Hardness of hinted ISIS from the space time hardness of lattice problems

For the novel lattice assumptions workshop, ICMS Edinburgh, July 2026

Martin R. Albrecht, Russell W. F. Lai, and **Eamonn W. Postlethwaite**

A space time barrier

Take your favourite worst case lattice problem (e.g. SIVP_α), how much does it cost to solve?

A space time barrier

Take your favourite worst case lattice problem (e.g. SIVP_α), how much does it cost to solve?

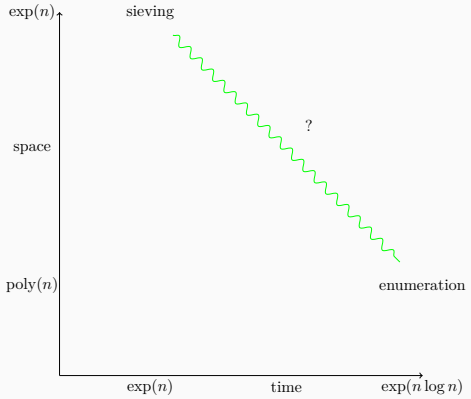

A space time barrier

Take your favourite worst case lattice problem (e.g. SIVP_α), how much does it cost to solve?

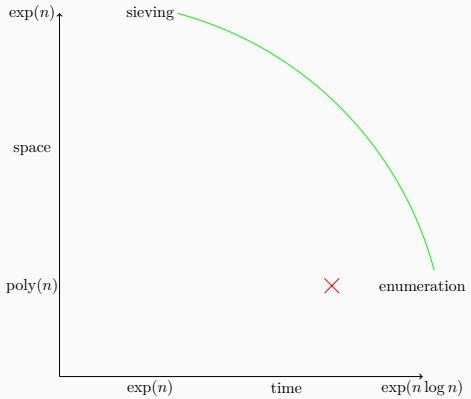

A space time barrier

Take your favourite worst case lattice problem (e.g. SIVP_α), how much does it cost to solve?

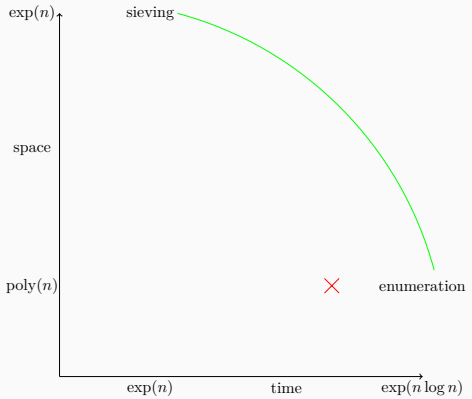

No known $(\mathsf{exp}(o(n \log n)), \mathrm{poly}(n))$: conjectured in [LV20, Conj. 2.1] for GapSVP_α with $\alpha \in \mathrm{poly}(n)$.

A space time barrier

Take your favourite worst case lattice problem (e.g. SIVP_α), how much does it cost to solve?

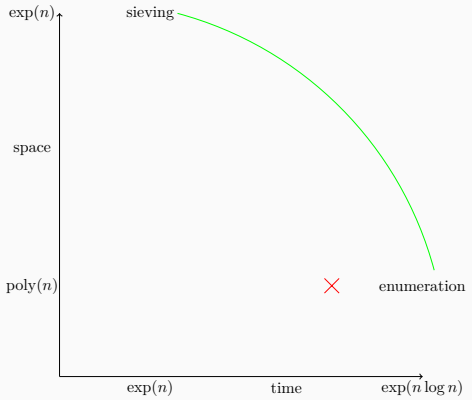

No known $(\mathsf{exp}(o(n \log n)), \mathrm{poly}(n))$: conjectured in [LV20, Conj. 2.1] for GapSVP_α with $\alpha \in \mathrm{poly}(n)$.

This is the ‘space time hardness’ of the title. We do not expect such algorithms to exist: ‘STH’.

SIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $\mathbf{A} \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m}\right)$. Given $\mathbf{A}$ return

$$\begin{aligned} \mathbf{x} \in P^+_{\beta,\mathbf{A}} &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \{\mathbf{u} \in \mathbb{Z}^m : \mathbf{A}\mathbf{u} = \mathbf{0}\} \\ &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \Lambda_q^\perp(\mathbf{A}). \end{aligned}$$

SIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $\mathbf{A} \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m}\right)$. Given $\mathbf{A}$ return

$$\begin{aligned} \mathbf{x} \in P^+_{\beta,\mathbf{A}} &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \{\mathbf{u} \in \mathbb{Z}^m : \mathbf{A}\mathbf{u} = \mathbf{0}\} \\ &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \Lambda_q^\perp(\mathbf{A}). \end{aligned}$$

For $m, \beta \in \mathrm{poly}(n)$ and $q \geq 8n\sqrt{m}\beta$ then $\mathrm{SIVP}_\alpha \leq \mathrm{SIS}_{m,q,\beta}$ for $\alpha \in 8\beta\sqrt{n} \cdot \omega(\sqrt{\log n})$ [MR07].

SIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $\mathbf{A} \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m}\right)$. Given $\mathbf{A}$ return

$$\begin{aligned} \mathbf{x} \in P^{+}_{\beta,\mathbf{A}} &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \{\mathbf{u} \in \mathbb{Z}^m : \mathbf{A}\mathbf{u} = \mathbf{0}\} \\ &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \Lambda_q^{\perp}(\mathbf{A}). \end{aligned}$$

For $m, \beta \in \mathrm{poly}(n)$ and $q \geq 8n\sqrt{m}\beta$ then $\mathrm{SIVP}_\alpha \leq \mathrm{SIS}_{m,q,\beta}$ for $\alpha \in 8\beta\sqrt{n} \cdot \omega(\sqrt{\log n})$ [MR07].

Since $\mathrm{GapSVP}_{n\alpha} \leq \mathrm{SIVP}_\alpha$ an $(\exp(o(n \log n)), \mathrm{poly}(n))$ adversary for $\mathrm{SIS}_{m,q,\beta}$ violates STH.

SIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $\mathbf{A} \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m}\right)$. Given $\mathbf{A}$ return

$$\begin{aligned} \mathbf{x} \in P^{+}_{\beta,\mathbf{A}} &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \{\mathbf{u} \in \mathbb{Z}^m : \mathbf{A}\mathbf{u} = \mathbf{0}\} \\ &= \left(\mathrm{B}^m(\beta) \setminus \{\mathbf{0}\}\right) \cap \Lambda_q^{\perp}(\mathbf{A}). \end{aligned}$$

For $m, \beta \in \mathrm{poly}(n)$ and $q \geq 8n\sqrt{m}\beta$ then $\mathrm{SIVP}_\alpha \leq \mathrm{SIS}_{m,q,\beta}$ for $\alpha \in 8\beta\sqrt{n} \cdot \omega(\sqrt{\log n})$ [MR07].

Since $\mathrm{GapSVP}_{n\alpha} \leq \mathrm{SIVP}_\alpha$ an $(\exp(o(n \log n)), \mathrm{poly}(n))$ adversary for $\mathrm{SIS}_{m,q,\beta}$ violates STH.

If $m \in o(n \log n)$ then an $(\exp(m), \mathrm{poly}(m))$ adversary for $\mathrm{SIS}_{m,q,\beta}$ violates STH.

Hinted lattice problems

Add ‘hints’ to a lattice problem to make it more expressive, but potentially easier to solve.

Hinted lattice problems

Add ‘hints’ to a lattice problem to make it more expressive, but potentially easier to solve.

ISIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $(\mathbf{A}, \mathbf{t}) \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n\right)$. Given $(\mathbf{A}, \mathbf{t})$ return

$$\mathbf{x} \in P_{\beta, \mathbf{A}, \mathbf{t}} = \mathrm{B}^m(\beta) \cap \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\}.$$

Hinted lattice problems

Add ‘hints’ to a lattice problem to make it more expressive, but potentially easier to solve.

ISIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $(\mathbf{A}, \mathbf{t}) \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n\right)$. Given $(\mathbf{A}, \mathbf{t})$ return

$$\mathbf{x} \in P_{\beta, \mathbf{A}, \mathbf{t}} = \mathrm{B}^m(\beta) \cap \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\}.$$

What might a hint be? Note $P_{\mathbf{A},\mathbf{t}} = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\} = \mathbf{v} + \Lambda_q^{\perp}(\mathbf{A})$.

Hinted lattice problems

Add ‘hints’ to a lattice problem to make it more expressive, but potentially easier to solve.

ISIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $(\mathbf{A}, \mathbf{t}) \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n\right)$. Given $(\mathbf{A}, \mathbf{t})$ return

$$\mathbf{x} \in P_{\beta, \mathbf{A}, \mathbf{t}} = \mathrm{B}^m(\beta) \cap \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\}.$$

What might a hint be? Note $P_{\mathbf{A}, \mathbf{t}} = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\} = \mathbf{v} + \Lambda_q^\perp(\mathbf{A})$.

Elements of $\Lambda_q^\perp(\mathbf{A})$, especially short ones, help solve arbitrary ISIS instances.

Hinted lattice problems

Add ‘hints’ to a lattice problem to make it more expressive, but potentially easier to solve.

ISIS: Let $(m, q, \beta) \in \mathbb{N}^2 \times \mathbb{R}_{>0}$ be functions of n and $(\mathbf{A}, \mathbf{t}) \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n\right)$. Given $(\mathbf{A}, \mathbf{t})$ return

$$\mathbf{x} \in P_{\beta, \mathbf{A}, \mathbf{t}} = \mathrm{B}^m(\beta) \cap \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\}.$$

What might a hint be? Note $P_{\mathbf{A},\mathbf{t}} = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} = \mathbf{t}\} = \mathbf{v} + \Lambda_q^\perp(\mathbf{A})$.

Elements of $\Lambda_q^\perp(\mathbf{A})$, especially short ones, help solve arbitrary ISIS instances.

H-ISIS: Let $(k, m, q, \beta) \in \mathbb{N}^3 \times \mathbb{R}_{>0}$ and DIST be functions of n with $\mathrm{DIST}(\mathbf{A})$ an r.v. over $\Lambda_q^\perp(\mathbf{A})$. Let $(\mathbf{A}, \mathbf{t}) \sim \mathrm{U}\left(\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n\right)$ and $\mathbf{U} = (\mathbf{u}_1, \ldots, \mathbf{u}_k) \sim \mathrm{DIST}(\mathbf{A})^k$. Given $(\mathbf{A}, \mathbf{t}, \mathbf{U})$ return $\mathbf{x} \in P_{\beta, \mathbf{A}, \mathbf{t}}$.

(e.g. $k \in \mathrm{poly}(n)$ and $\mathrm{DIST}(\mathbf{A}) = D_{\Lambda_q^\perp(\mathbf{A}), s}$.)

Reductions (many conditions omitted)

Let $\gamma > 1$ then $2^{O(m \log \gamma)}$ calls to a H-ISIS$_\beta$ adversary realises

Reductions (many conditions omitted)

Let $\gamma > 1$ then $2^{O(m \log \gamma)}$ calls to a H-ISIS$_\beta$ adversary realises

where H-SIS is the natural hinted variant of SIS (omitted problem parameters are equal).

Reductions (many conditions omitted)

Let $\gamma > 1$ then $2^{O(m \log \gamma)}$ calls to a H-ISIS$_\beta$ adversary realises

where H-SIS is the natural hinted variant of SIS (omitted problem parameters are equal).

Whether this is reduction is interesting depends on DIST, but cf. 2β from the 'standard'

$$\text{SIS}_{2\beta} \leq \text{ISIS}_\beta.$$

Reductions (many conditions omitted)

Let $\gamma > 1$ then $2^{O(m \log \gamma)}$ calls to a H-ISIS$_\beta$ adversary realises

where H-SIS is the natural hinted variant of SIS (omitted problem parameters are equal).

Whether this is reduction is interesting depends on DIST, but cf. 2β from the 'standard'

$$\text{SIS}_{2\beta} \leq \text{ISIS}_\beta.$$

N.b. the $2^{O(m \log \gamma)}$ calls are 'memory local': H-SIS$_{\beta/\gamma}$ memory equals that of H-ISIS$_\beta$.

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

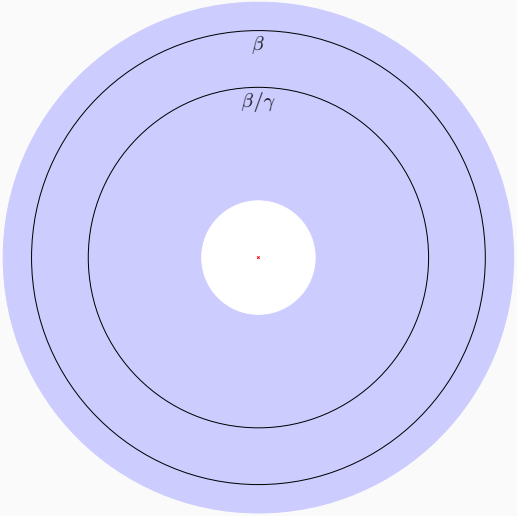

A trivial H-SIS$_{\beta/\gamma}$ instance.

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

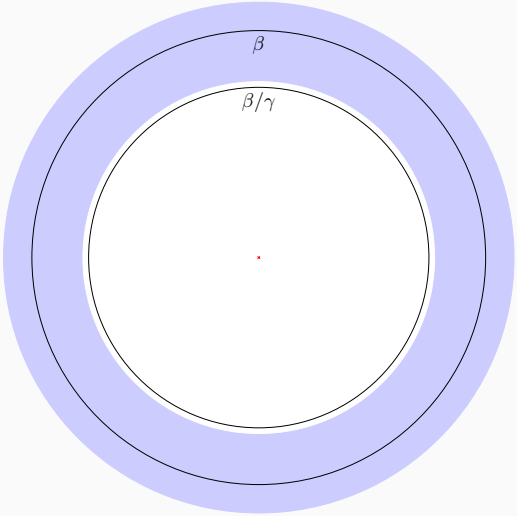

An apparently non trivial H-SIS$_{\beta/\gamma}$ instance.

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

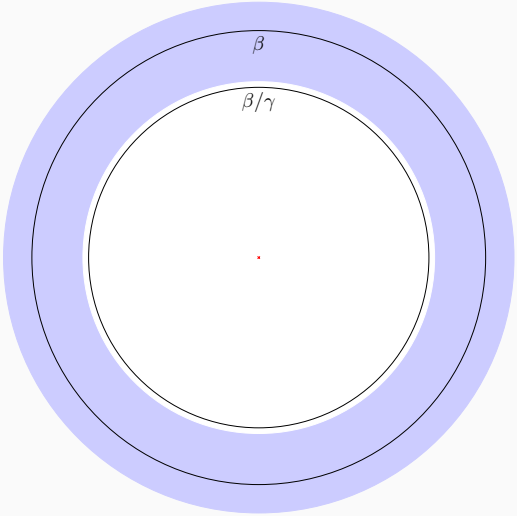

An apparently non trivial H-SIS$_{\beta/\gamma}$ instance.

Not short enough hints to ‘reduce’, not enough hints (or memory) to ‘sieve’.

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

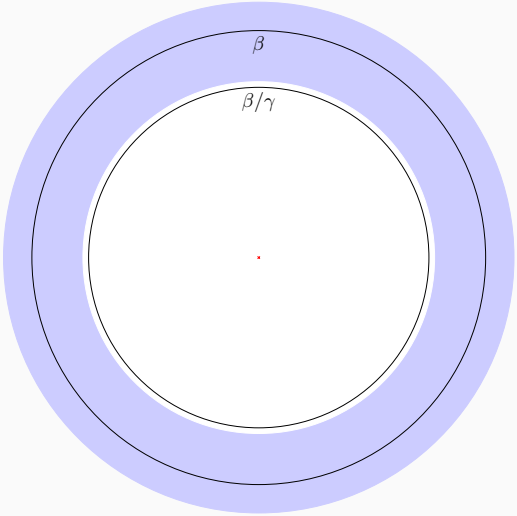

An apparently non trivial H-SIS$_{\beta/\gamma}$ instance.

Not short enough hints to 'reduce', not enough hints (or memory) to 'sieve'.

Returning non zero $\mathbf{x} \in \Lambda_q^{\perp}(\mathbf{A})$ even slightly below the norms of the hints seems space time hard.

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

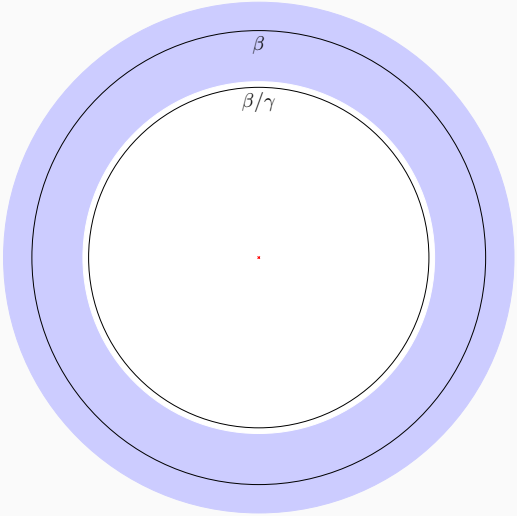

An apparently non trivial H-SIS$_{\beta/\gamma}$ instance.

Not short enough hints to 'reduce', not enough hints (or memory) to 'sieve'.

Returning non zero $\mathbf{x} \in \Lambda_q^{\perp}(\mathbf{A})$ even slightly below the norms of the hints seems space time hard.

If so, a polynomial H-ISIS$_\beta$ adversary and γ such that $\log \gamma \in O(1)$ give an interesting

$$\text{H-SIS}_{\beta/\gamma} \leq \text{H-ISIS}_\beta.$$

Interesting reduction?

Assume $k \in \text{poly}(n)$ hints and DIST has non negligible mass in shaded area.

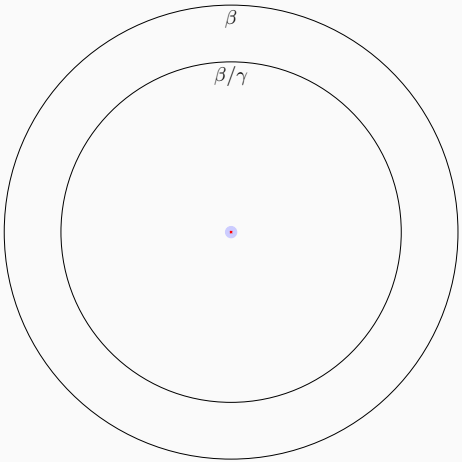

N.b. if hints can be efficiently computed from $\mathbf{A}$:

- some distribution over $q\mathbb{Z}^m \subset \Lambda_q^\perp(\mathbf{A})$,
- wide enough $D_{\Lambda_q^\perp(\mathbf{A}),s}$,

we have H-ISIS $=$ ISIS and H-SIS $=$ SIS.

Reduction overview

Open up the standard $\mathrm{SIS}_{2\beta} \leq \mathrm{ISIS}_{\beta}$ reduction and a new tool[1] from graph theory.

[1] to cryptography

Reduction overview

Open up the standard $\mathrm{SIS}_{2\beta} \leq \mathrm{ISIS}_{\beta}$ reduction and a new tool[1] from graph theory.

Let $\mathcal{A}$ be an adversary against H-ISIS_{β} and construct an adversary $\mathcal{B}$ against H-$\mathrm{SIS}_{\beta/\gamma}$.

[1] to cryptography

Reduction overview

Open up the standard $\mathrm{SIS}_{2\beta} \leq \mathrm{ISIS}_{\beta}$ reduction and a new tool[1] from graph theory.

Let $\mathcal{A}$ be an adversary against H-ISIS_{β} and construct an adversary $\mathcal{B}$ against H-$\mathrm{SIS}_{\beta/\gamma}$.

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$.

[1] to cryptography

Reduction overview

Open up the standard $\mathrm{SIS}_{2\beta} \leq \mathrm{ISIS}_{\beta}$ reduction and a new tool[1] from graph theory.

Let $\mathcal{A}$ be an adversary against H-ISIS_{β} and construct an adversary $\mathcal{B}$ against H-$\mathrm{SIS}_{\beta/\gamma}$.

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$. $\mathcal{B}$ does

1. repeat (†) in pairs until $(\mathbf{v}_\ell, \mathbf{v}'_\ell)$ and $(\mathbf{v}_r, \mathbf{v}'_r)$ are returned and have $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma$,
2. return $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r)$.

[1] to cryptography

Open up the standard $\mathrm{SIS}_{2\beta} \leq \mathrm{ISIS}_{\beta}$ reduction and a new tool[1] from graph theory.

Let $\mathcal{A}$ be an adversary against H-ISIS_{β} and construct an adversary $\mathcal{B}$ against H-$\mathrm{SIS}_{\beta/\gamma}$.

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m, s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta, \mathbf{A}, \mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$. $\mathcal{B}$ does

1. repeat (†) in pairs until $(\mathbf{v}_\ell, \mathbf{v}'_\ell)$ and $(\mathbf{v}_r, \mathbf{v}'_r)$ are returned and have $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma$,
2. return $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r)$.

For sufficiently large s: $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r) \sim_\varepsilon D_{\Lambda_q^\perp(\mathbf{A}), \sqrt{2}s, \mathbf{v}'_r - \mathbf{v}'_\ell}$ with norm (concentrated) at most

$$\sqrt{2m}s + \left\|\mathbf{v}'_r - \mathbf{v}'_\ell\right\|.$$

[1] to cryptography

Reduction overview

Open up the standard $\mathrm{SIS}_{2\beta} \leq \mathrm{ISIS}_{\beta}$ reduction and a new tool[1] from graph theory.

Let $\mathcal{A}$ be an adversary against H-ISIS_{β} and construct an adversary $\mathcal{B}$ against H-$\mathrm{SIS}_{\beta/\gamma}$.

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m, s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta, \mathbf{A}, \mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$. $\mathcal{B}$ does

1. repeat (†) in pairs until $(\mathbf{v}_\ell, \mathbf{v}'_\ell)$ and $(\mathbf{v}_r, \mathbf{v}'_r)$ are returned and have $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma$,
2. return $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r)$.

For sufficiently large s: $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r) \sim_\varepsilon D_{\Lambda_q^\perp(\mathbf{A}), \sqrt{2}s, \mathbf{v}'_r - \mathbf{v}'_\ell}$ with norm (concentrated) at most

$$\sqrt{2m}s + \left\|\mathbf{v}'_r - \mathbf{v}'_\ell\right\|.$$

Choosing $s \leq \beta/2\gamma\sqrt{2m}$ and 1. imply $\sqrt{2m}s + \|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma + \beta/2\gamma = \beta/\gamma$.

[1] to cryptography

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

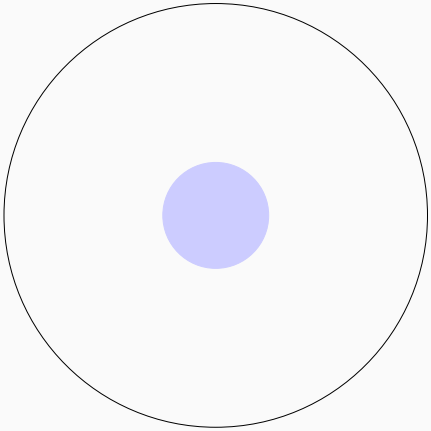

Large circle has radius β, small circle (just for scale) has radius $\beta/2\gamma$.

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

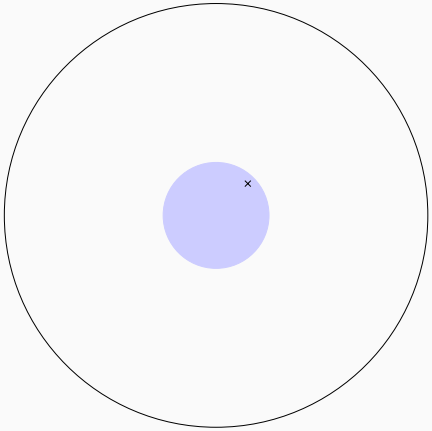

Having a single element in the support with mass one is bad: $\mathbf{Pr}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma] = 1$.

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

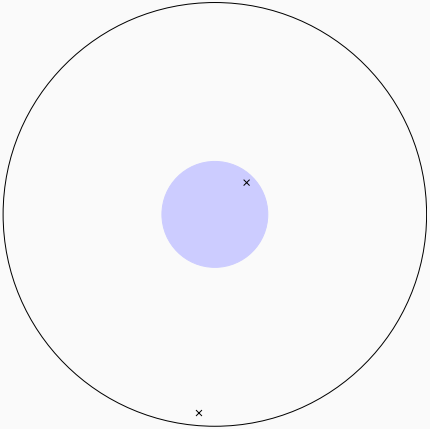

Having two far points, $\mathbf{x}_1, \mathbf{x}_2$, is better.

$$\begin{aligned}\Pr\left[\left\|\mathbf{x} - \mathbf{x}'\right\| \leq \beta/2\gamma\right] &= \Pr\left[\mathbf{x} = \mathbf{x}' = \mathbf{x}_1\right] + \Pr\left[\mathbf{x} = \mathbf{x}' = \mathbf{x}_2\right] \\ &= p_1^2 + p_2^2.\end{aligned}$$

This is minimised at maximum entropy:
$p_1 = p_2 = 1/2 \Rightarrow p_1^2 + p_2^2 = 1/2$.

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

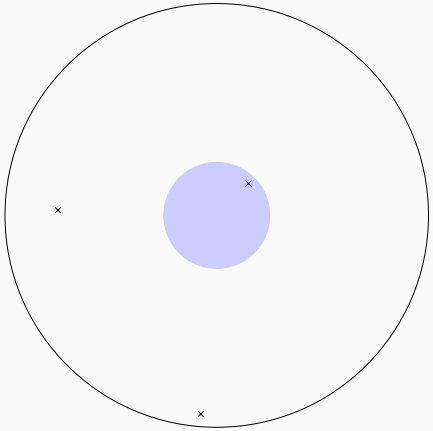

Similarly, three pairwise far points is better still.

$$\Pr\left[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma\right] = p_1^2 + p_2^2 + p_3^2.$$

Again, this is minimised at maximum entropy:
$p_i = 1/3 \Rightarrow \sum_i p_i^2 = 1/3$.

(If pairwise far then generalises to n points.)

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

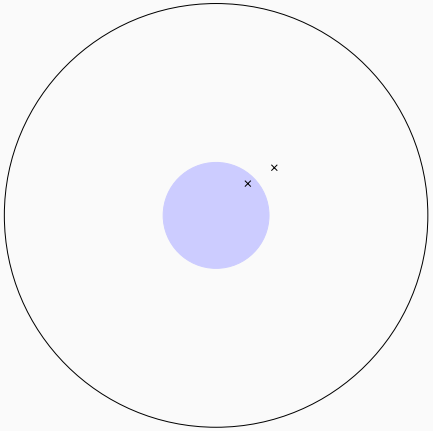

Having two close points is bad, they function as one in our game.

$\mathbf{Pr}[\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma] = 1.$

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

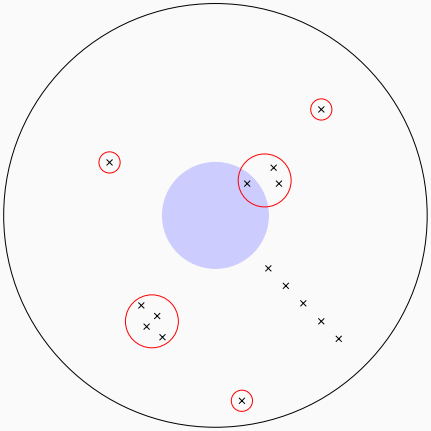

In fact, 'independent cliques' (in which all pairs of points are close and are close to no other points) should be considered as a single point in our game.

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

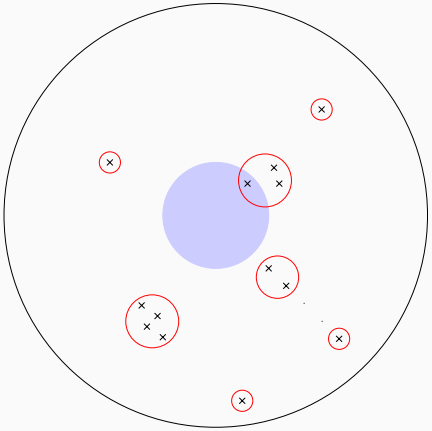

Of the remaining points, form a maximal number of independent cliques by removing points.

Conjecture that

$$1/\#\text{independent cliques}$$

lower bounds probability that $\mathbf{x}$ and $\mathbf{x}'$ are close.

Enter Motzkin–Strauss [MS65]

A game: you may choose any random variable X with finite support in $\mathrm{B}^m(\beta)$.
Which one minimises the probability that $\mathbf{x}, \mathbf{x}' \sim_{\text{i.i.d.}} X$ have $\|\mathbf{x} - \mathbf{x}'\| \leq \beta/2\gamma$?

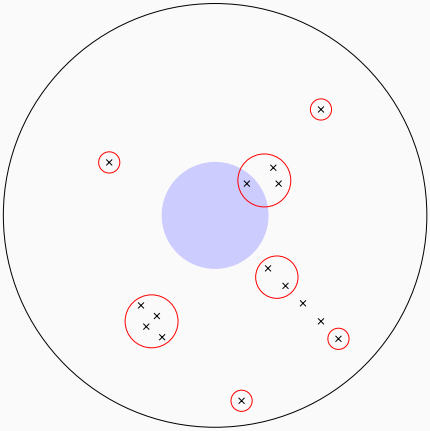

Motzkin–Strauss[a] proves this by

$$\min_{\mathbf{d} \in \Delta_N} \mathbf{d}^{\mathsf{t}} \mathbf{M} \mathbf{d} \geq 1/\alpha(\mathbf{M})$$

- $N = |\mathrm{Supp}(X)|$,
- $\Delta_N = \{\mathbf{x} \in [0,1]^N : \sum_i x_i = 1\}$,
- $\mathbf{M} \in \{0,1\}^N$ is an adjacency matrix with $M_{i,j} = 1$ when two points are close,
- $\alpha(\mathbf{M})$ is the independence number of $\mathbf{M}$: the largest set of points that are pairwise far. By simple geometry, $\alpha(\mathbf{M}) \in 2^{O(m \log \gamma)}$.

[a] this form is a simple corollary in our setting

Reduction recap

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$. $\mathcal{B}$ does

1. repeat (†) in pairs until $(\mathbf{v}_\ell, \mathbf{v}'_\ell)$ and $(\mathbf{v}_r, \mathbf{v}'_r)$ are returned and have $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma$,
2. return $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r)$.

Reduction recap

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$. $\mathcal{B}$ does

1. repeat (†) in pairs until $(\mathbf{v}_\ell, \mathbf{v}'_\ell)$ and $(\mathbf{v}_r, \mathbf{v}'_r)$ are returned and have $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma$,
2. return $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r)$.

By Motzkin–Strauss argument can repeat 1. some $2^{O(m \log \gamma)}/\varepsilon_{\mathcal{A}}^2$ times to succeed w.o.p.

Reduction recap

Subroutine (†): $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$, $\mathbf{v}' \leftarrow \mathcal{A}(1^n, \mathbf{A}, \mathbf{Av}, \mathbf{U})$, if $\mathbf{v}' \in P_{\beta,\mathbf{A},\mathbf{Av}}$ then return $(\mathbf{v}, \mathbf{v}')$. $\mathcal{B}$ does

1. repeat (†) in pairs until $(\mathbf{v}_\ell, \mathbf{v}'_\ell)$ and $(\mathbf{v}_r, \mathbf{v}'_r)$ are returned and have $\|\mathbf{v}'_r - \mathbf{v}'_\ell\| \leq \beta/2\gamma$,
2. return $(\mathbf{v}_\ell - \mathbf{v}'_\ell) - (\mathbf{v}_r - \mathbf{v}'_r)$.

By Motzkin–Strauss argument can repeat 1. some $2^{O(m \log \gamma)}/\varepsilon_{\mathcal{A}}^2$ times to succeed w.o.p.

Three conditions on s:

- large enough that $\mathbf{v} \leftarrow D_{\mathbb{Z}^m,s}$ in (†) satisfies an LHL,
- large enough that 2. is close to a discrete Gaussian with concentrated norm,
- small enough that $s\sqrt{2m} \leq \beta/2\gamma$.

The reduction fails when no satisfying s exists.

A chaining assumption

So far we have assumed STH of H-SIS$_{\beta/\gamma}$ with a sensible hint distribution. Can we iterate?

A chaining assumption

So far we have assumed STH of H-SIS$_{\beta/\gamma}$ with a sensible hint distribution. Can we iterate?

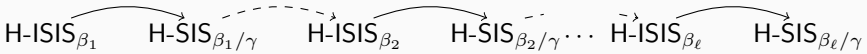

A chaining assumption

So far we have assumed STH of H-SIS$_{\beta/\gamma}$ with a sensible hint distribution. Can we iterate?

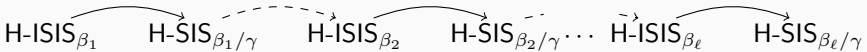

If there exist $\beta_1 > \beta_2 > \cdots > \beta_\ell$, H-ISIS$_{\beta_i}$ adversaries with hint distributions equal to the outputs of the H-SIS$_{\beta_{i-1}/\gamma}$ adversaries, and a shared set of 'good' A, then we realise the above diagram.

A chaining assumption

So far we have assumed STH of H-SIS$_{\beta/\gamma}$ with a sensible hint distribution. Can we iterate?

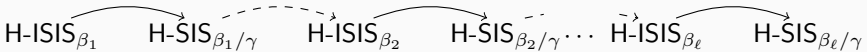

If there exist $\beta_1 > \beta_2 > \cdots > \beta_\ell$, H-ISIS$_{\beta_i}$ adversaries with hint distributions equal to the outputs of the H-SIS$_{\beta_{i-1}/\gamma}$ adversaries, and a shared set of 'good' A, then we realise the above diagram.

(We have some degree of control over the distributions of H-SIS$_{\beta_i/\gamma}$ solutions via ε-smooth rejection sampling [DFPS22].)

A chaining assumption

So far we have assumed STH of H-SIS$_{\beta/\gamma}$ with a sensible hint distribution. Can we iterate?

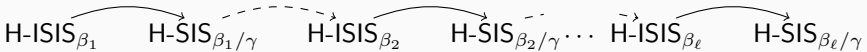

If there exist $\beta_1 > \beta_2 > \cdots > \beta_\ell$, H-ISIS$_{\beta_i}$ adversaries with hint distributions equal to the outputs of the H-SIS$_{\beta_{i-1}/\gamma}$ adversaries, and a shared set of 'good' $\mathbf{A}$, then we realise the above diagram.

(We have some degree of control over the distributions of H-SIS$_{\beta_i/\gamma}$ solutions via ε-smooth rejection sampling [DFPS22].)

If H-ISIS$_{\beta_1}$ has efficient to compute hints then we have a $\ell \cdot 2^{O(m \log \gamma)}$ time reduction giving

$$\mathrm{SIS}_{\beta_\ell/\gamma} \leq \mathrm{ISIS}_{\beta_1}.$$

A chaining assumption

So far we have assumed STH of H-SIS$_{\beta/\gamma}$ with a sensible hint distribution. Can we iterate?

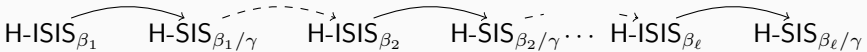

If there exist $\beta_1 > \beta_2 > \cdots > \beta_\ell$, H-ISIS$_{\beta_i}$ adversaries with hint distributions equal to the outputs of the H-SIS$_{\beta_{i-1}/\gamma}$ adversaries, and a shared set of 'good' A, then we realise the above diagram.

(We have some degree of control over the distributions of H-SIS$_{\beta_i/\gamma}$ solutions via ε-smooth rejection sampling [DFPS22].)

If H-ISIS$_{\beta_1}$ has efficient to compute hints then we have a $\ell \cdot 2^{O(m \log \gamma)}$ time reduction giving

$$\mathrm{SIS}_{\beta_\ell/\gamma} \leq \mathrm{ISIS}_{\beta_1}.$$

If the final norm bound β_ℓ/γ is small enough we need only assume STH of GapSVP$_\alpha$.

So what? Double signing in GPV signatures.

The seed of this project was the question: is it *actually* bad to ‘double sign’ in [GPV08] signatures?

[2] or something closely related

So what? Double signing in GPV signatures.

The seed of this project was the question: is it *actually* bad to 'double sign' in [GPV08] signatures?

Let $\mathbf{v}_1, \mathbf{v}_2 \in P_{\beta,\mathbf{A},\mathbf{t}}$, i.e. two signatures for the same message without salt: $\mathbf{t} = \mathrm{H}(m)$.

[2] or something closely related

So what? Double signing in GPV signatures.

The seed of this project was the question: is it *actually* bad to 'double sign' in [GPV08] signatures?

Let $\mathbf{v}_1, \mathbf{v}_2 \in P_{\beta,\mathbf{A},\mathbf{t}}$, i.e. two signatures for the same message without salt: $\mathbf{t} = \mathrm{H}(m)$.

Then $\mathbf{v}_1 - \mathbf{v}_2 \sim_\varepsilon D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$ is likely in $P^+_{\sqrt{2}\beta,\mathbf{A}}$ and definitely in $P_{2\beta,\mathbf{A}}$. A hint!

[2] or something closely related

So what? Double signing in GPV signatures.

The seed of this project was the question: is it *actually* bad to 'double sign' in [GPV08] signatures?

Let $\mathbf{v}_1, \mathbf{v}_2 \in P_{\beta,\mathbf{A},\mathbf{t}}$, i.e. two signatures for the same message without salt: $\mathbf{t} = \mathrm{H}(m)$.

Then $\mathbf{v}_1 - \mathbf{v}_2 \sim_\varepsilon D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$ is likely in $P^+_{\sqrt{2}\beta,\mathbf{A}}$ and definitely in $P_{2\beta,\mathbf{A}}$. A hint!

Does this matter? Not if H-SIS$_{\beta/\gamma}$ with $\mathrm{DIST}(\mathbf{A}) = D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}{}^2$ is STH:

$$\text{H-SIS}_{\beta/\gamma} \leq \text{H-ISIS}_\beta.$$

[2] or something closely related

So what? Double signing in GPV signatures.

The seed of this project was the question: is it *actually* bad to 'double sign' in [GPV08] signatures?

Let $\mathbf{v}_1, \mathbf{v}_2 \in P_{\beta,\mathbf{A},\mathbf{t}}$, i.e. two signatures for the same message without salt: $\mathbf{t} = \mathsf{H}(m)$.

Then $\mathbf{v}_1 - \mathbf{v}_2 \sim_\varepsilon D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$ is likely in $P^+_{\sqrt{2}\beta,\mathbf{A}}$ and definitely in $P_{2\beta,\mathbf{A}}$. A hint!

Does this matter? Not if H-SIS$_{\beta/\gamma}$ with $\mathrm{DIST}(\mathbf{A}) = D_{\Lambda_q^\perp(\mathbf{A}),\sqrt{2}s}$[2] is STH:

$$\text{H-SIS}_{\beta/\gamma} \leq \text{H-ISIS}_\beta.$$

N.b. we cannot follow the original reduction without salts randomising the targets because we cannot answer multiple signing queries on the same message.

[2] or something closely related

Open questions

- How hard, concretely or theoretically, is H-SIS when the hints are sensibly distributed?
- Both our bound on $\alpha(\mathbf{M})$ and our use of Motzkin–Strauss are loose, can we improve them?
- Can the chaining assumption be weakened or removed?
- Can we relate other hinted lattice assumptions to this framework? (We have omISIS but in the ‘wrong’ direction.)

Thanks!

`https://eprint.iacr.org/2026/187` to appear CRYPTO'26.

References

[LV20]: Lombardi A., Vaikuntanathan V.: Fiat–Shamir for Repeated Squaring with Applications to PPAD-Hardness and VDFs, CRYPTO'2020

[MR07]: Micciancio D., Regev O.: Worst-case to average-case reductions based on Gaussian measures, SIAM Journal on Computing, Vol. 37, Iss. 1 (2007)

[MS65]: Motzkin T.S., Straus E.G.: Maxima for graphs and a new proof of a theorem of Turán, Canad. J. Math. 17, 533–540 (1965)

[DFPS22]: Devevey J., Fawzi O., Passelègue A., Stehlé D.: On Rejection Sampling in Lyubashevsky's Signature Scheme, ASIACRYPT'22

[GPV08]: Gentry C., Peikert C., Vaikuntanathan V.: Trapdoors for Hard Lattices and New Cryptographic Constructions, STOC'08