THE MATHEMATICS OF LATTICE-BASED CRYPTOGRAPHY

5. SIS/LWE and Lattices

Alfred Menezes
cryptography101.ca

© Alfred Menezes
December 2024

Outline

1. The SIS lattice
2. Average-case hardness of SIS
3. The LWE lattice
4. Average-case hardness of LWE

 © *Alfred Menezes*

Definition of the SIS lattice

- **SIS(n, m, q, B)**. Given $A \in_R \mathbb{Z}_q^{n\times m}$ (where $m \gg n$) and $B \ll q/2$, find $z \in \mathbb{Z}_q^m$ such that $Az = 0 \pmod{q}$, where $z \neq 0$ and $z \in [-B, B]^m$.

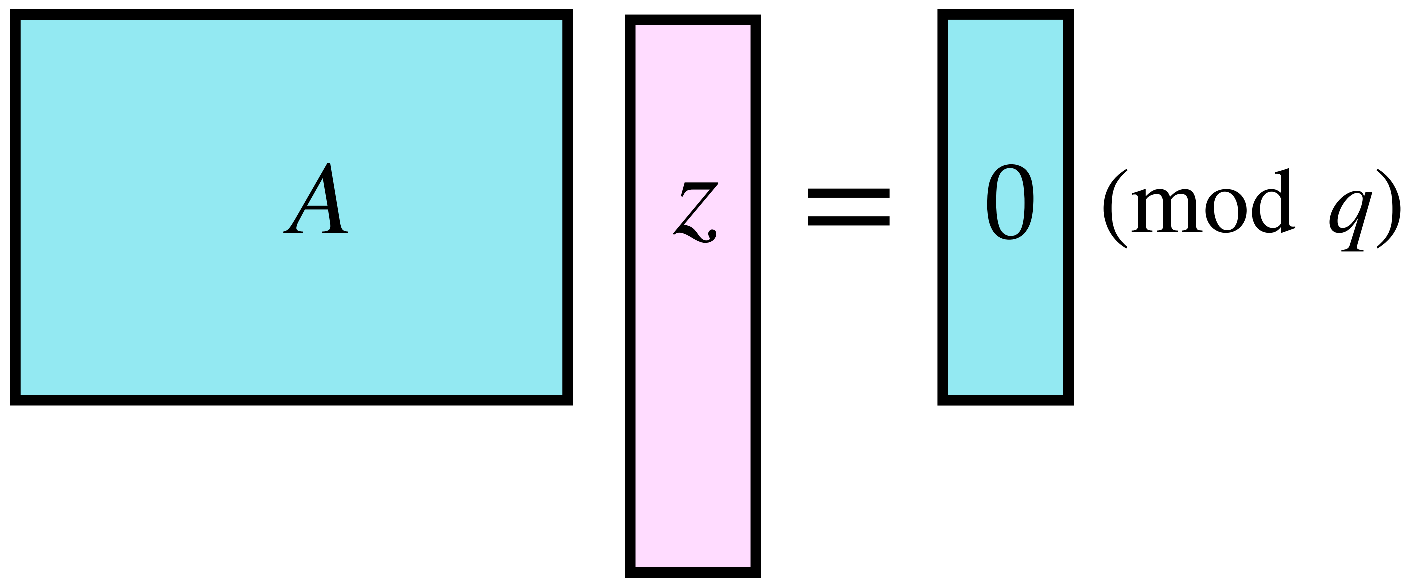

- Define the *SIS lattice* to be $L_A^{\perp} = \{z \in \mathbb{Z}^m : Az = 0 \pmod{q}\}$.
- **Claim 1.** $L_A^{\perp}$ is an integer lattice in $\mathbb{R}^m$.
- The claim can be easily proven using the following *equivalent* definition of a lattice.
- **Fact**. A lattice L is a discrete additive subgroup of $\mathbb{R}^m$.
  - L is an *additive subgroup* of $\mathbb{R}^m$ means that (i) L is non-empty subset of $\mathbb{R}^m$; and (ii) $x + y, -x \in L$ for all $x, y \in L$.
  - L is *discrete* means that for each $x \in L$, there exists $\epsilon > 0$ such that no element of L (other than x) is within distance ϵ of x.

Rank of the SIS lattice

- **Claim 2**. The SIS lattice $L_A^\perp = \{z \in \mathbb{Z}^m : Az = 0 \pmod{q}\}$ has full rank m.
- **Proof**. The lattice $q\mathbb{Z}^m$ is a sublattice of $L_A^\perp$.
  Now, the m vectors $(q,0,\ldots,0)$, $(0,q,\ldots,0)$, … , $(0,0,\ldots,q)$ are in $q\mathbb{Z}^m$ and are linearly independent (over $\mathbb{R}$).
  Thus, $q\mathbb{Z}^m$ is a full-rank lattice, and so $L_A^\perp$ is also a full-rank lattice. □
- **Notes**.
  1. $L_A^\perp$ is a q-ary lattice, i.e. for all $z \in \mathbb{Z}^m$ we have $z \in L_A^\perp$ if and only if $z \bmod q \in L_A^\perp$.
  2. A basis matrix for the lattice $q\mathbb{Z}^m$ is qI_m.
     Thus, $\mathrm{vol}(q\mathbb{Z}^m) = |\det(qI_m)| = q^m$ and hence $\mathrm{vol}(L_A^\perp) \le q^m$.

Volume of the SIS lattice

- **Claim 3**. The SIS lattice $L_A^\perp = \{z \in \mathbb{Z}^m : Az = 0 \pmod{q}\}$ has volume q^n (assuming that A has rank n over $\mathbb{Z}_q$.)

See Section 1.6 of *Algebraic Number Theory and Fermat's Last Theorem* (3rd edition), by Stewart and Tall.

- **Proof**. $\mathbb{Z}^m$ and $L_A^\perp$ are free (additive) abelian groups of rank m.
- Since $L_A^\perp$ is a subgroup of $\mathbb{Z}^m$, and they have the same rank, the quotient group $\mathbb{Z}^m/L_A^\perp$ is finite. Moreover, $\text{vol}(L_A^\perp) = |\mathbb{Z}^m/L_A^\perp|$. (This is Theorem 1.17 in Stewart & Tall's book.)
- So, to determine $\text{vol}(L_A^\perp)$, we need to compute $|\mathbb{Z}^m/L_A^\perp|$, the number of cosets of $L_A^\perp$ in $\mathbb{Z}^m$.
  - Now, let $x, y \in \mathbb{Z}^m$. Then $L_A^\perp + x = L_A^\perp + y \iff x - y \in L_A^\perp \iff A(x-y) = 0 \pmod{q} \iff Ax = Ay \pmod{q}$.
  - Assuming that A has rank n over $\mathbb{Z}_q$, its column space has dimension n over $\mathbb{Z}_q$.
  - Thus, the column space of A has size q^n, whence $|\mathbb{Z}^m/L_A^\perp| = q^n$. □

A basis of the SIS lattice

- **Claim 4**. Suppose that the first n columns of A are linearly independent over $\mathbb{Z}_q$, so A can be row-reduced to a matrix $\tilde{A} = [I_n \,|\, \overline{A}]$ (where $\overline{A} \in \mathbb{Z}_q^{n\times(m-n)}$).
Then $C = \begin{bmatrix} qI_n & -\overline{A} \\ 0 & I_{m-n} \end{bmatrix} \in \mathbb{Z}^{m\times m}$ is a basis matrix for the SIS lattice $L_A^{\perp}$.

- **Proof**. Since A and $\tilde{A}$ are row equivalent (over $\mathbb{Z}_q$), they have the same null space (mod q). Hence, $L_{\tilde{A}}^{\perp} = L_A^{\perp}$, so we will find a basis for $L_{\tilde{A}}^{\perp}$.
Now, each column v of C is in $L_{\tilde{A}}^{\perp}$ since $\tilde{A}v = 0 \pmod{q}$ [check this!].
Moreover, the columns of C are linearly independent over $\mathbb{R}$ since $\det(C) = q^n$.
Thus, C is a basis matrix for a full-rank sublattice L of $L_{\tilde{A}}^{\perp}$.
Since $\text{vol}(L) = q^n = \text{vol}(L_A^{\perp}) = \text{vol}(L_{\tilde{A}}^{\perp})$, we have $L_{\tilde{A}}^{\perp} = L$.
Thus, C is a basis matrix for the SIS lattice $L_A^{\perp}$. $\square$

Solving SIS

- **SIS(n, m, q, B)**. Given $A \in_R \mathbb{Z}_q^{n\times m}$ find $z \in \mathbb{Z}_q^m$ such that $Az = 0 \pmod q$, where $z \neq 0$ and $z \in [-B, B]^m$.
- An equivalent **lattice formulation** is:
  SIS(n, m, q, B): Given $A \in_R \mathbb{Z}_q^{n\times m}$, find a nonzero $z \in [-B, B]^m$ in the SIS lattice $L_A^\perp = L(C)$ where $C = \begin{bmatrix} qI_n & -\overline{A} \\ 0 & I_{m-n} \end{bmatrix}$.
- For $z \in \mathbb{R}^m$, the infinity norm of z is $\|z\|_\infty = \max_i |z_i|$.
  - So, an SIS solution $z \in \mathbb{Z}^m$ must satisfy $0 < \|z\|_\infty \leq B$.
- SIS hardness is usually studied using the Euclidean norm: $\|z\|_2 = \sqrt{z_1^2 + z_2^2 + \cdots + z_m^2}$.
- **Exercise**: Show that for all $z \in \mathbb{R}^m$, $\|z\|_\infty \leq \|z\|_2 \leq \sqrt{m}\,\|z\|_\infty$.

Solving SIS_2

- **$\mathrm{SIS}_2(n, m, q, \beta)$**. Given $A \in_R \mathbb{Z}_q^{n\times m}$ where $\beta \ll q$, find nonzero $z \in \mathbb{Z}_q^m$ such that $Az = 0 \pmod q$ and $\|z\|_2 \leq \beta$.
- An equivalent lattice formulation is:
SIS_2 (n, m, q, β): Given $A \in_R \mathbb{Z}_q^{n\times m}$, find nonzero z with $\|z\|_2 \leq \beta$ in the SIS lattice $L_A^\perp$.
    - By Minkowski's Theorem (slide 49), $\lambda_1(L_A^\perp) \leq \sqrt{m}\, q^{n/m}$.
    - We'll assume that $\beta \geq \sqrt{m}\, q^{n/m}$, whereby an SIS_2 solution is guaranteed to exist.
- Now, by the Gaussian heuristic (slide 49), $\lambda_1(L_A^\perp) \approx \sqrt{m/(2\pi e)}\, q^{n/m}$.
- Thus, SIS_2 can be seen as an instance of approximate-SVP (SVP_γ) in the SIS lattice $L_A^\perp$ with approximation factor $\gamma = \beta\sqrt{2\pi e}/(\sqrt{m}q^{n/m})$.
- **Exercise**: Show that $\mathrm{SIS}(n, m, q, B) \leq \mathrm{SIS}_2(n, m, q, B) \leq \mathrm{SIS}(n, m, q, B/\sqrt{m})$.

 © *Alfred Menezes*

Average-case hardness of SIS

- It's reasonable to conjecture that SIS is hard in the *worst case*.
- But, what can we say about the hardness of SIS *on average*?
- In 1996, Ajtai proved a striking *average case hardness result* for SIS:
  - If SIVP_γ is hard in the *worst-case*, then SIS is hard on *average*.
  - Such a reduction is called a *worst-case to average-case reduction.*
- Since the assumption that SIVP_γ is hard in the worst case is a reasonable assumption, we have a provable guarantee that SIS is hard on average.

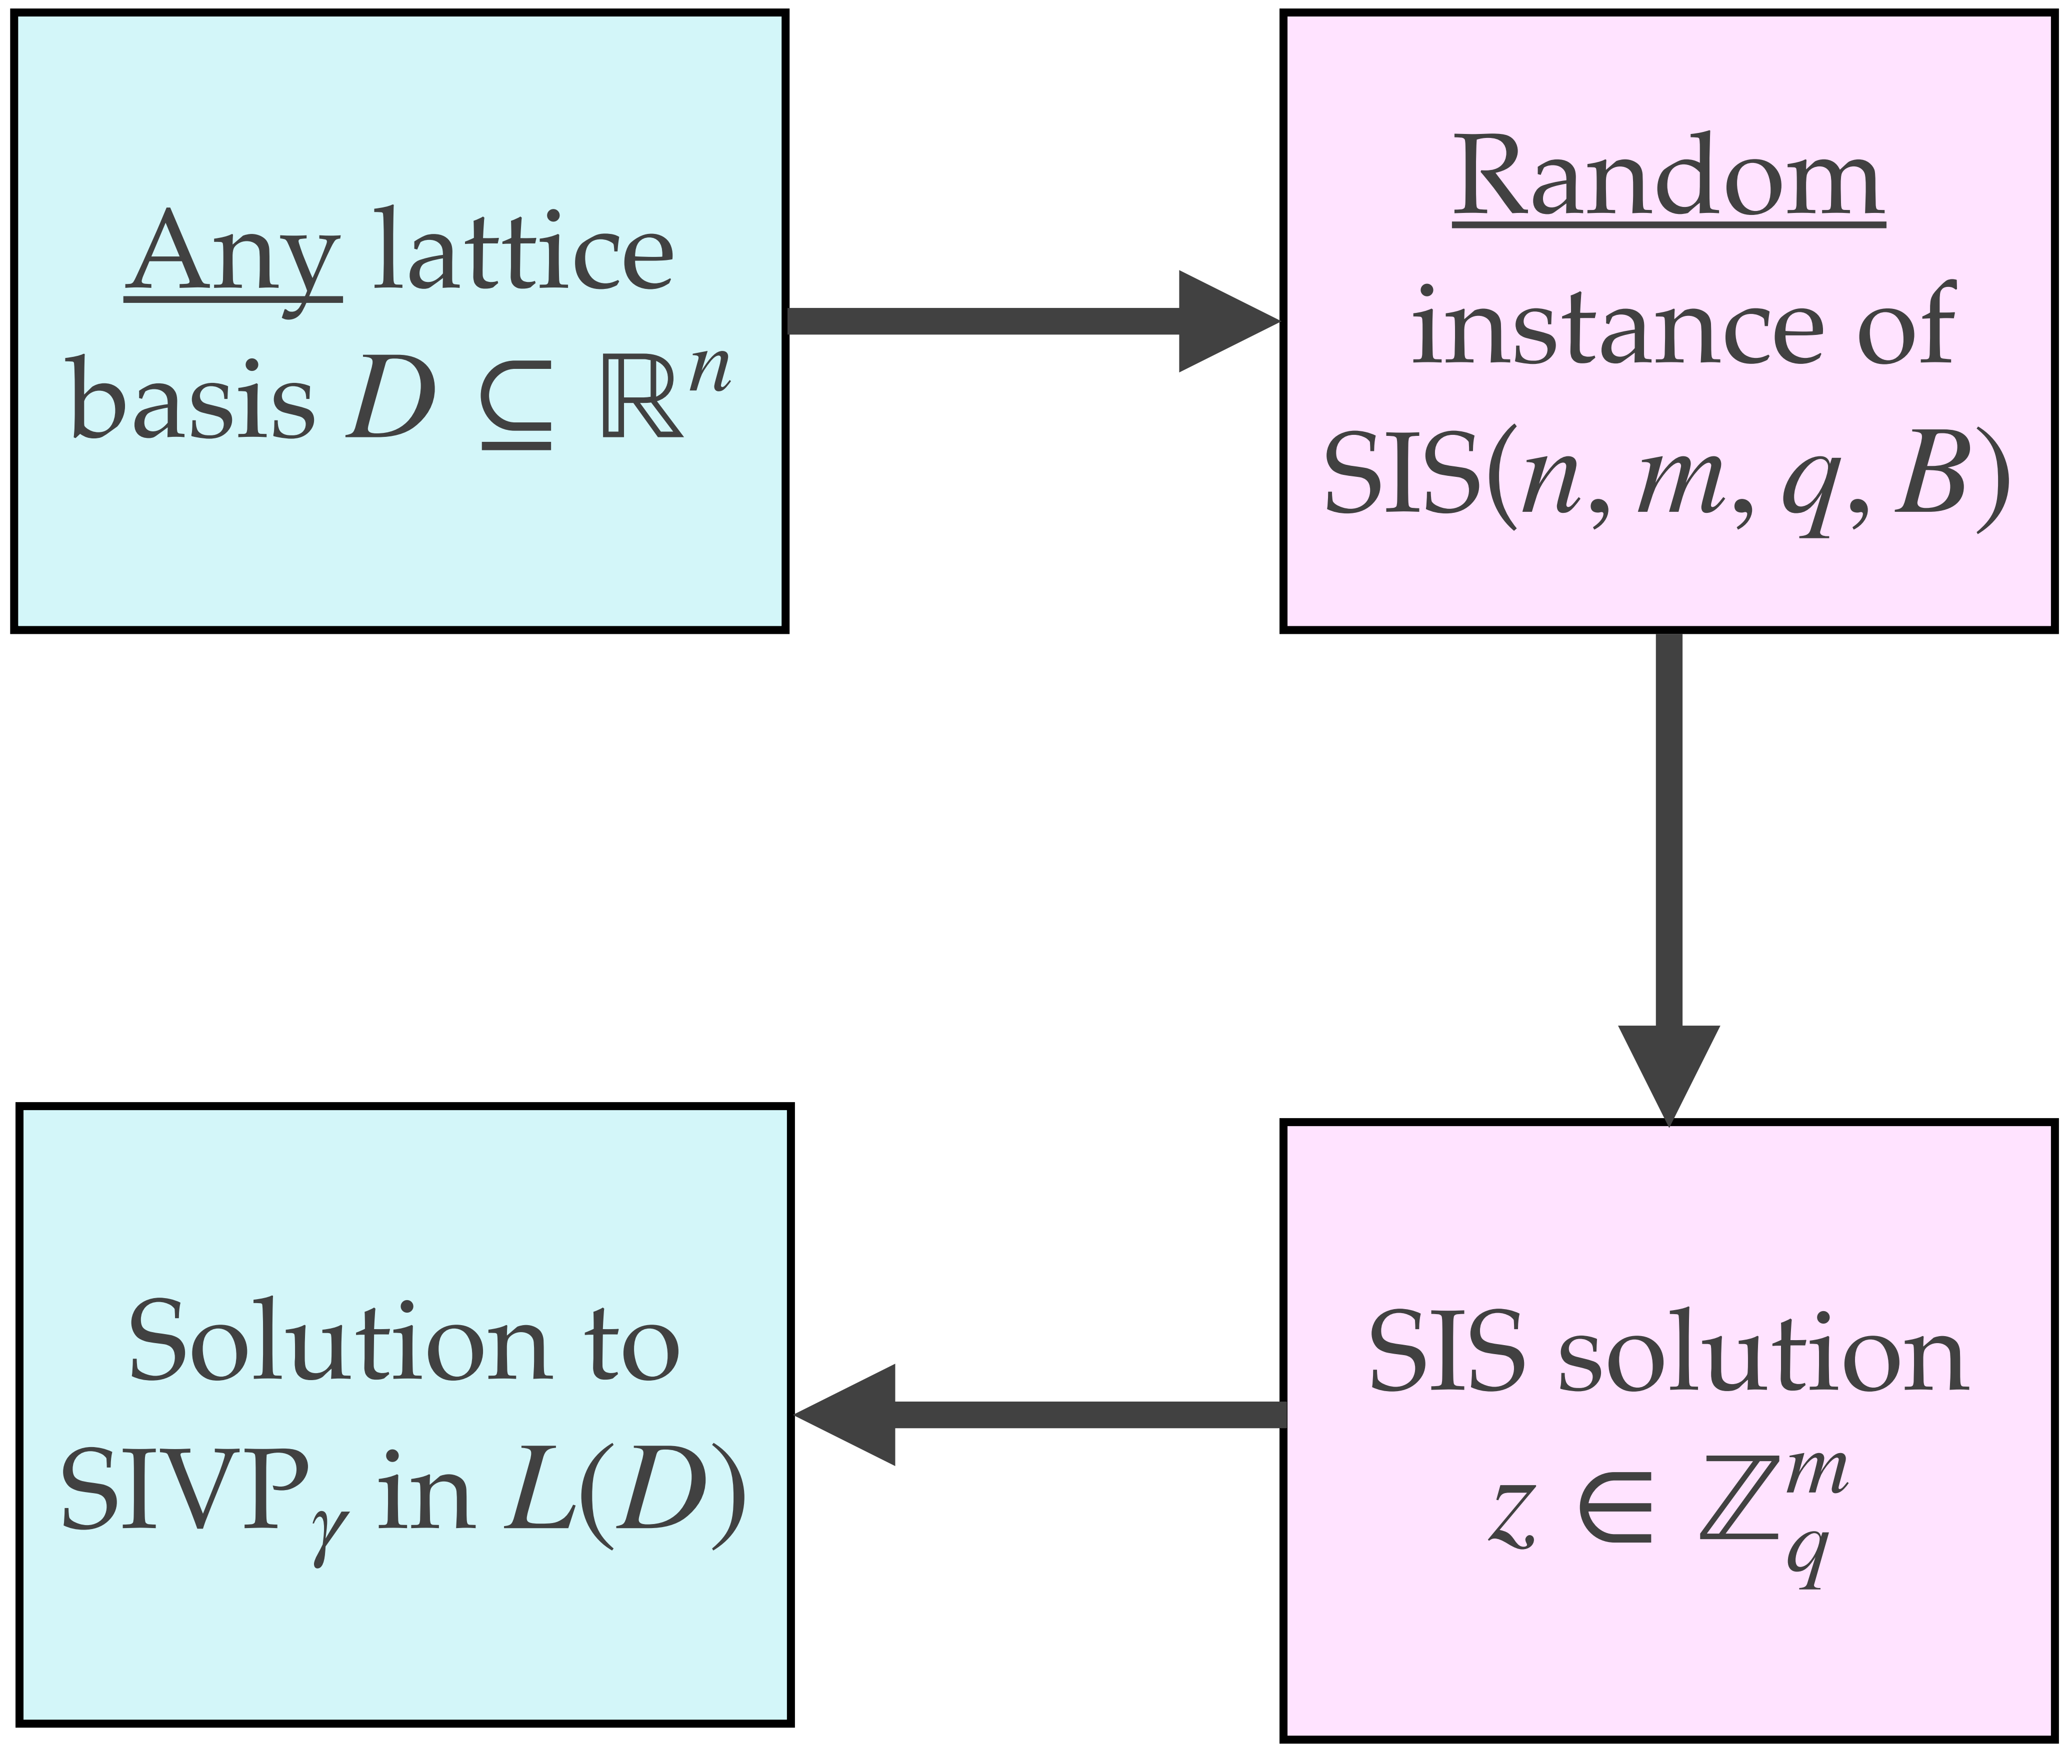

The worst-case to average-case reduction is asymptotic

- Although Ajtai's worst-case to average-case reduction provides a strong guarantee for the average-case hardness of SIS, the guarantee is an *asymptotic* one.
  - Also, the reduction is *highly non-tight*.
- In 2004, Micciancio & Regev proved the following:
  Theorem. For any $m(n) = \Theta(n \log n)$, there exists a $q(n) = O(n^2 \log n)$ such that for any function $\gamma(n) = \omega(n \log n)$, solving $\text{SIS}_2(n, m, q, \beta)$ on average with non-negligible probability is at least as hard as solving SIVP_γ in the worst case.

WORST-CASE TO AVERAGE-CASE REDUCTIONS BASED ON GAUSSIAN MEASURES*

DANIELE MICCIANCIO† AND ODED REGEV‡

SIS summary

SIS is considered a lattice problem for two reasons.

1. SIS is equivalent to solving SVP_γ in the SIS lattice.
   - The fastest algorithm known for solving SVP_γ is the Block-Korkine-Zolotarev (BKZ) algorithm, which has an exponential running time.
   - The running time of BKZ is used to select concrete parameters for SIS for a desired security level.
2. Solving SIS on average is provably at least as hard as solving SIVP_γ in the worst case.
   - This hardness guarantee is an asymptotic one, and its relevance to the hardness SIS in practice is not clear.

Definition of the LWE lattice

- **LWE(m, n, q, B)**. Let $s \in_R \mathbb{Z}_q^n$ and $e \in_R [-B, B]^m$. Given $A \in_R \mathbb{Z}_q^{m\times n}$ and $b = As + e \pmod{q}$, find s.

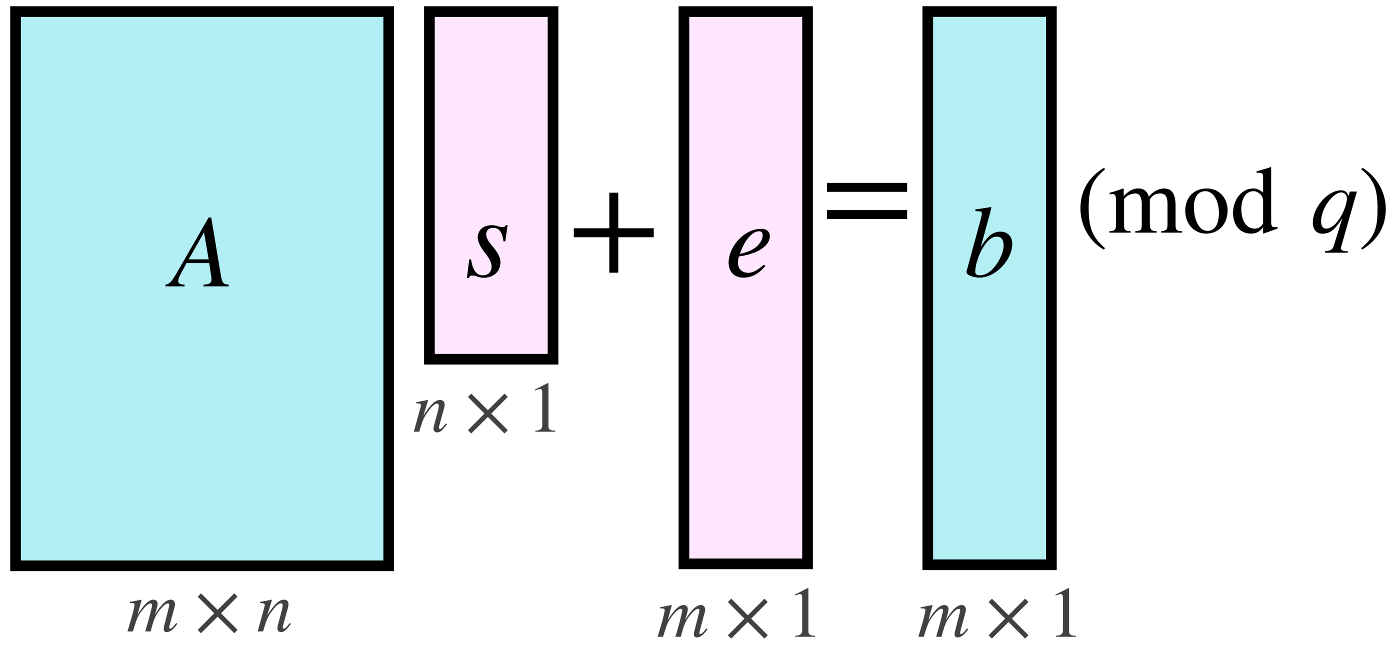

- Define the *LWE lattice* to be $L_A = \{y \in \mathbb{Z}^m : Az = y \pmod{q} \text{ for some } z \in \mathbb{Z}^n\} \subseteq \mathbb{R}^m$.
- **Claim 1.** L_A is a full-rank (integer) q-ary lattice in $\mathbb{R}^m$.
- **Proof**. L_A is a discrete additive subgroup of $\mathbb{R}^m$, and thus is a lattice. Since $q\mathbb{Z}^m$ is a sublattice of L_A, it follows that L_A is q-ary and has rank m. □

A basis of the LWE lattice

- **Claim 2.** Let $A = \begin{bmatrix} A_1 \\ A_2 \end{bmatrix}$ where $A_1 \in \mathbb{Z}_q^{n\times n}$ and $A_2 \in \mathbb{Z}_q^{(m-n)\times n}$, and suppose that A_1 is invertible mod q. Let $D_2 = A_2 A_1^{-1} \pmod{q}$. Then $D = \begin{bmatrix} I_n & 0 \\ D_2 & qI_{m-n} \end{bmatrix} \in \mathbb{Z}^{m\times m}$ is a basis matrix for L_A (and so $\mathrm{vol}(L_A) = q^{m-n}$).
- **Proof**. Since $\det(D) = q^{m-n}$, the columns of D are linearly independent over $\mathbb{R}$.
  Write $y \in \mathbb{Z}^m$ as $\begin{bmatrix} y_1 \\ y_2 \end{bmatrix}$ where $y_1 \in \mathbb{Z}^n$ and $y_2 \in \mathbb{Z}^{m-n}$.
  Now, $y \in L_A \iff y = Az \pmod{q}$ for some $z \in \mathbb{Z}^n \iff y_1 = A_1 z \pmod{q}$ and $y_2 = A_2 z \pmod{q}$ for some $z \in \mathbb{Z}^n \iff y_2 = A_2 A_1^{-1} y_1 \pmod{q} \iff y_2 = D_2 y_1 + qc$ for some $c \in \mathbb{Z}^{m-n}$.
  Observing that $y = D \begin{bmatrix} y_1 \\ c \end{bmatrix}$, it follows that the columns of D are a basis for L_A. $\square$

Solving LWE

- **LWE(m, n, q, B)**. Let $s \in_R \mathbb{Z}_q^n$ and $e \in_R [-B, B]^m$. Given $A \in_R \mathbb{Z}_q^{m \times n}$ and $b = As + e \pmod{q}$, find s.
- **LWE lattice**:
  $L_A = \{y \in \mathbb{Z}^m : As = y \pmod{q} \text{ for some } s \in \mathbb{Z}^n\} \subseteq \mathbb{R}^m$.
- Note that for an LWE instance (A, b, s, e), we have $y = As \bmod q \in L_A$, and $\|b - y\|_2 = \|e\|_2 \leq \sqrt{m}\, B$.
- Thus, LWE is a special instance of the following lattice problem:
  Bounded Distance Decoding (BDD$_\alpha$):
  Given a lattice $L = L(D) \subseteq \mathbb{R}^m$ and $b \in \mathbb{R}^m$ with the guarantee that there is a unique $y \in L$ within distance α of b, find y.

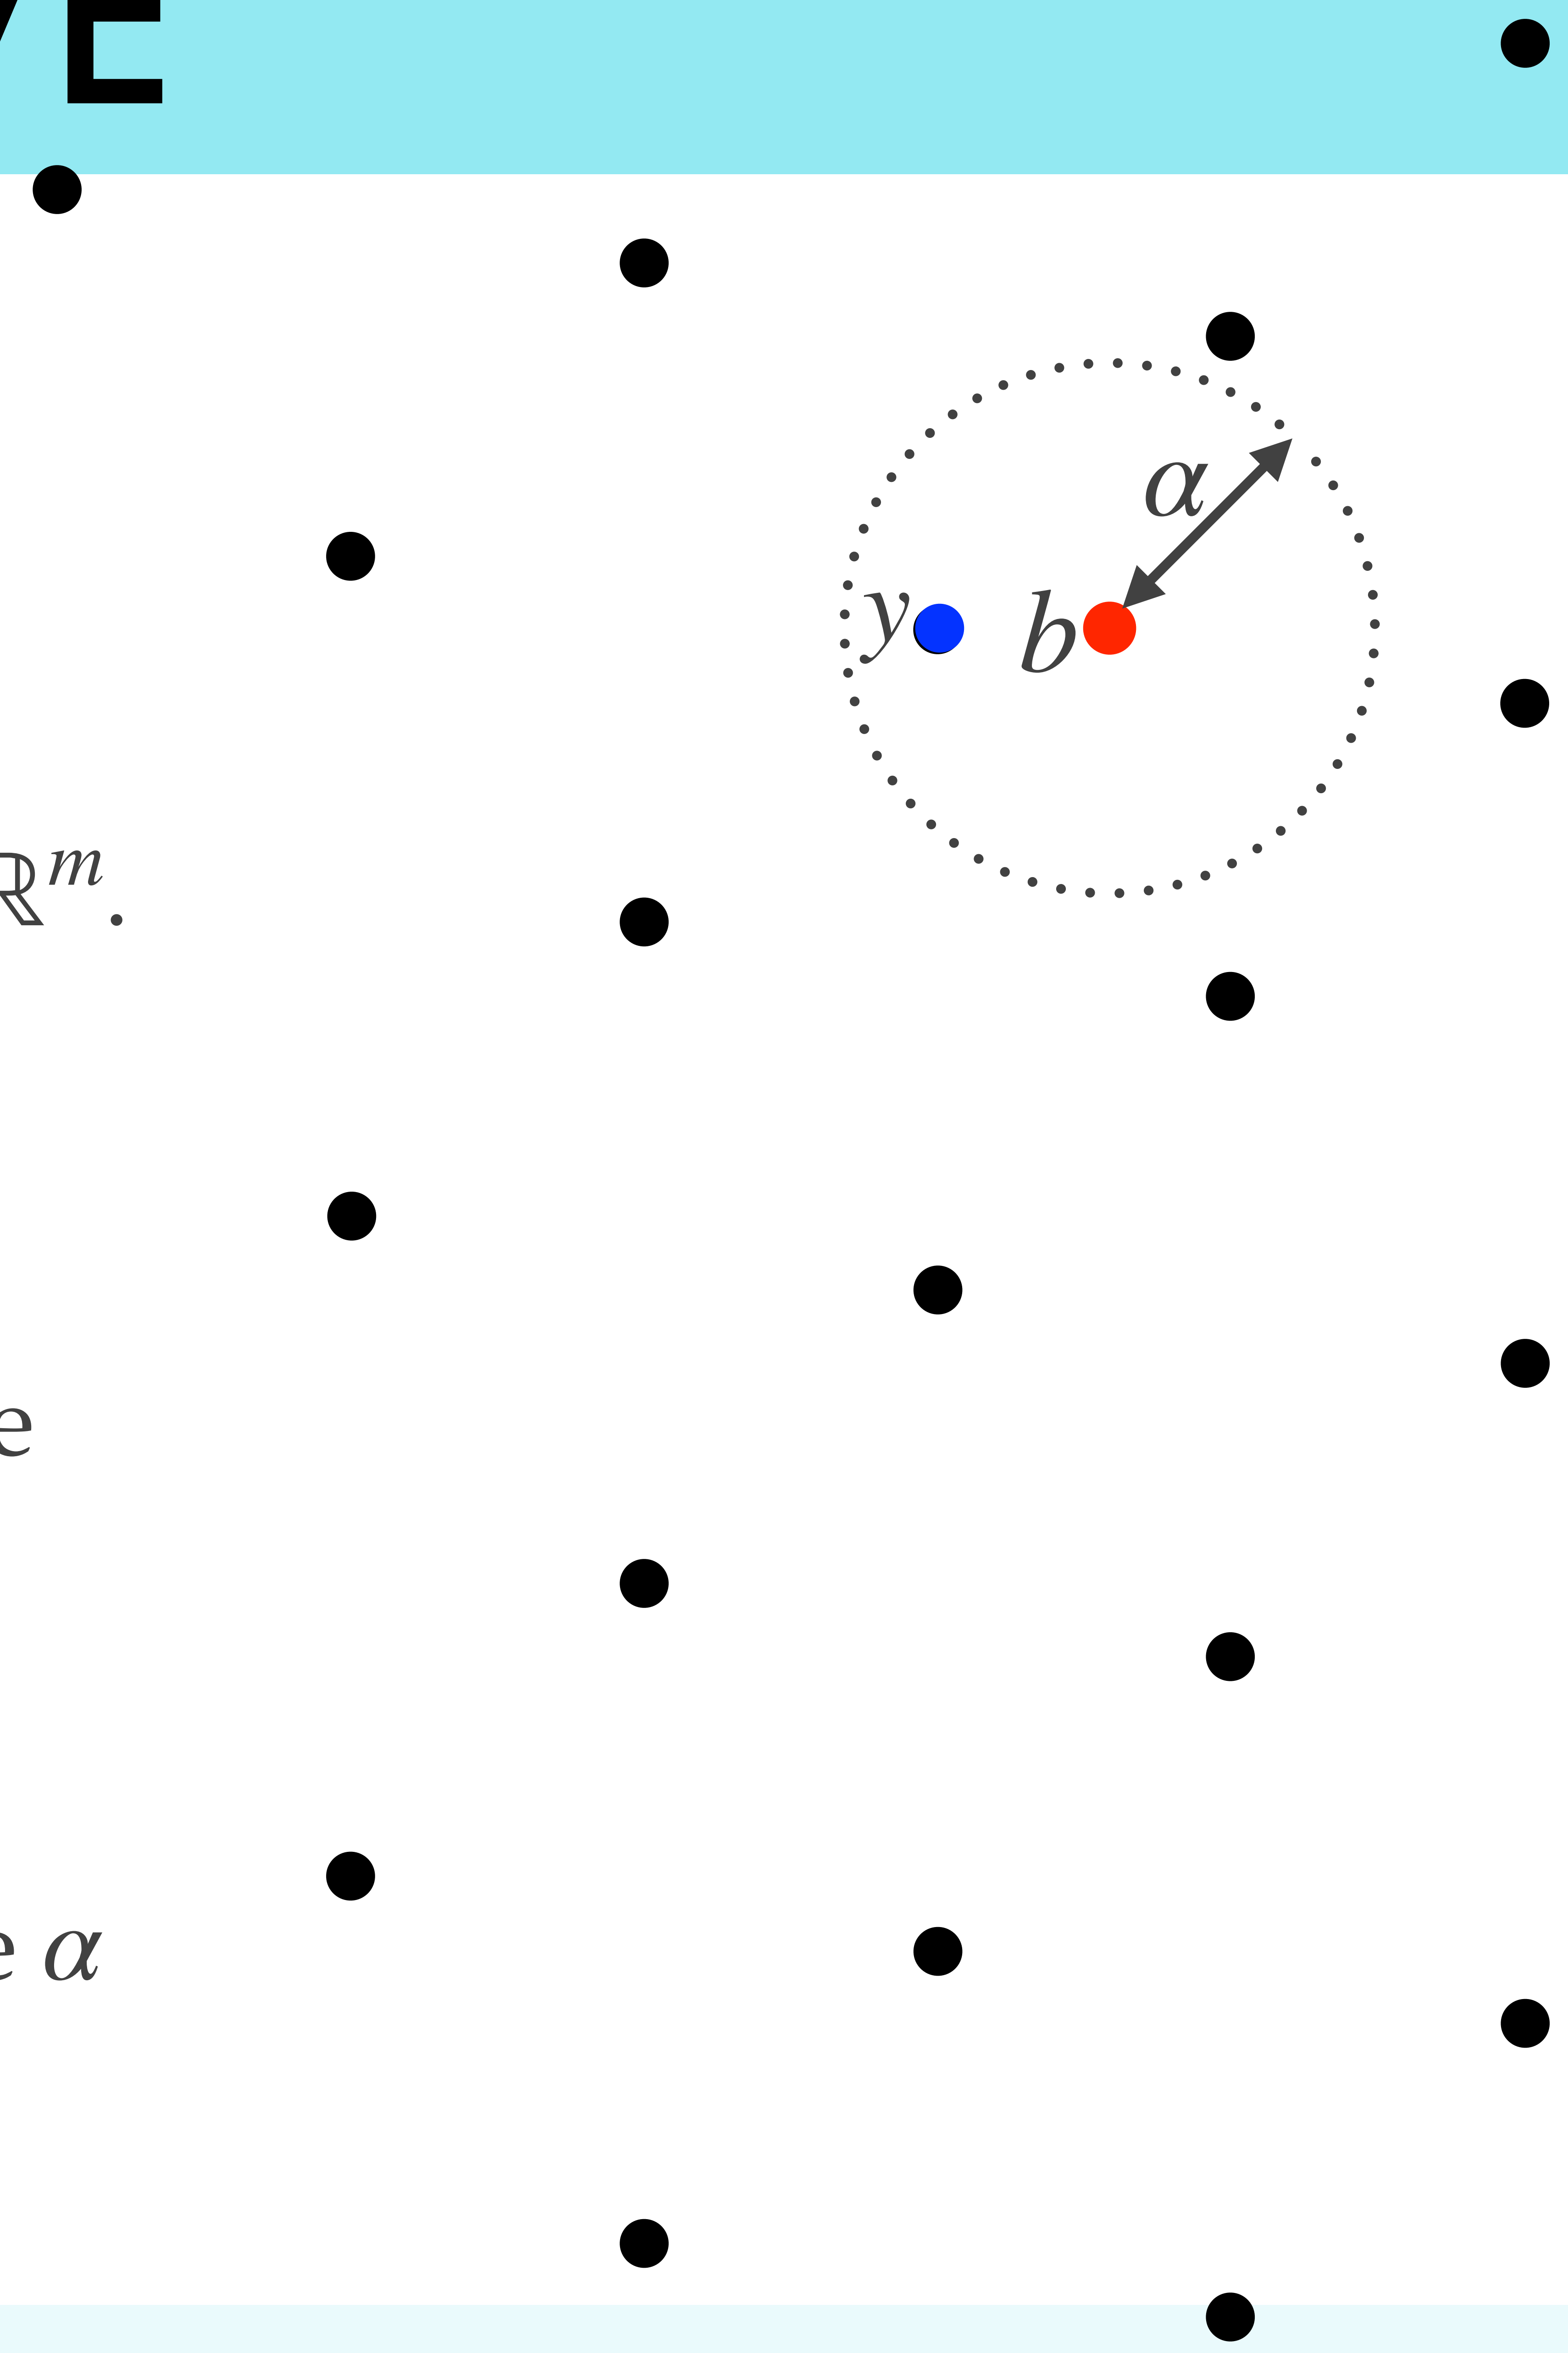

Reducing BDD to SVP (1)

- **BDD**$_\alpha$: Given a lattice $L = L(D) \subseteq \mathbb{R}^m$ and $b \in \mathbb{R}^m$ with the guarantee that there is a unique $y \in L$ within distance α of b, find y.
- We'll suppose that $\alpha < \lambda_1(L)/\sqrt{2}$.
- Let $D' = \begin{bmatrix} D & -b \\ 0 & \alpha \end{bmatrix} \in \mathbb{R}^{(m+1)\times(m+1)}$. Then $L' = L(D') = \left\{ \begin{bmatrix} v - cb \\ c\alpha \end{bmatrix} : v \in L(D) \text{ and } c \in \mathbb{Z} \right\}$.
- Notice that for $(v, c) = (y, 1)$, we have $\tilde{v} = \begin{bmatrix} y - b \\ \alpha \end{bmatrix} \in L'$ with $\|\tilde{v}\|_2 = \sqrt{\|y - b\|_2^2 + \alpha^2} \le \sqrt{2}\alpha$.
  Hence, $\lambda_1(L') \le \sqrt{2}\alpha < \lambda_1(L)$.
- Suppose now that $v' = \begin{bmatrix} v - cb \\ c\alpha \end{bmatrix} \in L'$ has length $\|v'\|_2 = \lambda_1(L')$.
- If $c = 0$, then $\|v'\|_2 = \|v\|_2 \ge \lambda_1(L) > \lambda_1(L')$, a contradiction.
- And, if $|c| \ge 2$, then $\|v'\|_2 \ge 2\alpha > \sqrt{2}\alpha \ge \lambda_1(L')$, a contradiction.
- Hence, we must have $c = \pm 1$.
  If $c = 1$, we have $v' = \begin{bmatrix} v - b \\ \alpha \end{bmatrix}$ for some $v \in L$.
- Now, if $v \neq y$, then $\|v - b\|_2 > \|y - b\|_2$, whence $\|v'\|_2 > \|\tilde{v}\|_2$, contradicting $\|v'\|_2 = \lambda_1(L')$.
- Hence $\pm\tilde{v}$ are the only vectors of length $\lambda_1(L')$ in L'.

 © *Alfred Menezes*

Reducing BDD to SVP (2)

- **BDD**$_\alpha$: Given a lattice $L = L(D) \subseteq \mathbb{R}^m$ and $b \in \mathbb{R}^m$ with the guarantee that there is a unique $y \in L$ within distance α of b, find y.
- **Summary**: We can solve the BBD$_\alpha$ instance by solving SVP for $L(D')$ where $D' = \begin{bmatrix} D & -b \\ 0 & \alpha \end{bmatrix}$.
- This method of solving LWE is called a "primal attack using a Kannan embedding".

Average-case hardness of LWE

- It's reasonable to conjecture that LWE is hard in the *worst case*.
- But, what can we say about the hardness of LWE *on average*?
- In 2005, Regev proved a striking *average-case hardness result* for LWE:
    - If SIVP_γ is quantumly hard in the *worst-case*, then LWE is hard on *average*.
- Since the assumption that SIVP_γ is quantumly hard in the worst case is a reasonable assumption, we have a provable guarantee that LWE is hard on average.
- However, as with Ajtai's worst-case to average-case reduction for SIS, Regev's reduction is *highly non-tight* (and also a quantum reduction).
    - For a concrete analysis of Regev's reduction, see Section 5 of:
      "Another look at tightness II: practical issues in cryptography"
      by Chatterjee, Koblitz, Menezes & Sarkar, https://eprint.iacr.org/2016/360.

Gaussian distributions

- I should note that in Regev's worst-case to average-case reduction, and also in much of the cryptographic literature on LWE-based protocols, the components of the LWE error vector e are drawn from certain Gaussian distributions (and not from uniform distributions)
- However, for the sake of simplicity, I didn't use Gaussians in my lectures.
- Also, Kyber and Dilithium use uniform distributions and central binomial distributions.

LWE summary

LWE is considered a lattice problem for two reasons.

1. LWE can be reduced to solving BDD_α in the LWE lattice, which in turn can be reduced to solving an instance of SVP.

   - The fastest algorithm known for solving SVP is the Block-Korkine-Zolotarev (BKZ) algorithm, which has an exponential running time.
   - The running time of BKZ can be used to select concrete parameters for LWE for a desired security level.

2. Solving LWE on average is provably at least as hard as (quantumly) solving SIVP_γ in the worst case.

   - This hardness guarantee is an asymptotic one, and its relevance to LWE in practice is not clear.