WORST-CASE TO AVERAGE-CASE REDUCTIONS

PHONG NGUYEN

http://www.di.ens.fr/~pnguyen

October 2024

TODAY

- SIS and LWE
- Variants of SIS and LWE
- The SIS Worst-Case to Average-Case Reduction
    - ➤ Proving that Short Vectors Exist
    - ➤ Mordell's Proof
    - ➤ The SIS Reduction
    - ➤ Sampling Lattice Vectors

SIS AND LWE

THE SIS PROBLEM (1996): SMALL INTEGER SOLUTIONS

- Let $(G,+)$ be a finite Abelian group: $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Ajtai96]. View G as a $\mathbf{Z}$-module.
- Pick $g_1,\ldots,g_m$ uniformly at random from G.

THE SIS PROBLEM (1996): SMALL INTEGER SOLUTIONS

- Let $(G,+)$ be a finite Abelian group: $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Ajtai96]. View G as a $\mathbf{Z}$-module.
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Goal: Find short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\sum_i x_i g_i = 0$, e.g. $||x|| \leq m\,(\#G)^{1/m}$.

THE SIS PROBLEM (1996): SMALL INTEGER SOLUTIONS

- Let $(G,+)$ be a finite Abelian group: $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Ajtai96]. View G as a $\mathbf{Z}$-module.
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Goal: Find short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$, e.g. $||x|| \leq m\,(\#G)^{1/m}$.

➤ This is essentially finding a short vector in a (uniform) random lattice of $L_m(G) = \{$ lattices $L\subseteq\mathbf{Z}^m$ s.t. $\mathbf{Z}^m/L \sim G\}$.

WORST-CASE TO AVERAGE-CASE REDUCTION

WORST-CASE TO AVERAGE-CASE REDUCTION

- [Ajtai96]: If one can efficiently solve SIS for $G=(\mathbf{Z}/q_n\mathbf{Z})^n$ on the average, then one can efficiently find short vectors in every n-dim lattice.

WORST-CASE TO AVERAGE-CASE REDUCTION

- [Ajtai96]: If one can efficiently solve SIS for $G=(\mathbf{Z}/q_n\mathbf{Z})^n$ on the average, then one can efficiently find short vectors in every n-dim lattice.

- [GINX16]: This can be generalized to any sequence (G_n) of finite abelian groups, provided that $\#G_n$ is sufficiently large $\geq n^{\Omega(\max(n,\mathrm{rank}(G)))}$ and m too. Ex: $(\mathbf{Z}/2\mathbf{Z})^n$ is not.

WORST-CASE TO AVERAGE-CASE REDUCTION

Generating Hard Instances of Lattice Problems

Extended abstract

M. Ajtai
IBM Almaden Research Center
650 Harry Road, San Jose, CA, 95120
e-mail: ajtai@almaden.ibm.com

ABSTRACT. We give a random class of lattices in $\mathbf{Z}^n$ so that, if there is a probabilistic polynomial time algorithm which finds a short vector in a random lattice with a probability of at least $\frac{1}{2}$ then there is also a probabilistic polynomial time algorithm which solves the following three lattice problems in *every* lattice in $\mathbf{Z}^n$ with a probability exponentially close to one. (1) Find the length of a shortest nonzero vector in an n-dimensional lattice, approximately, up to a polynomial factor. (2) Find the shortest nonzero vector in an n-dimensional lattice L where the shortest vector v is unique in the sense that any other vector whose length is at most $n^c\|v\|$ is parallel to v, where c is a sufficiently large absolute constant. (3) Find a basis $b_1, ..., b_n$ in the n-dimensional lattice L whose length, defined as $\max_{i=1}^n \|b_i\|$, is the smallest possible up to a polynomial factor.

THE SIS ONE-WAY FUNCTION

THE SIS ONE-WAY FUNCTION

- Let (G,+) be a finite Abelian group

THE SIS ONE-WAY FUNCTION

- Let $(G,+)$ be a finite Abelian group
- Pick $g_1,\ldots,g_m$ uniformly at random from G.

THE SIS ONE-WAY FUNCTION

- Let $(G,+)$ be a finite Abelian group
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Let f: short $(x_1,\ldots,x_m)\in\mathbf{Z}^m \longmapsto \Sigma_i\, x_i\, g_i \in G$.

THE SIS ONE-WAY FUNCTION

- Let $(G,+)$ be a finite Abelian group
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Let f: short $(x_1,\ldots,x_m)\in\mathbf{Z}^m \longmapsto \Sigma_i\, x_i\, g_i \in G$.
    - ➤ f is many-to-one.

THE SIS ONE-WAY FUNCTION

- Let $(G,+)$ be a finite Abelian group
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Let f: short $(x_1,\ldots,x_m)\in\mathbf{Z}^m \longmapsto \Sigma_i\, x_i\, g_i \in G$.
  - f is many-to-one.
  - Given $h = \Sigma_i\, y_i\, g_i \in G$, finding a short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $h=\Sigma_i\, x_i\, g_i \in G$ is as hard as SIS.

DUALITY

- Remember the SIS lattice:
  - $g_1,\ldots,g_m$ in some finite Abelian group $(G,+)$
  - $L=\{x=(x_1,\ldots,x_m)\in\mathbf{Z}^m \text{ s.t. } \Sigma_i\, x_i\, g_i = 0\}$

DUALITY

- Remember the SIS lattice:
  - $g_1,\ldots,g_m$ in some finite Abelian group $(G,+)$
  - $L=\{x=(x_1,\ldots,x_m)\in\mathbf{Z}^m \text{ s.t. } \Sigma_i\, x_i\, g_i = 0\}$
- The dual lattice of L is related to the dual group G^x of (additive) characters of G: morphisms from G to $\mathbf{T}=\mathbf{R}/\mathbf{Z}$

DUALITY

- Remember the SIS lattice:
  - $g_1,\ldots,g_m$ in some finite Abelian group $(G,+)$
  - $L=\{x=(x_1,\ldots,x_m)\in\mathbf{Z}^m \text{ s.t. } \Sigma_i\, x_i\, g_i = 0\}$
- The dual lattice of L is related to the dual group G^x of (additive) characters of G: morphisms from G to $\mathbf{T}=\mathbf{R}/\mathbf{Z}$
  - $L^x=\{(y_1,\ldots,y_m)\in\mathbf{R}^m \text{ s.t. for some } s\in G^x,$ for all i $y_i \equiv s(g_i) \pmod 1\}$

WHY?

- Let $x \in L$: $x=(x_1,\ldots,x_m) \in \mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$
- For any character s in G^x, we have $s(\Sigma_i\, x_i\, g_i)=0$ in $\mathbf{T}=\mathbf{R}/\mathbf{Z}$ so $\Sigma_i\, x_i\, s(g_i)=0$
- Hence, for any $(y_1,\ldots,y_m) \in \mathbf{R}^m$ s.t. for all i $y_i \equiv s(g_i) \pmod 1$, we have $\sum_{i=1}^{m} x_i y_i \in \mathbb{Z}$ so $(y_1,\ldots,y_m) \in L^x$
- Reciprocally if $\sum_{i=1}^{m} x_i y_i \in \mathbb{Z}$ whenever $\Sigma_i\, x_i\, g_i = 0$, there is s in G^x s.t. for all i $y_i \equiv s(g_i) \pmod 1$. This defines s over the subgroup generated by the g_i's: if $\sum_{i=1}^{m} a_i g_i = \sum_{i=1}^{m} b_i g_i$ then $s(\sum_{i=1}^{m} a_i g_i) = s(\sum_{i=1}^{m} b_i g_i)$ because $\sum_{i=1}^{m} (a_i - b_i) g_i = 0$ implies that $\sum_{i=1}^{m} (a_i - b_i) y_i \in \mathbb{Z}$

THE LWE PROBLEM: LEARNING (A CHARACTER) WITH ERRORS

- Let $(G,+)$ be any finite Abelian group.
  - ➤ e.g. $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Re05]
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Pick a random character s in G^x.

THE LWE PROBLEM: LEARNING (A CHARACTER) WITH ERRORS

- Let (G,+) be any finite Abelian group.
  - ➤ e.g. $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Re05]
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Pick a random character s in G^x.
- Goal: recover s given $g_1,\ldots,g_m$ and noisy approximations of $s(g_1),\ldots, s(g_m)$.
  - ➤ Ex: Gaussian noise

GAUSSIAN NOISE OVER R

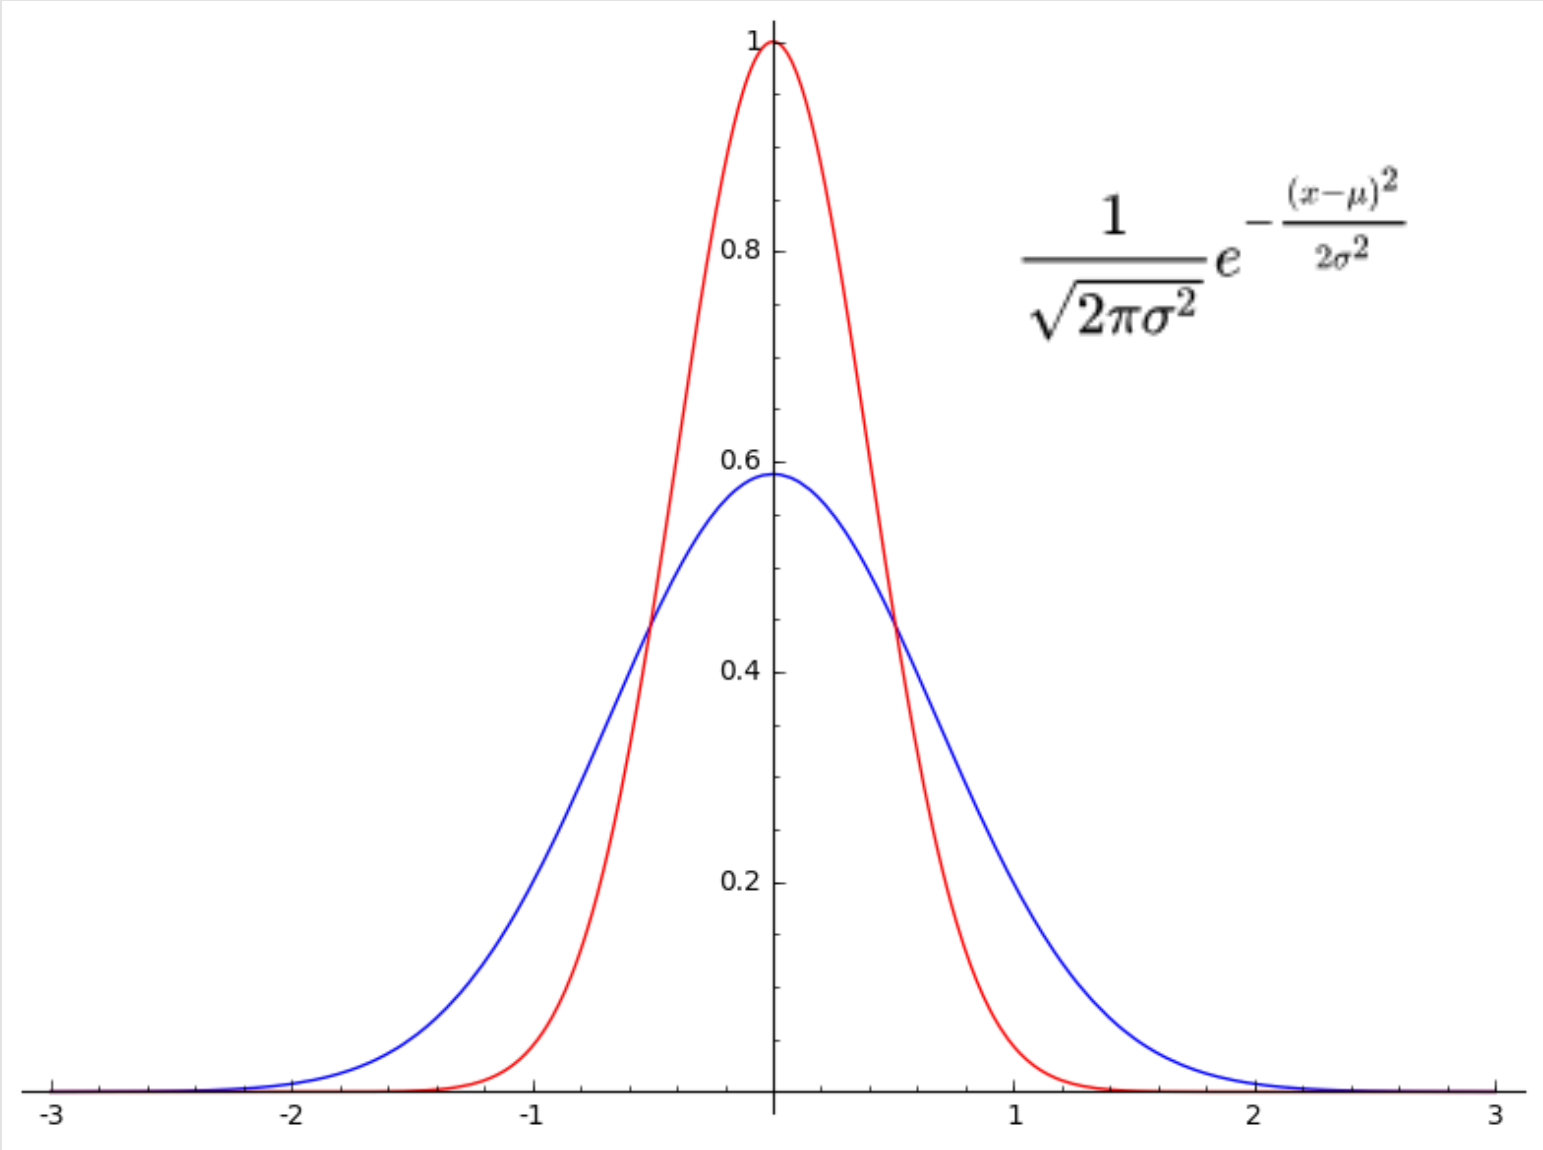

GAUSSIAN NOISE OVER THE TORUS R/Z

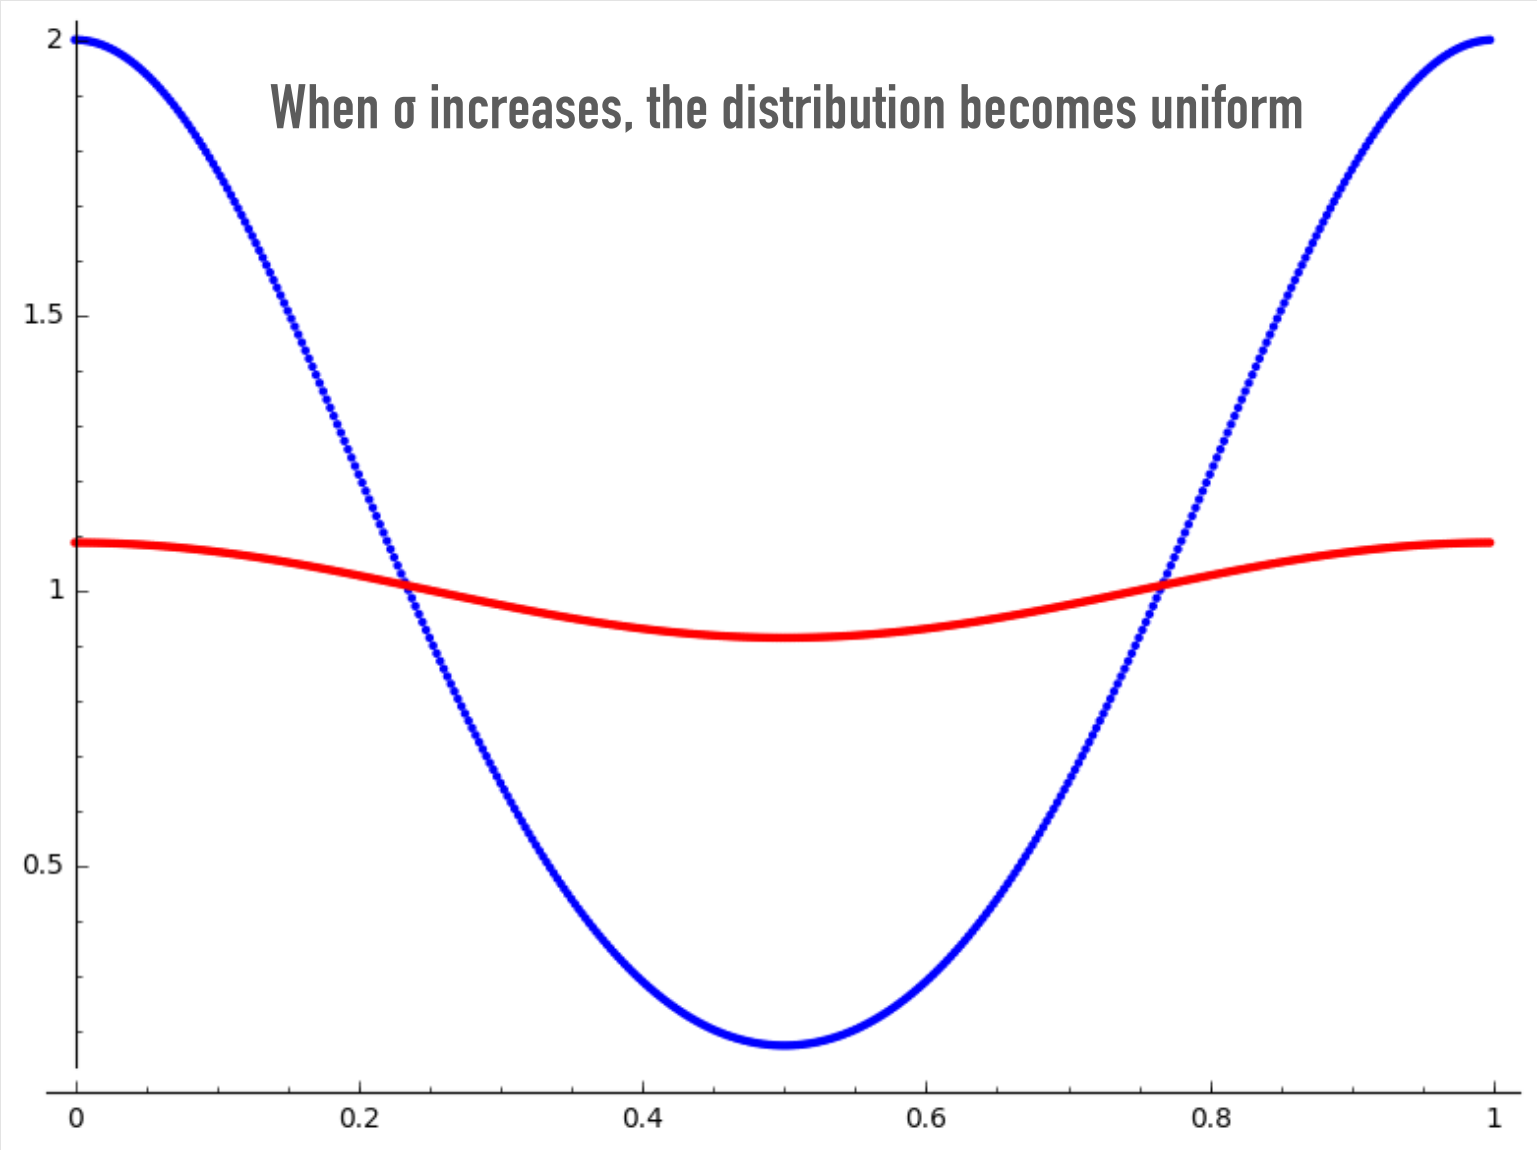

EX: CYCLIC G

- Let $G = \mathbf{Z}/q\mathbf{Z}$. Then $G^{\times} = \left\{ x \mapsto \frac{ax}{q} \bmod 1 \right\}_{a \in \mathbb{Z}/q}$
- Pick $g_1,\ldots,g_m$ uniformly at random mod q.
- Goal: recover $s \in \mathbf{Z}$ given $g_1,\ldots,g_m$ and randomized approximations of sg_1 mod q,..., sg_m mod q.

EX: CYCLIC G

- Let $G = \mathbf{Z}/q\mathbf{Z}$. Then $G^\times = \left\{ x \mapsto \frac{ax}{q} \bmod 1 \right\}_{a \in \mathbb{Z}/q}$
- Pick $g_1,\ldots,g_m$ uniformly at random mod q.
- Goal: recover $s \in \mathbf{Z}$ given $g_1,\ldots,g_m$ and randomized approximations of sg_1 mod q,..., sg_m mod q.

➤ This is exactly a randomized variant of Boneh-Venkatesan's Hidden Number Problem from CRYPTO '96.

HARDNESS OF LWE

- [Regev05]: If one can efficiently solve LWE for $G=(\mathbf{Z}/q_n\mathbf{Z})^n$ on the average, then one can quantum-efficiently find short vectors in every n-dim lattice.
- [GINX16]: This can be generalized to any sequence (G_n) of finite abelian groups, provided that $\#G_n$ is sufficiently large.

DECISIONAL-LWE

- Let $(G,+)$ be any finite Abelian group.
  - ➤ e.g. $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Re05]
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Pick a random character s in G^v.

DECISIONAL-LWE

- Let $(G,+)$ be any finite Abelian group.
  - ➤ e.g. $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Re05]
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Pick a random character s in G^v.
- Goal: Distinguish $(g_1,\ldots,g_m,$ noisy approximations of $s(g_1),\ldots,s(g_m))$ and uniform samples of $G^m x(\mathbf{R}/\mathbf{Z})^m$

SIS AND DECISIONAL-LWE

- Suppose that one finds a short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$.
- What can you say about $\Sigma_i\, x_i\, y_i$

SIS AND DECISIONAL-LWE

- Suppose that one finds a short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$.
- What can you say about $\Sigma_i\, x_i\, y_i$
  - ➤ If $y_1,\ldots,y_m$ are random in $\mathbf{R}/\mathbf{Z}$.

SIS AND DECISIONAL-LWE

- Suppose that one finds a short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$.
- What can you say about $\Sigma_i\, x_i\, y_i$
  - ➤ If $y_1,\ldots,y_m$ are random in $\mathbf{R}/\mathbf{Z}$.
  - ➤ If $y_1,\ldots,y_m$ are approximations of $s(g_1),\ldots, s(g_m)$ with a small Gaussian noise

VARIANTS OF SIS AND LWE

VARIANTS OF SIS AND LWE

- Replace the $\mathbf{Z}$-module by an $\mathbf{R}$-module
- Change the distribution of:
  - ➤ the LWE noise
  - ➤ the secret character

RING TRADEOFFS

- NTRU [HPS98] proposed to use special lattices: better efficiency, yet stronger hardness assumption.

RING TRADEOFFS

- NTRU [HPS98] proposed to use special lattices: better efficiency, yet stronger hardness assumption.

- Starting with [Mi02], one can obtain 'restricted' worst-case to average-case reductions:

RING TRADEOFFS

- NTRU [HPS98] proposed to use special lattices: better efficiency, yet stronger hardness assumption.

- Starting with [Mi02], one can obtain 'restricted' worst-case to average-case reductions:

➤ The worst-case now refers to a special class of lattices, e.g. ideal lattices or module lattices.

M-SIS: SIS OVER MODULES

- Let M be a finite R-module for some ring R: R=$\mathbf{Z}$ in SIS.
- Pick $g_1,\ldots,g_m \in M$ uniformly at random.

M-SIS: SIS OVER MODULES

- Let M be a finite R-module for some ring R: R=**Z** in SIS.
- Pick $g_1,\ldots,g_m \in M$ uniformly at random.
- Goal: find short $(x_1,\ldots,x_m) \in R^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$.

M-SIS: SIS OVER MODULES

- Let M be a finite R-module for some ring R: R=**Z** in SIS.
- Pick $g_1,\ldots,g_m \in M$ uniformly at random.
- Goal: find short $(x_1,\ldots,x_m) \in R^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$.

➤ If R^m is a lattice, this is finding a short vector in some random (module) sublattice of R^m.

M-SIS: SIS OVER MODULES

- Let M be a finite R-module for some ring R: R=**Z** in SIS.
- Pick $g_1,\ldots,g_m \in M$ uniformly at random.
- Goal: find short $(x_1,\ldots,x_m) \in R^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$.

➤ If R^m is a lattice, this is finding a short vector in some random (module) sublattice of R^m.

➤ Ex: NTRU used m=2, R=**Z**[X]/(X^N-1) and M=**Z**[X]/(q,X^N-1) but g_1=public key, g_2=-1.

WORST-CASE TO AVERAGE-CASE REDUCTIONS FOR MODULES

- [LaSt14]: If one can efficiently solve M-SIS for $M=(R/qR)^d$ where R is the ring of integers of a cyclotomic field, then one can efficiently find short vectors in every module lattice of R^d.

WORST-CASE TO AVERAGE-CASE REDUCTIONS FOR MODULES

- [LaSt14]: If one can efficiently solve M-SIS for $M=(R/qR)^d$ where R is the ring of integers of a cyclotomic field, then one can efficiently find short vectors in every module lattice of R^d.
- This generalizes previous ideal-lattice reductions for d=1 [Mi02,LyMi06].

WORST-CASE TO AVERAGE-CASE REDUCTIONS FOR MODULES

- [LaSt14]: If one can efficiently solve M-SIS for $M=(R/qR)^d$ where R is the ring of integers of a cyclotomic field, then one can efficiently find short vectors in every module lattice of R^d.
- This generalizes previous ideal-lattice reductions for d=1 [Mi02,LyMi06].
- Similar results for M-LWE [LaSt14] generalizing Ring-LWE hardness [LPR10].

THE SIS WORST-CASE TO AVERAGE-CASE REDUCTION

PROVING THAT SHORT VECTORS EXIST

SHORT LATTICE VECTORS: MINKOWSKI'S INEQUALITY

SHORT LATTICE VECTORS: MINKOWSKI'S INEQUALITY

- [Minkowski]: Any d-dim lattice L has at least one non-zero vector of norm $\leq 2\frac{\Gamma(1+d/2)^{1/d}}{\sqrt{\pi}}\mathrm{covol}(\mathrm{L})^{1/d} \leq \sqrt{d}\ \mathrm{covol}(\mathrm{L})^{1/d}$

SHORT LATTICE VECTORS: MINKOWSKI'S INEQUALITY

- [Minkowski]: Any d-dim lattice L has at least one non-zero vector of norm $\leq 2\frac{\Gamma(1+d/2)^{1/d}}{\sqrt{\pi}}\mathrm{covol}(\mathrm{L})^{1/d} \leq \sqrt{d}\ \mathrm{covol}(\mathrm{L})^{1/d}$

➤ This is Minkowski's inequality.

Hermann Minkowski

FOUR PROOFS OF MINKOWSKI'S INEQUALITY

- Blichfeldt's proof: «continuous» pigeon-hole principle.

FOUR PROOFS OF MINKOWSKI'S INEQUALITY

- Blichfeldt's proof: «continuous» pigeon-hole principle.

- Minkowski's original proof: sphere packings.

FOUR PROOFS OF MINKOWSKI'S INEQUALITY

- Blichfeldt's proof: «continuous» pigeon-hole principle.
- Minkowski's original proof: sphere packings.
- Siegel's proof: Poisson summation.

FOUR PROOFS OF MINKOWSKI'S INEQUALITY

- Blichfeldt's proof: «continuous» pigeon-hole principle.
- Minkowski's original proof: sphere packings.
- Siegel's proof: Poisson summation.
- Mordell's proof: pigeon-hole principle.

MORDELL'S PROOF (1933)

REMEMBER BLICHFELDT'S PROOF

REMEMBER BLICHFELDT'S PROOF

- The short lattice vector is some u-v where $u,v \in F$ for a well-chosen convex (infinite) set F.

Hans Frederick Blichfeldt

REMEMBER BLICHFELDT'S PROOF

- The short lattice vector is some u-v where $u,v \in F$ for a well-chosen convex (infinite) set F.

Hans Frederick Blichfeldt

- Mordell's proof uses a finite F.

MORDELL'S PROOF (1933)

BEATS THE WORLD AT MATHEMATICS

LEWIS J. MORDEL

Lewis J. Mordel, High School Graduate, Wins Scholarship in Cambridge Over Competitors from Many Countries.

Lewis J. Mordel, a graduate of the Central High School, brought additional honors to his alma mater yesterday, when he was awarded a three-year scholarship in mathematics by St. John's College, Cambridge, England.

Mordel went to Cambridge with nothing but his High School training and competed against graduates of schools and colleges in every part of the world. The examinations were open to all competitors, but for the first time a High School graduate was entered against college men. His entry created laughter instead of serious consideration, but at the conclusion of the examinations, which lasted four days, he stood No. 1 of 250 applicants, with an average of a trifle below 100.

At the Central High School Mordel's ability along mathematical lines was regarded by the members of the faculty as phenomenal. In his Sophomore year he had completed the mathematical course provided for the four-year course and during his last two years in the school he took up the higher mathematics.

To support himself he devoted seven hours of every day to coaching his fellow-students, and on one occasion stood at a blackboard for forty-eight hours in an endeavor to pull a student through an examination. And the examination was passed. At the end of his Senior year he devoted all his time to coaching, having no examinations to take, and in this manner earned enough money to take him to England.

Mordel's present aim is to cover his three years' work sufficiently well to entitle him to a fellowship for four additional years.

MORDELL'S PROOF (1933)

- For q∈**N**, let $\overline{L} = q^{-1}L$ then $[\overline{L} : L] = q^d$.
- Among $>q^d$ points $v_1,\ldots,v_m$ in $\overline{L}$, ∃i≠j s.t. v_i-v_j∈L.

MORDELL'S PROOF (1933)

- For q∈**N**, let $\overline{L} = q^{-1}L$ then $[\overline{L} : L] = q^d$.
- Among $>q^d$ points $v_1,\ldots,v_m$ in $\overline{L}$, ∃i≠j s.t. v_i-v_j∈L.
- There are enough points in a large ball of radius r (r is close to Minkowski's bound in L, but large for $\overline{L}$)

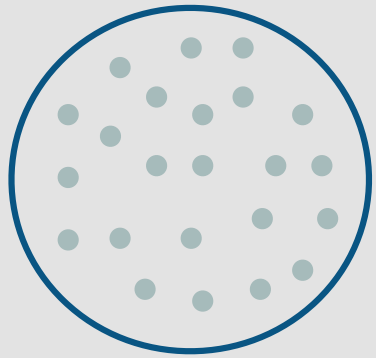

MORDELL'S PROOF (1933)

- For q∈**N**, let $\overline{L} = q^{-1}L$ then $[\overline{L} : L] = q^d$.
- Among $>q^d$ points $v_1,\ldots,v_m$ in $\overline{L}$, ∃i≠j s.t. v_i-v_j∈L.
- There are enough points in a large ball of radius r (r is close to Minkowski's bound in L, but large for $\overline{L}$)

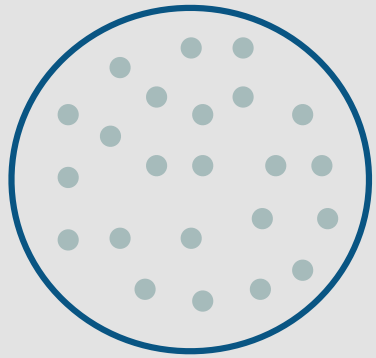

➤ We obtain a short non-zero point in L: norm ≤ 2r.

KEY POINT

- Mordell proved the existence of short lattice vectors by using the existence of short vectors in a special class of higher-dimensional integer lattices.

Louis Joel Mordell

KEY POINT

- Mordell proved the existence of short lattice vectors by using the existence of short vectors in a special class of higher-dimensional integer lattices.

Louis Joel Mordell

- Let distinct $v_1,\ldots,v_m \in \overline{L} = q^{-1}L$.
- Consider the integer lattice L′ formed by all $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\sum_i x_i v_i \in L$.
  - ➤ If $m>q^d$, $\lambda_1(L')\leq\sqrt{2}$.

AN ALGORITHM FROM MORDELL'S PROOF

AN ALGORITHM FROM MORDELL'S PROOF

- Mordell's proof gives an (inefficient) algorithm:

AN ALGORITHM FROM MORDELL'S PROOF

- Mordell's proof gives an (inefficient) algorithm:
  - ➤ Need to generate $>q^d$ lattice points in $\overline{L}$.

AN ALGORITHM FROM MORDELL'S PROOF

- Mordell's proof gives an (inefficient) algorithm:
  - Need to generate $>q^d$ lattice points in $\overline{L}$.
  - Among these exponentially many lattice points, find a difference in L, possibly by exhaustive search.

AN ALGORITHM FROM MORDELL'S PROOF

- Mordell's proof gives an (inefficient) algorithm:
  - Need to generate $>q^d$ lattice points in $\overline{L}$.
  - Among these exponentially many lattice points, find a difference in L, possibly by exhaustive search.
  - Both steps are expensive.

AN ALGORITHM FROM MORDELL'S PROOF

- Mordell's proof gives an (inefficient) algorithm:
  - Need to generate $>q^d$ lattice points in $\overline{L}$.
  - Among these exponentially many lattice points, find a difference in L, possibly by exhaustive search.
  - Both steps are expensive.
- [BGJ14] and [ADRS15] are more efficient randomized variants of Mordell's algorithm: sampling over $\overline{L}$ may allow to sample over L.

WISHFUL THINKING

- To apply the pigeon-hole principle, we need an exponential number m of lattice vectors in $\overline{L}$.

WISHFUL THINKING

- To apply the pigeon-hole principle, we need an exponential number m of lattice vectors in $\overline{L}$.

- Can we get away with a small polynomial number m and make the algorithm efficient? (unlike [BGJ14] and [ADRS15])

WISHFUL THINKING

- To apply the pigeon-hole principle, we need an exponential number m of lattice vectors in $\overline{L}$.

- Can we get away with a small polynomial number m and make the algorithm efficient? (unlike [BGJ14] and [ADRS15])

➤ Maybe if we could find short vectors in certain higher-dimensional random lattices.

THE SIS REDUCTION

REMEMBER SIS

- Let $(G,+)$ be a finite Abelian group: $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Ajtai96]. View G as a $\mathbf{Z}$-module.
- Pick $g_1,\ldots,g_m$ uniformly at random from G.

REMEMBER SIS

- Let $(G,+)$ be a finite Abelian group: $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Ajtai96]. View G as a $\mathbf{Z}$-module.
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Goal: Find short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$, e.g. $||x|| \leq m\,(\#G)^{1/m}$.

REMEMBER SIS

- Let $(G,+)$ be a finite Abelian group: $G=(\mathbf{Z}/q\mathbf{Z})^n$ in [Ajtai96]. View G as a $\mathbf{Z}$-module.
- Pick $g_1,\ldots,g_m$ uniformly at random from G.
- Goal: Find short $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$, e.g. $||x|| \leq m\,(\#G)^{1/m}$.

➤ This is essentially finding a short vector in a (uniform) random lattice of $L_m(G) = \{$ lattices $L\subseteq\mathbf{Z}^m$ s.t. $\mathbf{Z}^m/L \sim G\}$.

OVERLATTICES AND GROUPS

OVERLATTICES AND GROUPS

- If L is n-dim, $\overline{L} = q^{-1}L$ and G=$(\mathbf{Z}/q\mathbf{Z})^n$ then $\overline{L}/L \simeq G$.

OVERLATTICES AND GROUPS

- If L is n-dim, $\overline{L} = q^{-1}L$ and G=$(\mathbf{Z}/\mathrm{q}\mathbf{Z})^{\mathrm{n}}$ then $\overline{L}/L \simeq G$.
- There is an exact sequence: a surjective morphism $\varphi : \bar{L} \to G$ such that $L = \mathrm{Ker}(\varphi)$

$$0 \to L \xrightarrow{1} \bar{L} \xrightarrow{\varphi} G \to 0$$

OVERLATTICES AND GROUPS

- If L is n-dim, $\overline{L} = q^{-1}L$ and G=$(\mathbf{Z}/\mathrm{q}\mathbf{Z})^{\mathrm{n}}$ then $\overline{L}/L \simeq G$.
- There is an exact sequence: a surjective morphism $\varphi : \bar{L} \to G$ such that $L = \mathrm{Ker}(\varphi)$

$$0 \to L \xrightarrow{1} \bar{L} \xrightarrow{\varphi} G \to 0$$

- ϕ is efficiently computable: which ϕ?

OVERLATTICES AND GROUPS

- If L is n-dim, $\overline{L} = q^{-1}L$ and $G=(\mathbf{Z}/q\mathbf{Z})^n$ then $\overline{L}/L \simeq G$.
- There is an exact sequence: a surjective morphism $\varphi : \bar{L} \to G$ such that $L = \mathrm{Ker}(\varphi)$

$$0 \to L \xrightarrow{1} \bar{L} \xrightarrow{\varphi} G \to 0$$

- ϕ is efficiently computable: which ϕ?
- Let $v_1,\ldots,v_m \in \overline{L}$ and define $g_1,\ldots,g_m \in G$ by $g_i=\phi(v_i)$.
- $\Sigma_i\, x_i\, g_i = 0$ for $(x_1,\ldots,x_m)\in\mathbf{Z}^m$ iff $\Sigma_i\, x_i\, v_i \in L$.

WORST-TO-AVERAGE REDUCTION FROM MORDELL'S PROOF

- Sample short $v_1,\ldots,v_m \in \overline{L}$ from a suitable distribution, so that $g_i=\phi(v_i)$ has uniform distribution over $G=(\mathbf{Z}/q\mathbf{Z})^n$

- Call the SIS-oracle on $(g_1,\ldots,g_m)$ to find a short $x=(x_1,\ldots,x_m)\in\mathbf{Z}^m$ s.t. $\Sigma_i\, x_i\, g_i = 0$ in G, i.e. $\Sigma_i\, x_i\, v_i \in L$.

- Return $\Sigma_i\, x_i\, v_i \in L$.

GENERALIZED SIS REDUCTION

GENERALIZED SIS REDUCTION

- The SIS reduction is based on this crucial fact: If B is a reduced basis of a lattice L, then $q^{-1}B$ is a reduced basis of the overlattice $\overline{L} = q^{-1}L$.

GENERALIZED SIS REDUCTION

- The SIS reduction is based on this crucial fact: If B is a reduced basis of a lattice L, then q^{-1}B is a reduced basis of the overlattice $\overline{L} = q^{-1}L$.

- But if G is an arbitrary finite Abelian group, we need to find a reduced basis of some overlattice $\overline{L} \supseteq$ L s.t. $\overline{L}/L \simeq G$, so that we can sample short vectors in $\overline{L}$.

STRUCTURAL LATTICE REDUCTION

STRUCTURAL LATTICE REDUCTION

- In classical lattice reduction, we try to find a good basis of a given lattice.

STRUCTURAL LATTICE REDUCTION

- In classical lattice reduction, we try to find a good basis of a given lattice.

- In structural lattice reduction [GINX16], given a lattice L and a (sufficiently large) finite Abelian group G, we find a good basis of some overlattice $\overline{L}$ s.t. $\overline{L}/L \simeq G$.

EASY CASES

EASY CASES

- If $G=(\mathbf{Z}/q\mathbf{Z})^n$, any basis B of a full-rank lattice L in $\mathbf{Z}^n$ can be transformed into a basis $q^{-1}B$ of $\overline{L} = q^{-1}L$, which is $q=\#G^{1/n}$ times shorter.

EASY CASES

- If $G=(\mathbf{Z}/q\mathbf{Z})^n$, any basis B of a full-rank lattice L in $\mathbf{Z}^n$ can be transformed into a basis $q^{-1}B$ of $\overline{L} = q^{-1}L$, which is $q=\#G^{1/n}$ times shorter.

- If $G=\mathbf{Z}^n/L$, the canonical basis of $\overline{L} = \mathbf{Z}^n$ is a short basis, typically $\#G^{1/n}$ times shorter than a short basis of L.

KEYPOINT: FOURIER ANALYSIS

Joseph Fourier

KEYPOINT: FOURIER ANALYSIS

- Fourier analysis shows that if $\vec{v}_1, \ldots, \vec{v}_m \in \bar{L}$ are chosen from a suitable (short) distribution, $g_i = \varphi(\vec{v}_i)$ has uniform distribution over G.

- Any probability mass function f over $\bar{L}$ s.t. for any $\vec{x} \in \bar{L}$, $\sum_{\vec{y} \in L} f(\vec{x} + \vec{y}) \approx \frac{1}{\#G}$.
  - ➤ Ex: discrete Gaussian distribution.

KEYPOINT: FOURIER ANALYSIS

- Fourier analysis shows that if $\vec{v}_1, \ldots, \vec{v}_m \in \bar{L}$ are chosen from a suitable (short) distribution, $g_i = \varphi(\vec{v}_i)$ has uniform distribution over G.

- Any probability mass function f over $\bar{L}$ s.t.
  for any $\vec{x} \in \bar{L}$, $\sum_{\vec{y} \in L} f(\vec{x} + \vec{y}) \approx \frac{1}{\#G}$.
  - ➤ Ex: discrete Gaussian distribution.

➤ This is a key step: transforming a worst-case into an average-case.

WHY FOURIER ANALYSIS?

- Recall Poisson's summation: if $f : \mathbb{R} \to \mathbb{C}$ is a "nice" function, then $\sum_{n\in\mathbb{Z}} f(n) = \sum_{n\in\mathbb{Z}} \hat{f}(n)$ where $\hat{f}(n) = \int_{-\infty}^{+\infty} f(t)e^{-2i\pi nt}dt$ and more generally $\forall x \in \mathbb{R} \ \sum_{n\in\mathbb{Z}} f(x+n) = \sum_{n\in\mathbb{Z}} \hat{f}(n)e^{2i\pi nx}$
- In arbitrary dimension, if $f : \mathbb{R}^n \to \mathbb{C}$ is a "nice" function and L is a full-rank lattice in $\mathbb{R}^n$, then $\sum_{\vec{x}\in L} f(\vec{x}) = \frac{1}{\mathrm{covol}(L)} \sum_{\vec{y}\in L^\times} \hat{f}(\vec{y})$ and more generally $\forall \vec{x} \in \mathbb{R}^n \ \sum_{\vec{y}\in L} f(\vec{x}+\vec{y}) = \sum_{\vec{y}\in L^\times} \hat{f}(\vec{y})e^{2i\pi\langle\vec{x},\vec{y}\rangle}$

WHY FOURIER ANALYSIS?

- We'd like $\forall \vec{x} \in \mathbb{R}^n \sum_{\vec{y} \in L} f(\vec{x}+\vec{y}) = \sum_{\vec{y} \in L^\times} \hat{f}(\vec{y}) e^{2i\pi\langle \vec{x}, \vec{y}\rangle}$ to be nearly constant over $\vec{x}$.
- $\sum_{\vec{y} \in L^\times} \hat{f}(\vec{y}) e^{2i\pi\langle \vec{x}, \vec{y}\rangle} = \hat{f}(0) + \sum_{\vec{y} \in L^\times \setminus \{0\}} \hat{f}(\vec{y}) e^{2i\pi\langle \vec{x}, \vec{y}\rangle}$ so we want $\sum_{\vec{y} \in L^\times \setminus \{0\}} \hat{f}(\vec{y}) e^{2i\pi\langle \vec{x}, \vec{y}\rangle}$ to be uniformly small.
  - $\lambda_1(L^\times)$ is related to $1/\lambda_n(L)$
  - The Fourier transform of a Gaussian function is Gaussian with inverse parameter: if $f(\vec{x}) = e^{-\pi \|\vec{x}\|^2/s^2}$ then $\hat{f}(\vec{x}) = e^{-\pi s^2 \|\vec{x}\|^2}$
  - If s is big enough, then $\hat{f}(\vec{y})$ will be very small.
  - $\|\vec{y}\| \geq \lambda_1(L^\times) \geq 1/\lambda_n(L)$ so we need s somewhat larger than $\lambda_n(L)$.

SAMPLING LATTICE VECTORS

GAUSSIAN MEASURE

GAUSSIAN MEASURE

- Though lattices are infinite, there is a natural probability distribution over lattice points, introduced by [Ba1993] for transference.

Wojciech Banaszczyk

GAUSSIAN MEASURE

- Though lattices are infinite, there is a natural probability distribution over lattice points, introduced by [Ba1993] for transference.
- This Gaussian measure was implicitly used in [Klein00]'s randomized variant of Babai's nearest-plane algorithm to solve BDD.

Philip N. Klein

Wojciech Banaszczyk

GAUSSIAN MEASURE

- Though lattices are infinite, there is a natural probability distribution over lattice points, introduced by [Ba1993] for transference.
- This Gaussian measure was implicitly used in [Klein00]'s randomized variant of Babai's nearest-plane algorithm to solve BDD.
- [Regev2005] noted that the Gaussian measure could sometimes be sampled.

Oded Regev

GAUSSIAN MEASURE

- Though lattices are infinite, there is a natural probability distribution over lattice points, introduced by [Ba1993] for transference.
- This Gaussian measure was implicitly used in [Klein00]'s randomized variant of Babai's nearest-plane algorithm to solve BDD.
- [Regev2005] noted that the Gaussian measure could sometimes be sampled.
- [GPV2008] rediscovered [Klein00] and showed that it samples from the Gaussian measure.

GAUSSIAN MEASURE

- Center c, parameter s
- Mass of x∈L proportional to

$$\rho_{s,\vec{c}}(\vec{x}) = e^{-\pi \| \frac{\vec{x}-\vec{c}}{s} \|^2}$$

GAUSSIAN MEASURE

- Center c, parameter s
- Mass of x∈L proportional to

$$\rho_{s,\vec{c}}(\vec{x}) = e^{-\pi \| \frac{\vec{x}-\vec{c}}{s} \|^2}$$

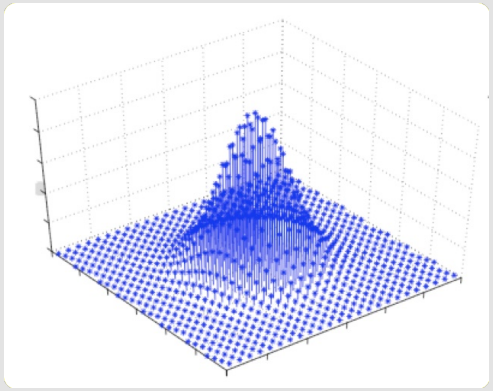

- The distribution is independent of the basis.

GAUSSIAN MEASURE

- Center c, parameter s
- Mass of x∈L proportional to

$$\rho_{s,\vec{c}}(\vec{x}) = e^{-\pi \| \frac{\vec{x}-\vec{c}}{s} \|^2}$$

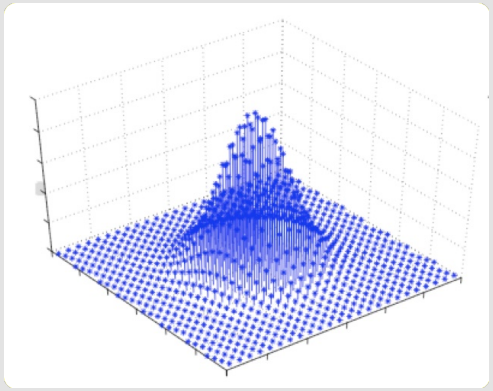

- The distribution is independent of the basis.

- Introduced in [Ba93], then used in cryptography in [Cai99,Regev03,MiRe04,…]

INTUITION

- Beyond a so-called smoothing parameter, the discrete Gaussian distribution behaves like a continuous Gaussian.
- This smoothing parameter is a bit larger than $\lambda_n(L)$.

BABAI'S NEAREST PLANE ALGORITHM

- Input: a basis $(\vec{b}_1, \ldots, \vec{b}_n)$ of a lattice L and a target $\vec{t}$ in span(L).
- Output: a lattice point $\vec{u}$ such that $\vec{t} - \vec{u} \in \left\{ \sum_{i=1}^{n} x_i \vec{b}_i^{\star}, -1/2 \leq x_i < 1/2 \right\}$ where the $\vec{b}_i^{\star}$'s are the Gram-Schmidt orthogonalization.

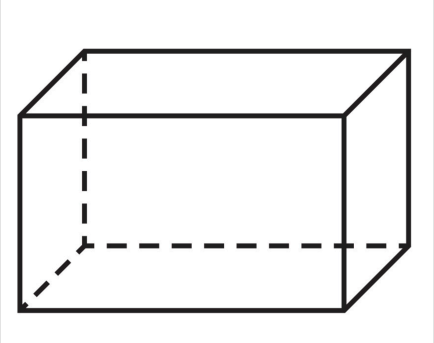

- Return

$$\vec{u} = \sum_{i=1}^{n} \lfloor \mu_i \rceil \vec{b}_i$$

BABAI'S NEAREST PLANE ALGORITHM

- Input: a basis $(\vec{b}_1, \ldots, \vec{b}_n)$ of a lattice L and a target $\vec{t}$ in span(L).
- Output: a lattice point $\vec{u}$ such that $\vec{t} - \vec{u} \in \left\{ \sum_{i=1}^{n} x_i \vec{b}_i^{\star}, -1/2 \leq x_i < 1/2 \right\}$ where the $\vec{b}_i^{\star}$'s are the Gram-Schmidt orthogonalization.

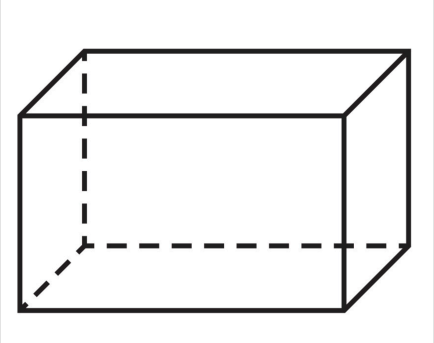

- For i=n downto 1
  - Compute $\mu_i = \dfrac{\langle \vec{t}, \overrightarrow{b_i}^{\star} \rangle}{\| \overrightarrow{b_i}^{\star} \|^2}$
  - $\vec{t} \leftarrow \vec{t} - \lfloor \mu_i \rceil \vec{b}_i$
- Return $\vec{u} = \sum_{i=1}^{n} \lfloor \mu_i \rceil \vec{b}_i$

RANDOMIZING BABAI'S NEAREST PLANE ALGORITHM

- Input: a basis $(\vec{b}_1, \ldots, \vec{b}_n)$ of a lattice L and a target $\vec{t}$ in span(L).
- Output: a "random" lattice point $\vec{u}$ "close" to $\vec{t}$.

- Return

$$\vec{u} = \sum_{i=1}^{n} x_i \vec{b}_i$$

RANDOMIZING BABAI'S NEAREST PLANE ALGORITHM

- Input: a basis $(\vec{b}_1, \ldots, \vec{b}_n)$ of a lattice L and a target $\vec{t}$ in span(L).
- Output: a "random" lattice point $\vec{u}$ "close" to $\vec{t}$.

- For i=n downto 1
  - Compute $\mu_i = \dfrac{\langle \vec{t}, \overrightarrow{b_i}^{\star} \rangle}{\| \overrightarrow{b_i}^{\star} \|^2}$ and

    $x_i = \text{RandomizedRounding}(\mu_i) = \text{ShortElement}(\mu_i + \mathbb{Z})$
  - $\vec{t} \leftarrow \vec{t} - x_i \vec{b}_i$
- Return

$$\vec{u} = \sum_{i=1}^{n} x_i \vec{b}_i$$

GOAL

- We want to output $\vec{u} \in L$ with probability proportional to $\rho_{s,\vec{t}}(\vec{u}) = e^{-\pi\|\vec{u}-\vec{t}\|^2/s^2}$
- Here, $\vec{t} - \vec{u} = \sum_{i=1}^{n} (\mu_i - x_i)\vec{b}_i^{\star}$