Hardness of the Shortest Vector Problem: A Simplified Proof and a Survey

HUCK BENNETT (OREGON STATE UNIVERSITY)

SIMONS INSTITUTE, JUNE 16TH 2022

JOINT WORK WITH CHRIS PEIKERT (UNIVERSITY OF MICHIGAN / ALGORAND)

The Decisional Shortest Vector Problem γ-GapSVP

Def. A *lattice* is the set $\mathcal{L} = \{\sum_{i=1}^{n} a_i \boldsymbol{b}_i : a_1, \ldots a_n \in \mathbb{Z}\}$ for linearly independent $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n \in \mathbb{R}^m$.

Def. The ℓ_p *norm* of $\boldsymbol{x} \in \mathbb{R}^n$ is $\|\boldsymbol{x}\|_p := \left(\sum_{i=1}^{n} |x_i|^p\right)^{1/p}$ for $p \in (1, \infty)$, $\|\boldsymbol{x}\|_\infty = \max_{i \in [n]} |x_i|$.

Def. The *minimum distance* of a lattice $\mathcal{L}$ is $\lambda_1(\mathcal{L}) := \min_{\boldsymbol{x} \in L \setminus \{\boldsymbol{0}\}} \|\boldsymbol{x}\|$.

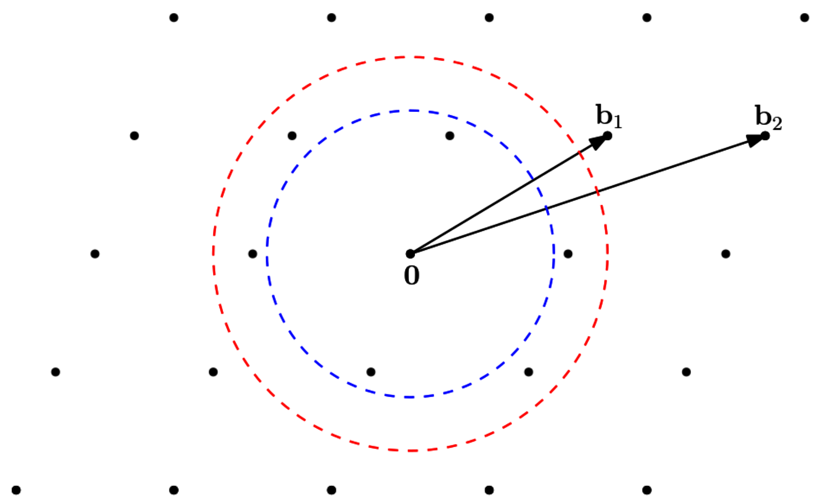

Def. γ-GapSVP for $\gamma = \gamma(n) \geq 1$.

Input: A basis $B = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ of a lattice $\mathcal{L}$ and $r > 0$.

Goal: Decide which of the following the input satisfies:

- **YES** instance: $\lambda_1(\mathcal{L}) \leq r$,
- **NO** instance: $\lambda_1(\mathcal{L}) > \gamma r$.

Simplified Complexity of γ-GapSVP

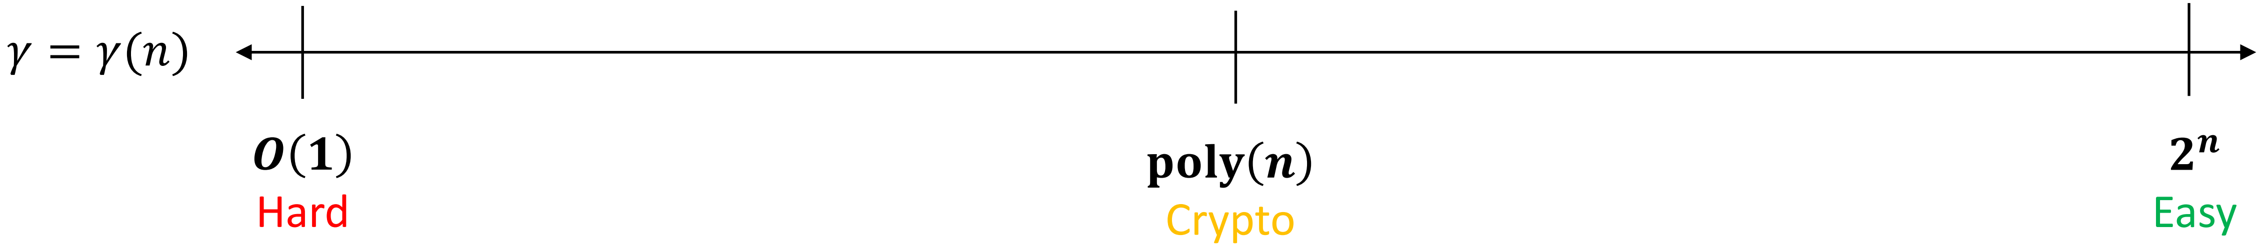

Complexity of γ-GapSVP

Problem: Known hardness results are all under *randomized* assumptions.

We don't even know that *exact* GapSVP is deterministically NP-hard!

Computational Complexity

Sunday, June 12, 2022

I am surprised that the Shortest Vector Problem is not known to be NP-hard, but perhaps I am wrong

$\gamma = \gamma(n)$

Hard

- $\mathbf{1+\varepsilon}$: $2^{\Omega(n)}$-hard under n.u. Gap-ETH
- $\mathbf{O(1)}$: $n^{\omega(1)}$-hard if NP $\not\subseteq$ RP
- $\mathbf{n^{1/\log\log n}}$: $n^{\omega(1)}$-hard if NP $\not\subseteq$ RSUBEXP

[Ajtai \`98, Micciancio \`01, Khot \`05, Haviv-Regev \`12, Aggarwal-(Stephens-Davidowitz) \`18]

Hardness Barriers

- $\sqrt{\mathbf{n/\log n}}$: in coAM
- $\sqrt{\mathbf{n}}$: in coNP

[Goldreich-Goldwasser \`00, Aharonov-Regev \`04]

Crypto

- $\mathbf{n}$: if hard, SIS is hard
- $\mathbf{n^{3/2}}$: if hard, LWE is hard

[Ajtai \`96, Micciancio-Regev \`04, Regev \`09, Lyubashevsky-Micciancio \`09]

Easy

- $\mathbf{2^{n\log\log n/\log n}}$: in P

[Lenstra-Lenstra-Lovász \`82, Schnorr \`87, Gama-Nguyen \`08]

Our Work (**B**-Peikert `22)

What we tried to do:

- Prove deterministic NP-hardness of GapSVP.

What we did do:

- Gave a **simpler randomized NP-hardness reduction**.
  - Key new ingredient: gadget lattices built from **Reed-Solomon codes**.
- Gave concrete **approaches for derandomization**.
- Gave **applications and connections:**
  - Matched the best family of lattices/algorithm for **decoding near Minkowski's bound**.
  - Approach for improved **list-decoding lower bounds** for Reed-Solomon codes.

Derandomization?
No dice.

The Ajtai-Micciancio Approach for Proving NP-Hardness of GapSVP

AS EASY AS STEPS 1-2-3

Step 1: Reducing from γ-GapCVP'

Def. For a vector $\boldsymbol{t}$ and lattice $\mathcal{L}$, $\mathrm{dist}(\boldsymbol{t}, \mathcal{L}) := \min_{\boldsymbol{x} \in L} \|\boldsymbol{x} - \boldsymbol{t}\|$.

Def. Variant of the Closest Vector Problem, $\gamma\text{-GapCVP}'$.

Input: A basis $B = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ of a lattice $\mathcal{L}$, a target vector $\boldsymbol{t}$, and $r > 0$.

Goal: Decide which of the following the input satisfies:

- **YES** instance: There exists $\boldsymbol{x} \in \{0, 1\}^n$ such that $\|B\boldsymbol{x} - \boldsymbol{t}\| \leq r$,
- **NO** instance: For all $w \in \mathbb{Z} \setminus \{0\}$, $\mathrm{dist}(w\boldsymbol{t}, \mathcal{L}) > \gamma r$.

Theorem (Arora-Babai-Stern-Sweedyk '97): γ-GapCVP' is NP-hard for any constant $\gamma \geq 1$.

Step 2: Kannan's Embedding

γ-GapCVP' → GapSVP Attempt 1: Kannan's embedding

$$B, \boldsymbol{t} \mapsto B' := \begin{pmatrix} B & -\boldsymbol{t} \\ 0 & u \end{pmatrix} \text{ for some } u > 0.$$

Analysis: Look at $\|B'\boldsymbol{x}'\|^2 = \|B\boldsymbol{x} - y\boldsymbol{t}\|^2 + |y|^2 u^2$ for $\boldsymbol{x}' = (\boldsymbol{x}, y) \in \mathbb{Z}^{n+1}$.

YES → YES: Consider $\boldsymbol{x}' = (\boldsymbol{x}, 1)^T$ with $\boldsymbol{x} \in \{0,1\}^n$ such that $\|B\boldsymbol{x} - \boldsymbol{t}\|^2 \leq r^2$.

- $\|B\boldsymbol{x} - y\boldsymbol{t}\|^2 = \|B\boldsymbol{x} - \boldsymbol{t}\|^2$ is small.

NO → NO: For $\boldsymbol{x}' = (\boldsymbol{x}, y) \in \mathbb{Z}^{n+1}$

- Case 1, $y \neq 0$: $\|B\boldsymbol{x} - y\boldsymbol{t}\|^2$ is large.
- Case 2, $y = 0$: $\|B\boldsymbol{x} - y\boldsymbol{t}\|^2 = \|B\boldsymbol{x}\|^2$ depends on $\lambda_1(\mathcal{L}(B))$.

Reduction works iff $\lambda_1(L(B))$ is large!

Step 3a: Locally Dense Lattices (LDLs)

α-Locally dense lattices: Lattice/target pairs $\mathcal{L}, \boldsymbol{s}$ with $N \geq 2^{n^\varepsilon}$ vectors in $\mathcal{L}$ at distance $\leq \alpha \cdot \lambda_1(\mathcal{L})$ to $\boldsymbol{s}$ for some consants $\varepsilon > 0, \alpha \in [1/2, 1)$.

The key to showing hardness of $(1/\alpha)$-GapSVP and α-BDD.

- [Ajtai `98, Micciancio `01, Liu-Lyubashevsky-Micciancio `06]
- Also interesting objects in their own right.

Main use of randomness in hardness reductions is constructing LDLs.

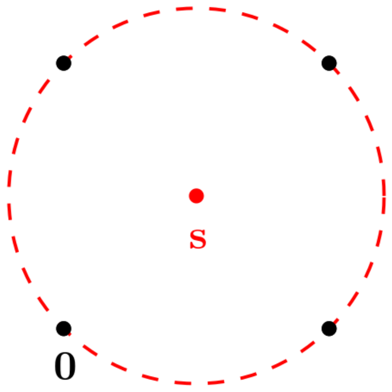

Ex. $\mathcal{L} = \mathbb{Z}^2, \boldsymbol{s} = (1/2,1/2)^T$
$\alpha = 1/\sqrt{2}, N = 4$

Step 3b: Locally Dense Lattices

γ -GapCVP' → GapSVP: Kannan's embedding with locally dense lattice $\mathcal{L}(A), \boldsymbol{s}$.

$$B, \boldsymbol{t} \mapsto B' := \begin{pmatrix} B & -\boldsymbol{t} \\ \beta A & -\beta \boldsymbol{s} \\ 0 & u \end{pmatrix} \text{ for some } \beta, u > 0.$$

Example: GapCVP' → GapSVP in ℓ_∞ with $(A := I_n, \boldsymbol{s} := 1/2 \cdot \mathbf{1})$:

$$B, \boldsymbol{t}, r \mapsto B' := \begin{pmatrix} B & -\boldsymbol{t} \\ 2rI_n & -r\mathbf{1} \\ 0 & r \end{pmatrix}, r' := r$$

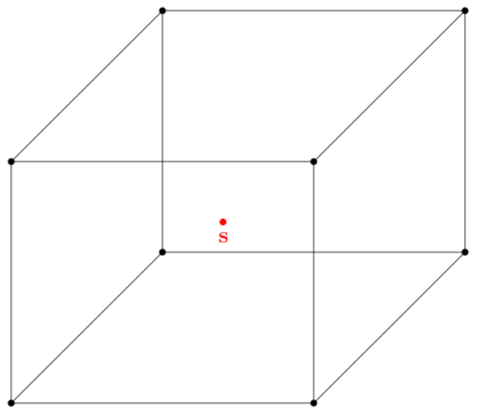

Observation: Reduction worked because $A\boldsymbol{x}$ close to $\boldsymbol{s}$ for each (candidate) coefficient vector $\boldsymbol{x} \in \{0,1\}^n$ of a (candidate) close vector $B\boldsymbol{x}$ to $\boldsymbol{t}$.

Remaining issue: In general, need a correspondence between close vectors in $\mathcal{L}(A)$ to $\boldsymbol{s}$ and in $\mathcal{L}(B)$ to $\boldsymbol{t}$.

- Done using a *random* linear map T.

(Randomized) Constructions of α-locally dense lattices in ℓ_p norms

Construction	Smallest $\alpha = \alpha(p)$	Reference	Notes
Prime Number Lattices	$1/2^{1/p}$	[Ajtai `98, Cai-Nerurkar `99, Micciancio `01]	Derandomizable under strong number-theoretic conjecture
BCH Code "Construction A"	$(1/2 + 1/2^p)^{1/p}$	[Khot '09, Haviv-Regev '12]	Tensors nicely
BCH Code Construction D	$(2/3)^{1/p}$	[Micciancio `12]	Tensors nicely
Sparsified $\mathbb{Z}^n$	$\alpha(p, C)$ with $\lim_{p\to\infty} \alpha(p, C) = 1/2$	[Aggarwal-(Stephens-Davidowitz) `18, **B**-Peikert `20]	2^{Cn} many close vectors, α decreases with p
Exponential Kissing Number Lattices	$\alpha < 0.985$	[Aggarwal-(Stephens-Davidowitz) `18, Vlăduţ `18, **B**-Peikert-Tang `22]	$2^{\Omega(n)}$ many close vectors, non-uniform construction
Reed-Solomon Code Construction A	$1/2^{1/p}$	[**B**-Peikert `22]	Simple. Derandomizable?

Our Locally Dense Lattice Construction

Parity-Check Lattices and Reed-Solomon Codes

Let q be a prime and let $k = q^{\varepsilon}$ for constant $\varepsilon \in (0,1)$.

Key "parity-check" matrix H:

$$H = H_q(k) \coloneqq \begin{pmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 0 & 1 & 2 & 3 & \cdots & q-1 \\ 0 & 1 & 2^2 & 3^2 & \cdots & (q-1)^2 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 1 & 2^{k-1} & 3^{k-1} & \cdots & (q-1)^{k-1} \end{pmatrix} \in \mathbb{F}_q^{k \times q}.$$

Corresponding "parity-check" lattice:

$$\mathcal{L}^{\perp}(H) \coloneqq \{\boldsymbol{x} \in \mathbb{Z}^q : H\boldsymbol{x} \bmod q = \boldsymbol{0}\}$$

Fact: $\mathcal{L}^{\perp}(H) = \mathrm{RS}\left[\mathbb{F}_q, q-k\right] + q\mathbb{Z}^q$.

Parameters and Dense Cosets of $\mathcal{L} = \mathcal{L}^{\perp}\left(H_q(k)\right)$

Minimum distance: For $k < q/2$:

- ℓ_0-minimum distance of $\mathrm{RS}\left[\mathbb{F}_q, q-k\right] = k+1$.
- ℓ_1-minimum distance of $\mathrm{RS}\left[\mathbb{F}_q, q-k\right] = \lambda_1^{(1)}(\mathcal{L}) \geq 2k$ (!!!).
- **Proof [Roth-Siegel `94, Conway-Sloane `99]:** via Newton's identities.

$$H_q(k) := \begin{pmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 0 & 1 & 2 & 3 & \cdots & q-1 \\ 0 & 1 & 2^2 & 3^2 & \cdots & (q-1)^2 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 1 & 2^{k-1} & 3^{k-1} & \cdots & (q-1)^{k-1} \end{pmatrix}$$

Determinant = (# of integer cosets): $\det(\mathcal{L}) = |\mathbb{Z}^q/\mathcal{L}| = q^k$.

Def. $B_{q,h} := \{\boldsymbol{x} \in \{0,1\}^q : \|\boldsymbol{x}\|_1 = h\}$.

Idea (in ℓ_1): Find $\boldsymbol{s} \in \mathbb{Z}^q$ such that $|B_{q,h} \cap (\mathcal{L} - \boldsymbol{s})|$ is subexponentially large.

- Need $h := \alpha \cdot (2k) \leq \alpha \cdot \lambda_1^{(1)}(\mathcal{L})$ to get an ℓ_1 α-LDL.

Pigeonhole principle: When $\alpha > 1/2$ there exists $\boldsymbol{s} \in \mathbb{Z}^q$ such that

$$\mu := |B_{q,h} \cap (\mathcal{L} - \boldsymbol{s})| \geq \binom{q}{h} / q^k \approx q^{(2\alpha - 1)k} = q^{\Omega(q^{\varepsilon})}.$$

Randomized version: $\Pr_{\boldsymbol{s} \sim B_{q,h}}\left[|B_{q,h} \cap (\mathcal{L} - \boldsymbol{s})| \geq \mu/100\right] \geq 0.99$.

Towards Derandomization

Goal: Want explicit center $\boldsymbol{s} \in \mathbb{F}_q^q$ such that $\left|B_{q,h} \cap \left(\mathrm{RS}\left[\mathbb{F}_q, q-k\right] - \boldsymbol{s}\right)\right|$ is subexponentially large for some $h \coloneqq \alpha \cdot (2k) \leq \alpha \cdot \lambda_1^{(1)}(\mathcal{L})$ with $\alpha \in [1/2, 1)$.

- More generally, want explicit-center Reed-Solomon list-decoding lower bounds in ℓ_1/ℓ_p.

Theorem [**B**-Peikert, Kopparty]**:** Would imply improved explicit-center Reed-Solomon list-decoding lower bounds in ℓ_0.

Approach: Discrete Fourier analysis/Weil bound.

- **Used to show:** Best-known explicit (Hamming) Reed-Solomon list-decoding lower bounds [Cheng-Wan `04, Guruswami-Rudra `06].
- **Used to show:** Deterministic MDP hardness [Cheng-Wan `12].

Approach: Point-counting via Gaussian mass.

Summary

- Showing deterministic NP-hardness of GapSVP is a beautiful (still) open question.
- We gave a *simpler, hopefully derandomizable* NP-hardness proof for GapSVP using Reed-Solomon codes.

Hardness of GapSVP: Open Problems

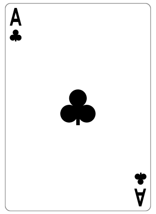

Prove deterministic NP-hardness of GapSVP.

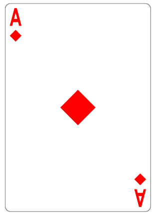

Reduce factoring and discrete log to n^{10}-GapSVP.

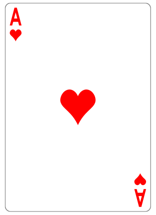

Show $2^{n/c}$-hardness of exact GapSVP for small constant $c > 0$ under a standard complexity assumption.

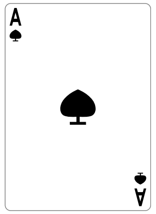

Show superpolynomial hardness of n^{10}-GapSVP under a standard complexity assumption.

Parting Words of Wisdom: Ajtai on Locally Dense Lattices

“[It] may easily happen that other, perhaps in some sense simpler, lattices also have the properties that are required from L to complete the proof… There are different reasons which may motivate the search for such a lattice: to make the proof **deterministic**; to **improve the factor in the approximation result**; to make the proof **simpler**.”

Miklós Ajtai
“The shortest vector problem in L_2 is *NP*-hard for randomized reductions”
STOC, 1998