}(B)$.▲
- If e is a valid solution and $e_j = \alpha$, then $\|v - t\| \leq \|S\|/g$.
  Proof:
  $e_j = \alpha \Rightarrow t = \frac{Ce}{q}$.
  For each i exists $w_i \in \mathcal{L}(B)$ $w_i + c_i = y_i$. Then
  $v - t = \frac{1}{q}(W + C)e - t = \frac{1}{q}(W + C) - \frac{1}{q}Ce = \frac{1}{q}We$.
  It is a combination of vectors sampled by $D_{\mathcal{L}(B),s,0}$. Since
  $s \geq \eta_\varepsilon$, $\|W\| \leq s\sqrt{n}$.
  $\|v - t\| \leq \frac{\|e\|\|W\|}{q} \leq \frac{\beta s\sqrt{n}}{q} = \frac{\gamma s\sqrt{n}}{qg\sqrt{n}} = \frac{\|S\|}{g}$.▲
- 1.-10. runs in polynomial time.

Termination: We showed that with non-negligible probability 1.-10. return a valid v. Let $\varepsilon = 2^{-\delta}$ with $\delta \gg 0$. We can approximate the probability to give $\mathcal{O}$ a valid input by $1 - 2^{-\delta}$, and the probability that $\|W\| \leq s\sqrt{n}$, too. Then the total success probabily is lower bonded by

$$(1 - 2^{-\delta})^2(1/2m\beta).$$

Thus we expected at most $poly(n)$ loops.

Correctness: δ-Correctness+Termination.

□

Termination: We showed that with non-negligible probability 1.-10. return a valid v. Let $\varepsilon = 2^{-\delta}$ with $\delta \gg 0$. We can approximate the probability to give $\mathcal{O}$ a valid input by $1 - 2^{-\delta}$, and the probability that $\|W\| \leq s\sqrt{n}$, too. Then the total success probabily is lower bonded by

$$(1 - 2^{-\delta})^2(1/2m\beta).$$

Thus we expected at most $poly(n)$ loops.

Correctness: δ**-Correctness+Termination**.

□

Question: Are we cheating?

SIVP to IncIVD

"Lattice-preserving reductions allow to reduce a (worst-case) lattice problem over a given class of lattices to another (worst-case) lattice problem over the same class of lattices."[Mic07]

Theorem

For any $\gamma(n) \geq 1$ there exists a reduction from SIVP_γ to $\mathsf{IncIVD}^{\lambda_n}_{\gamma,4}$.

Proof:
Given B a basis of a full-rank n-lattice we want to construct a set S of n linearly independent vectors such that $\|S\| \leq \gamma\lambda_n(B)$.
Let $\mathcal{O}$ be an oracle that solves IncIVD.
As input we set $S = B$.

1. Find $s \in S \setminus \{s\}$ such that $\|s\| = \|S\|$
2. Select t a vector orthogonal to $S \setminus \{s\}$ such that $\|t\| = \frac{\|S\|}{2}$
3. $v \leftarrow: \mathcal{O}(B, S, t, \gamma, 4)$
4. if $v = \perp$ return S
5. $S \leftarrow (S \setminus \{s\}) \cup \{v\}$
6. goto 1.

□

Correctness*:

- If the oracle fails $\|S\| \leq \gamma \lambda_n(B)$.

Correctness*:

- If the oracle fails $\|S\| \leq \gamma\lambda_n(B)$.
- If $\|v - t\| \leq \frac{\|S\|}{4}$ then $|\|v\| - \|t\|| \leq \frac{\|S\|}{4}$.

$$\frac{\|S\|}{4} \leq \|v\| \leq \frac{3\|S\|}{4}$$

Correctness*:

- If the oracle fails $\|S\| \leq \gamma\lambda_n(B)$.
- If $\|v - t\| \leq \frac{\|S\|}{4}$ then $|\|v\| - \|t\|| \leq \frac{\|S\|}{4}$.

$$\frac{\|S\|}{4} \leq \|v\| \leq \frac{3\|S\|}{4}$$

  - $\|(S \setminus \{s\}) \cup \{v\}\| \leq \|S\|$

Correctness*:

- If the oracle fails $\|S\| \leq \gamma\lambda_n(B)$.
- If $\|v - t\| \leq \frac{\|S\|}{4}$ then $|\|v\| - \|t\|| \leq \frac{\|S\|}{4}$.

$$\frac{\|S\|}{4} \leq \|v\| \leq \frac{3\|S\|}{4}$$

- $\|(S \setminus \{s\}) \cup \{v\}\| \leq \|S\|$
- v is linearly independent respect to $S \setminus \{s\}$:

$$\|t\| - \|v - t\| > 0.$$

Correctness*:

- If the oracle fails $\|S\| \leq \gamma\lambda_n(B)$.
- If $\|v - t\| \leq \frac{\|S\|}{4}$ then $|\|v\| - \|t\|| \leq \frac{\|S\|}{4}$.

$$\frac{\|S\|}{4} \leq \|v\| \leq \frac{3\|S\|}{4}$$

- $\|(S \setminus \{s\}) \cup \{v\}\| \leq \|S\|$
- v is linearly independent respect to $S \setminus \{s\}$:

$$\|t\| - \|v - t\| > 0.$$

Correctness*:

- If the oracle fails $\|S\| \leq \gamma\lambda_n(B)$.
- If $\|v - t\| \leq \frac{\|S\|}{4}$ then $|\|v\| - \|t\|| \leq \frac{\|S\|}{4}$.

$$\frac{\|S\|}{4} \leq \|v\| \leq \frac{3\|S\|}{4}$$

- $\|(S \setminus \{s\}) \cup \{v\}\| \leq \|S\|$
- v is linearly independent respect to $S \setminus \{s\}$:

$$\|t\| - \|v - t\| > 0.$$

Termination: Up to an LLL reduction, we can suppose $\|B\| \leq 2^n\lambda_n(B)$. The quantity $\log \prod_{s\in S} \|s\|$ decreases by a constant at each iteration. Then after at most $O(n^2)$ iterations it the algorithm terminates.

Remark: We can use $\phi = \eta_\varepsilon$ and reduce to $\mathsf{SIVP}_\gamma^{\eta_\varepsilon}$.

Remark: We can use $\phi = \eta_\varepsilon$ and reduce to $\mathsf{SIVP}_\gamma^{\eta_\varepsilon}$.

Theorem [Mic07]

For any n-dimensional lattice $\mathcal{L}$ and $\varepsilon > 0$,

$$\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\frac{\ln(2n(1+e^{-1}))}{\pi}}\lambda_n(\mathcal{L}).$$

In particular, for any function $\omega(\log(n))$ there exists $\varepsilon(n)$ such that $\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\omega}(\log(n))$.

Remark: We can use $\phi = \eta_\varepsilon$ and reduce to $\mathsf{SIVP}_\gamma^{\eta_\varepsilon}$.

Theorem [Mic07]

For any n-dimensional lattice $\mathcal{L}$ and $\varepsilon > 0$,

$$\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\frac{\ln(2n(1+e^{-1}))}{\pi}}\lambda_n(\mathcal{L}).$$

In particular, for any function $\omega(\log(n))$ there exists $\varepsilon(n)$ such that $\eta_\varepsilon(\mathcal{L}) \leq \sqrt{\omega}(\log(n))$.

Thank you!
Questions?

References

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 197–206, 2008.

[Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *computational complexity*, 16:365–411, 2007.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, April 2007.

[Per18] Hilder Vítor Lima Pereira. The SIS problem, 2018. `https://hilder-vitor.github.io/notes.html`,.