CS 258: Quantum Cryptography

Mark Zhandry

Midterm Logistics

Available on Gradescope from 10/25 – 10/28

Any 2 hour increment

Completed individually (including no AI)

Handwritten ok. Open computer, notes, internet

Material through group actions + Kuperberg (no lattices)

No Class Monday 10/27!

Next Class: Wednesday 10/29

Previously…

Lattices

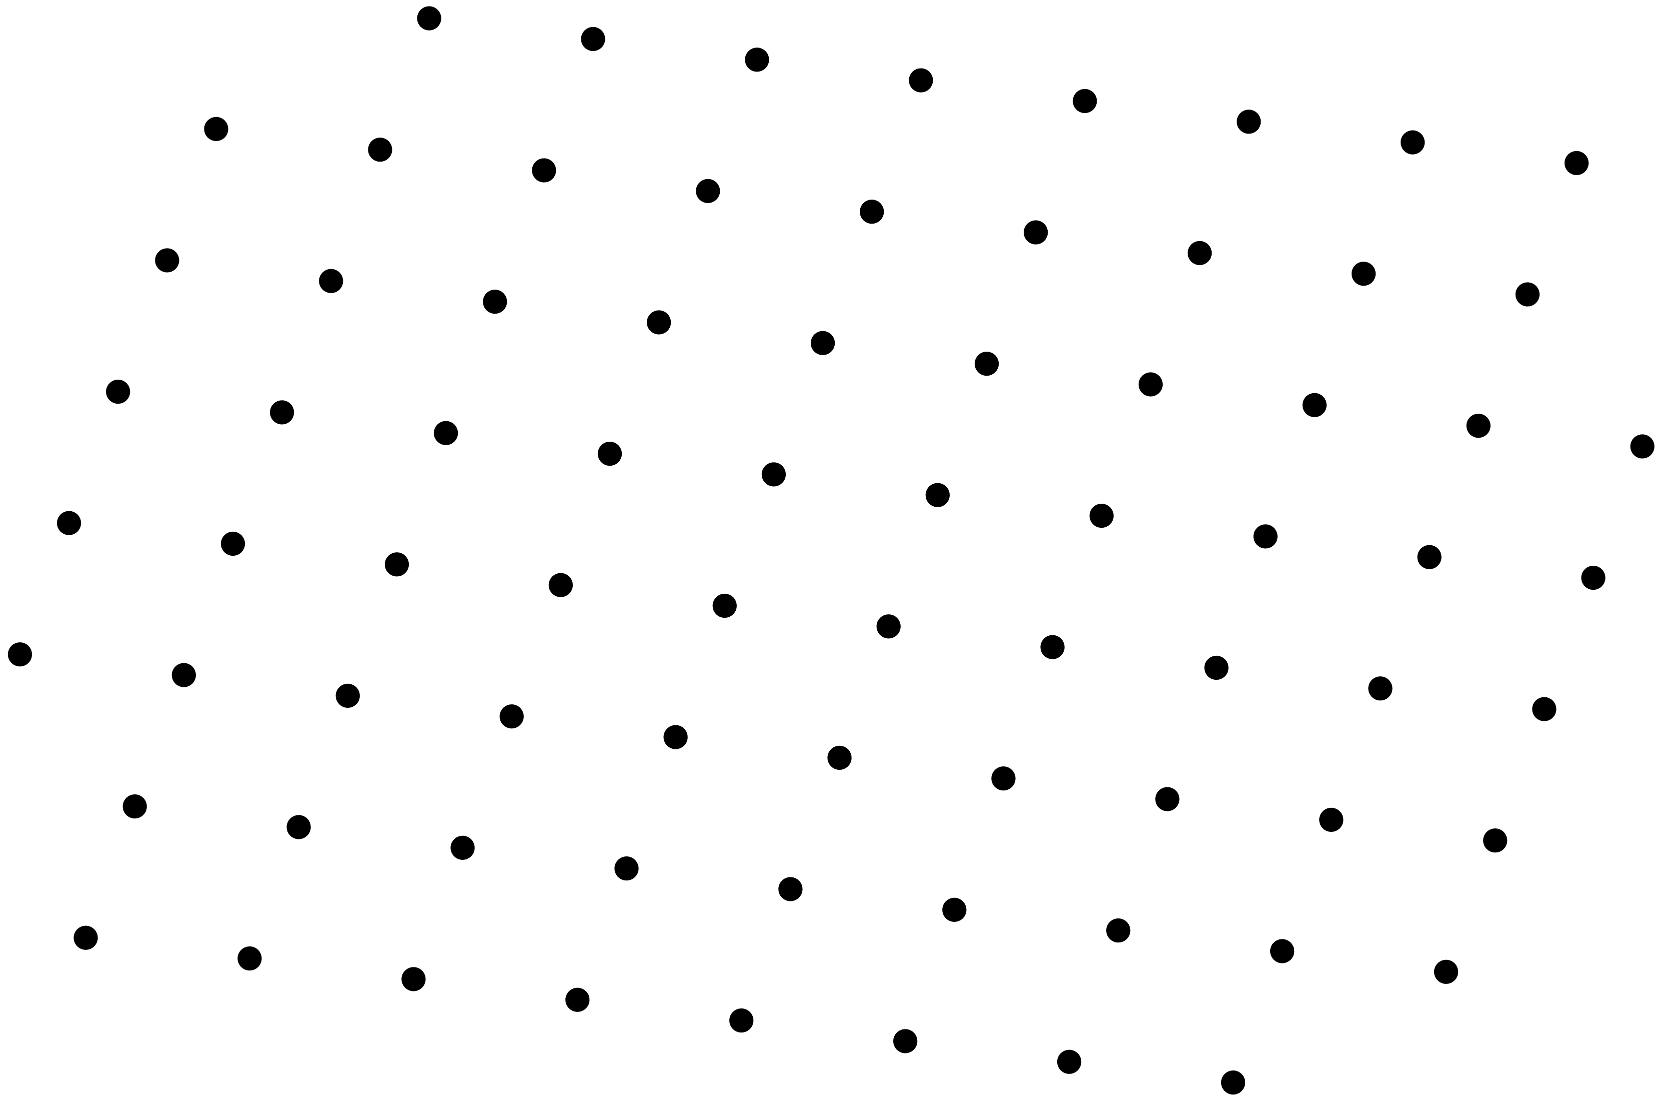

Imagine dimension in the 100s

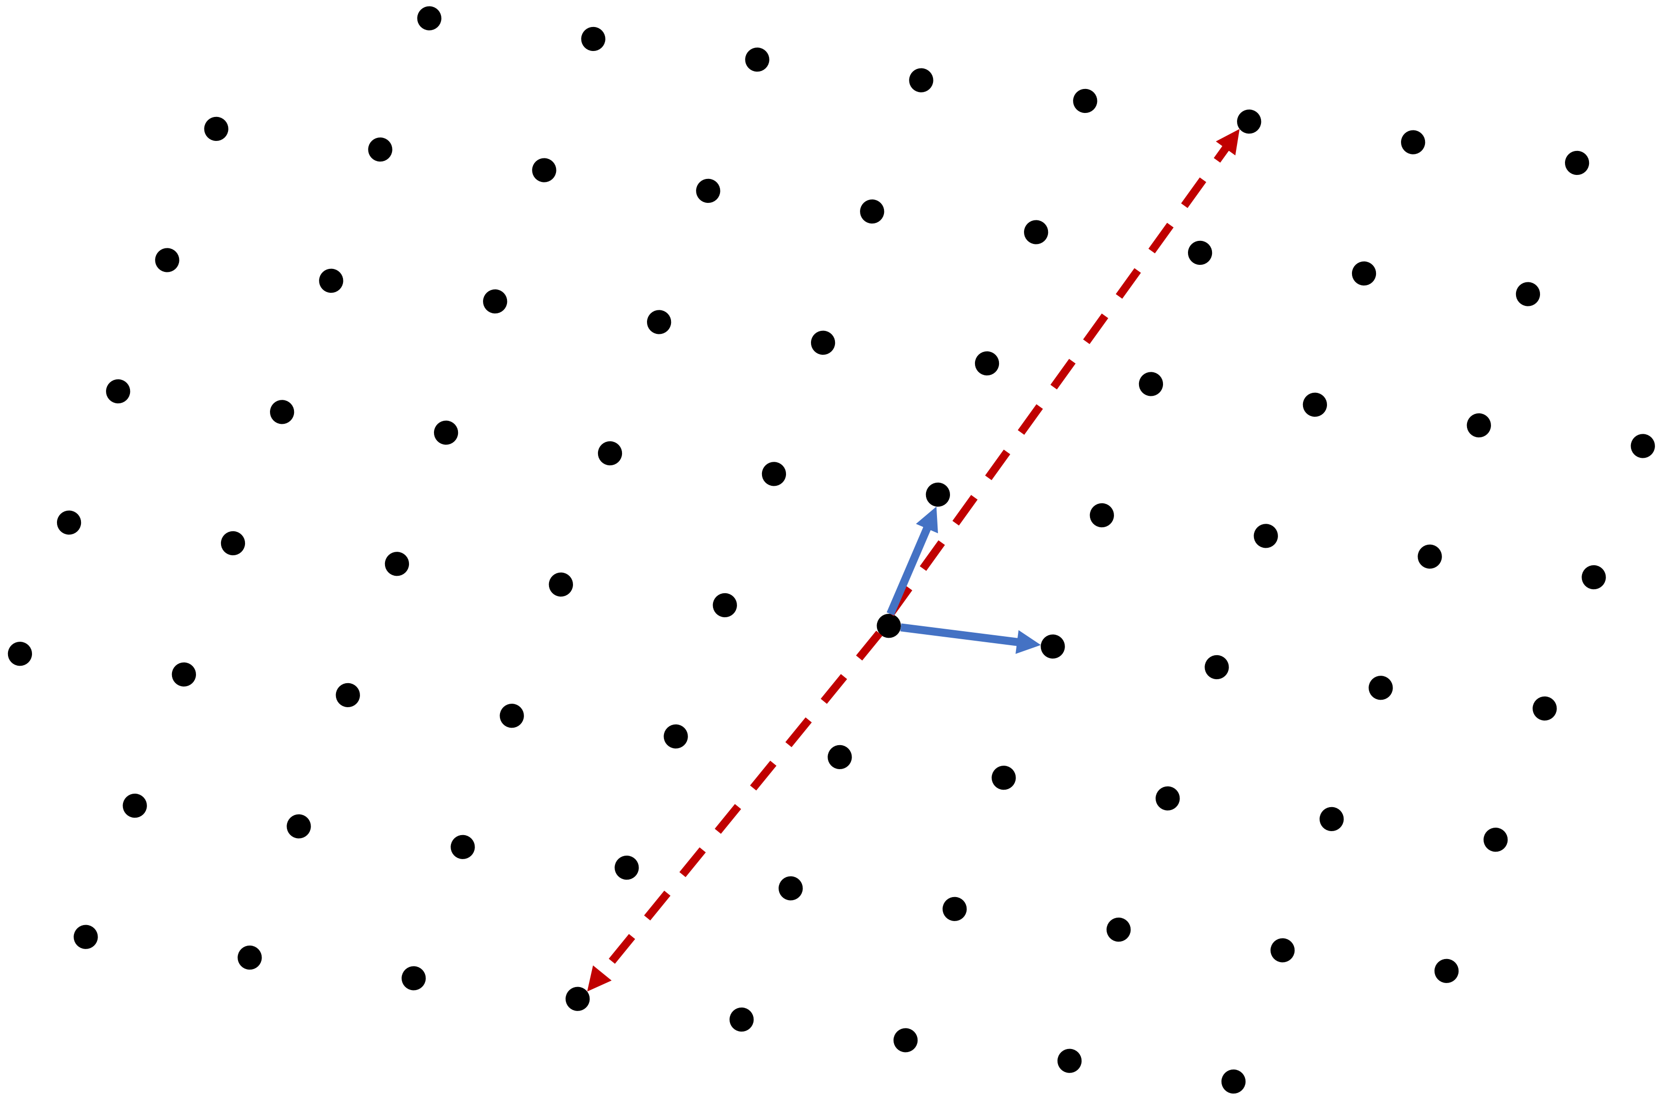

Different Bases

Different Bases

For vector spaces: two bases $\mathbf{B}_1, \mathbf{B}_2$ generate the same vector space if and only if there is an invertible $\mathbf{U}$ such that $\mathbf{B_2} = \mathbf{B}_1 \cdot \mathbf{U}$

For lattices: two bases $\mathbf{B}_1, \mathbf{B}_2$ generate the same lattice if and only if there is a **unimodular** $\mathbf{U}$ such that $\mathbf{B_2} = \mathbf{B}_1 \cdot \mathbf{U}$

Def: $\mathbf{U}$ is unimodular if $\mathbf{U} \in \mathbb{Z}^{n \times n}$ and $\det(\mathbf{U}) \in \{+1, -1\}$

Determinant of lattice

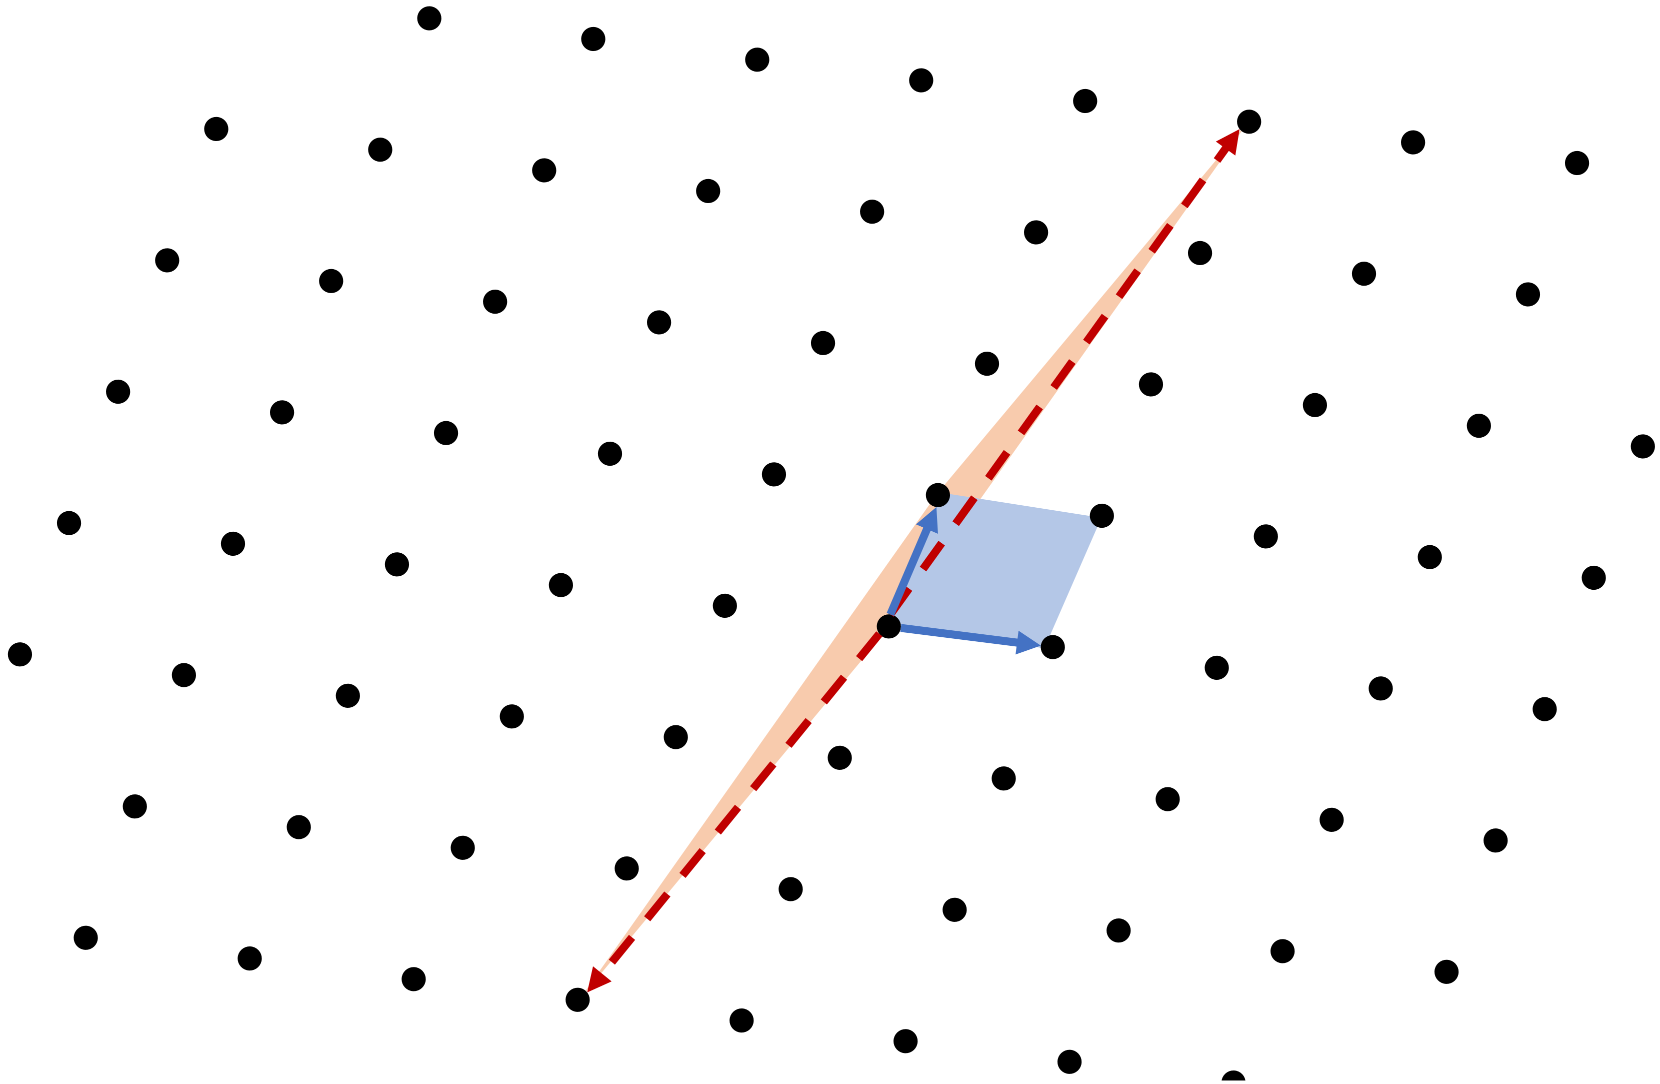

Measure of how dense the lattice is

Full-rank lattice: $\mathsf{span}(\mathbf{B}) = \mathbb{R}^n \iff \mathbf{B} \in \mathbb{R}^{n \times n}$

Integer lattice: $\mathbf{B} \in \mathbb{Z}^{m \times n}$

We will generally consider only full-rank integer lattices

Note that for integer lattices, can consider spanning set that is not full-rank, and still guarantee discreteness

SVP

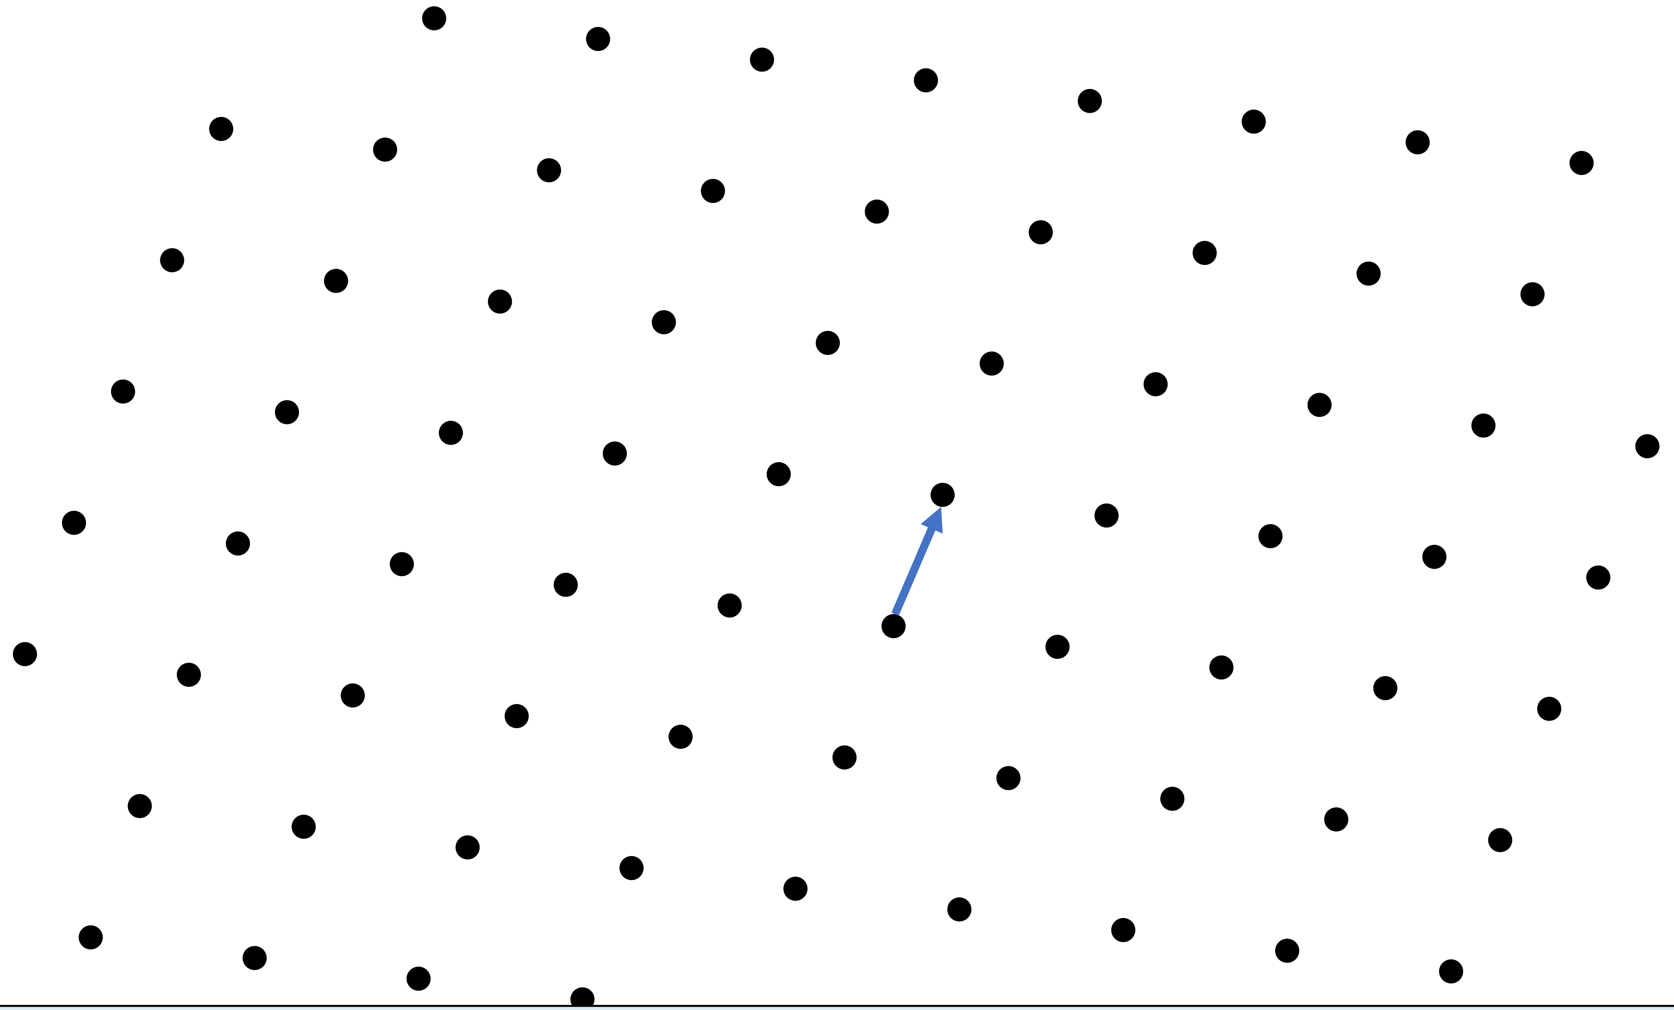

(Approx.) shortest vector problem (SVP): given lattice (described by some basis), find (approx.) shortest vector

SIVP

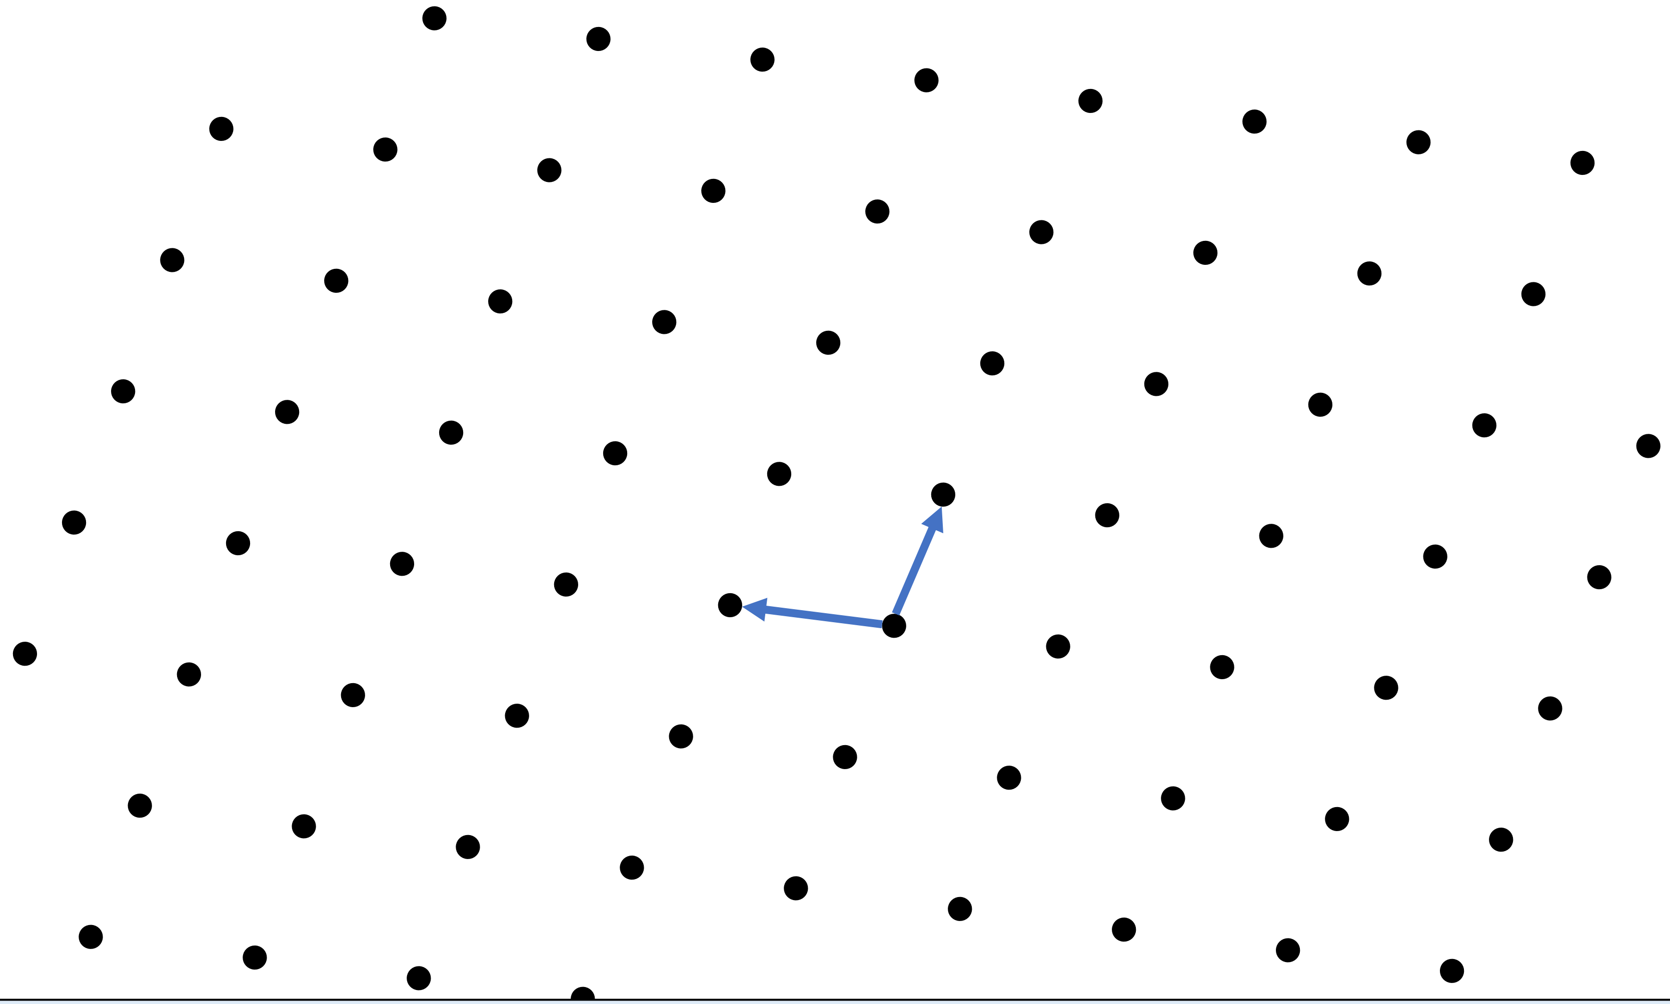

(Approx.) shortest independent vector problem (SIVP): given lattice (described by some basis), find (approx.) shortest basis

CVP

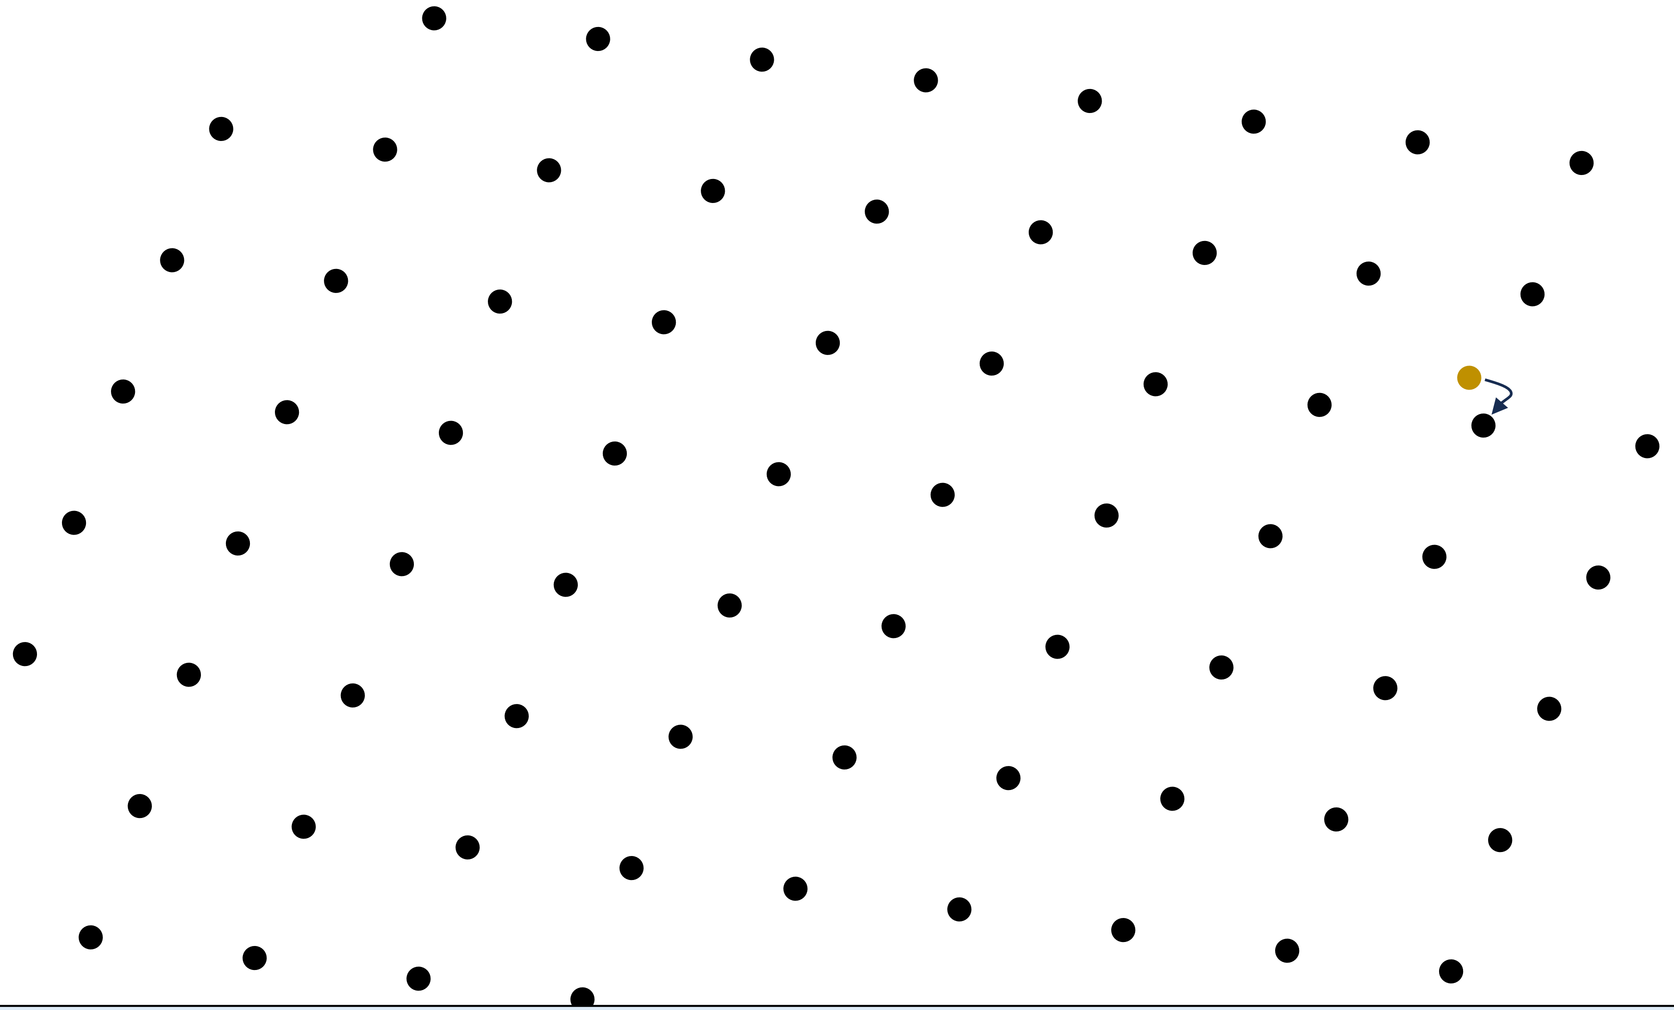

(Approx.) closest vector problem (CVP): given lattice and point off lattice, find (approx.) closest lattice point

Gram-Schmidt Orthogonalization

(no normalization)

$$\mathbf{B} = \left(\; \mathbf{b}_1 \;\middle|\; \mathbf{b}_2 \;\middle|\; \cdots \; \right)$$

$$\tilde{\mathbf{b}}_1 = \mathbf{b_1}$$

Note: $\tilde{\mathbf{b}}_i$ not in lattice

$$\tilde{\mathbf{b}}_2 = \mathbf{b}_2 - \frac{\tilde{\mathbf{b}}_1 \cdot \mathbf{b}_2}{|\tilde{\mathbf{b}}_1|^2} \tilde{\mathbf{b}}_1$$

$$\tilde{\mathbf{b}}_3 = \mathbf{b}_3 - \frac{\tilde{\mathbf{b}}_1 \cdot \mathbf{b}_3}{|\tilde{\mathbf{b}}_1|^2} \tilde{\mathbf{b}}_1 - \frac{\tilde{\mathbf{b}}_2 \cdot \mathbf{b}_3}{|\tilde{\mathbf{b}}_2|^2} \tilde{\mathbf{b}}_2$$

...

Lemma: Babai's nearest plane alg produces lattice point whose distance from target vector is at most

$$\frac{1}{2}\sqrt{\sum_i |\tilde{\mathbf{b}}_i|^2}$$

Today: SIS and LWE

Motivating question: what distribution over lattices to choose?

Short Integer Solution (SIS)

Parameterized by 4 quantities n, m, q, β

Last 3 are usually functions of first

n intuitively plays role of security parameter

q typically $q = O(n^c)$, but can also make exponential

m typically $m = \Omega(n \log q)$, but sometimes much bigger

β typically $\beta \geq \sqrt{m}$ but certainly $\beta \ll q$

Short Integer Solution (SIS)

Input: $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$ (short, wide)

Chosen uniformly at random

Goal: find vector $\mathbf{x} \in \mathbb{Z}^m$ such that:

$$\mathbf{A} \cdot \mathbf{x} \bmod q = 0$$

$$0 < |\mathbf{x}| \leq \beta$$

Claim: for $m > n \log q$ and $\beta \geq \sqrt{m}$, solution exists

Proof: consider $f_{\mathbf{A}} : \{0,1\}^m \to \mathbb{Z}_q^n$ defined as

$$f_{\mathbf{A}}(\mathbf{x}) = \mathbf{A} \cdot \mathbf{x} \bmod q$$

Domain size = 2^m Range size = $q^n < 2^m$

➡ Must exist distinct $\mathbf{x}_0, \mathbf{x}_1 \in \{0,1\}^m$ s.t.

$$f_{\mathbf{A}}(\mathbf{x}_0) = f_{\mathbf{A}}(\mathbf{x}_1)$$

➡ Let $\mathbf{x} = \mathbf{x}_0 - \mathbf{x}_1 \in \{-1,0,1\}^m$

SIS is a special case of SVP

$$\Lambda_q^{\perp}(\mathbf{A}) := \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} \bmod q = 0\}$$

Full-rank integer lattice

Approximate SVP in $\Lambda_q^{\perp}(\mathbf{A})$ for a random $\mathbf{A}$ is exactly SIS

Collision-resistance from SIS

$$f_{\mathbf{A}} : \{0,1\}^m \to \mathbb{Z}_q^n$$
$$f_{\mathbf{A}}(\mathbf{x}) = \mathbf{A} \cdot \mathbf{x} \bmod q$$

Collision = distinct $\mathbf{x}_0, \mathbf{x}_1 \in \{0,1\}^m$ s.t.
$$f_{\mathbf{A}}(\mathbf{x}_0) = f_{\mathbf{A}}(\mathbf{x}_1)$$

Security proof: let $\mathbf{x} = \mathbf{x}_0 - \mathbf{x}_1 \in \{-1,0,1\}^m$

Why the SIS distribution?

Atjai proved that SIS (on average) is as hard as approximate SVP in the worst case

That is, if you can solve SIS in polynomial-time on average, then you can solve approximate SVP in polynomial time on **any** lattice

Hardness of SIS

For polynomial-time attacks, best algorithm is typically LLL or variants

Works when $m \geq \Omega(\sqrt{n \log q})$, $\beta = 2^{O(\sqrt{n \log q})}$

Going forward, reducing mod q will produce a point in the interval

$$(-q/2, q/2]$$

Things close to 0 (positive or negative) don't get reduced

Learning with Errors (LWE)

Parameterized by 4 quantities n, m, q, σ

Last 3 are usually functions of first

n intuitively plays role of security parameter

q typically $q = O(n^c)$, but can also make exponential

m typically $m = \Omega(n \log q)$, but sometimes much bigger

σ typically $\sigma = \Omega(\sqrt{n})$ but certainly $\sigma \ll q$

Search LWE

Input: $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ (short, wide) Chosen uniformly at random

$\mathbf{u} = \mathbf{A}^T \cdot \mathbf{s} + \mathbf{e} \bmod q$ where

$\mathbf{s}$ uniform in $\mathbb{Z}_q^n$

$\mathbf{e} \in \mathbb{Z}^m$ “short”

Output: $\mathbf{s}$ (in this regime, $\mathbf{s}$ is whp unique)

The Distribution on e: Discrete Gaussians

D_σ = distribution over $\mathbb{Z}$ where

$$\Pr[x \leftarrow D_\sigma] \propto e^{-\pi x^2/\sigma^2}$$

Exact normalization constant is a big infinite sum, but for large σ can be approximated as

$$\Pr[x \leftarrow D_\sigma] \approx \frac{1}{\sigma} e^{-\pi x^2/\sigma^2}$$

D_σ^m = vector of m iid samples from D_σ

Search LWE

Input: $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$ (short, wide) Chosen uniformly at random

$\mathbf{u} = \mathbf{A}^T \cdot \mathbf{s} + \mathbf{e} \bmod q$ where

$\mathbf{s}$ uniform in $\mathbb{Z}_q^n$

$\mathbf{e} \leftarrow D_\sigma^m$

Output: $\mathbf{s}$ (in this regime, $\mathbf{s}$ is whp unique)

Decision LWE

Input: $\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}$ (short, wide) Chosen uniformly at random

Case 1: $\mathbf{u} = \mathbf{A}^T \cdot \mathbf{s} + \mathbf{e} \bmod q$ where

$\mathbf{s}$ uniform in $\mathbb{Z}_q^n$

$\mathbf{e} \leftarrow D_\sigma^m$

Case 2: $\mathbf{u}$ is random

Output: guess which case

LWE is a special case of CVP

$$\Lambda_q(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \exists \mathbf{s} \in \mathbb{Z}^n \text{ s.t. } \mathbf{x} = \mathbf{A}^T \cdot \mathbf{s} (\bmod q)\}$$

Full-rank integer lattice

LWE = CVP under, for random lattice and random target promised to be close to lattice

Public Key Encryption from LWE

$$\mathsf{pk} = (\mathbf{A}, \mathbf{u} = \mathbf{A}^T \cdot \mathbf{s} + \mathbf{e} \bmod q)$$
$$\mathsf{sk} = (\mathbf{s}, \mathbf{e})$$

$\mathbf{s}$ uniform in $\mathbb{Z}_q^n$

$\mathbf{e} \leftarrow D_\sigma^m$

$\mathsf{Enc}(\mathsf{pk}, m \in \{0,1\})$: Sample $\mathbf{r}$ uniform in $\{0,1\}^m$

Output $(\mathbf{v}^T = \mathbf{r}^T \mathbf{A}^T \;\; , \;\; w = \mathbf{r}^T \mathbf{u} + m\lfloor q/2 \rceil \bmod q)$

$\mathsf{Dec}(\mathsf{sk}, (\mathbf{v}, w))$: Compute

$$w - \mathbf{v}^T \cdot \mathbf{s} \bmod q = (\mathbf{r}^T \mathbf{A}^T \mathbf{s} + \mathbf{r}^T \mathbf{e} + m\lfloor q/2 \rceil) - \mathbf{r}^T \mathbf{A}^T \mathbf{s} \bmod q$$
$$= \mathbf{r}^T \mathbf{e} + m\lfloor q/2 \rceil \bmod q$$

Public Key Encryption from LWE

$$w - \mathbf{v}^T \cdot \mathbf{s} \bmod q = \mathbf{r}^T \mathbf{e} + m\lfloor q/2 \rceil \bmod q$$

$\mathbf{r} \in \{0,1\}^m$ $\mathbf{e}$ Gaussian of width σ

$\mathbf{r}^T \mathbf{e}$ is Guassian of width at most $\sigma\sqrt{m}$

With all but negligible probability, $|\mathbf{r}^T \mathbf{e}| \leq \sigma m$

$$\Rightarrow \mathbf{r}^T \mathbf{e} + m\lfloor q/2 \rceil \bmod q \approx \begin{cases} 0 & \text{if } m = 0 \\ \pm q/2 & \text{if } m = 1 \end{cases}$$

Decryption errors

Technically, there is a tiny chance that $\mathbf{r}^T\mathbf{e}$ is huge

In this case, decryption fails

Technically, scheme doesn't satisfy definition
we saw on first day of class

Def (PKE, syntax): A public key encryption scheme is a triple of algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ satisfying the following:

- $\mathsf{Gen}(1^\lambda)$: probabilistic polynomial-time (classical) procedure which takes as input a security parameter λ (represented in unary), and samples a secret/key public pair $(\mathsf{sk}, \mathsf{pk})$
- $\mathsf{Enc}(\mathsf{pk}, m)$: PPT procedure which takes as input the public key pk and message m, and samples a ciphertext c
- $\mathsf{Dec}(\mathsf{sk}, c)$: Deterministic PT procedure which takes as input the secret key sk and ciphertext c, and outputs a message m
- **Correctness:** $\forall \lambda, (\mathsf{sk}, \mathsf{pk})$ in support of $\mathsf{Gen}(1^\lambda), \forall m \in \{0,1\}^*$

$$\Pr[\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, m)) = m] = 1$$

Decryption errors

Solution 1: Truncate discrete Gaussian so that $\mathbf{e} \in [-B, B]^m$

$$B = \sigma\sqrt{m}$$

$\Rightarrow$ $|\mathbf{r}^T \mathbf{e}| \leq mB$ always

Generally results in larger error bounds → larger modulus
→ less efficient

Solution 2: Relax correctness definition to allow negligible probability of decryption errors

Sometimes (rarely) approximate correctness is insufficient

Lemma: Assuming decisional LWE, encryption scheme is CPA secure

Proof: Let $\mathcal{A}$ be a supposed adversary for the CPA-security of the encryption scheme

Define $W_b(\lambda)$ as the event that $\mathcal{A}$ outputs 1 in the following:

- Run $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Gen}(1^\lambda)$, give pk to $\mathcal{A}$
- $\mathcal{A}$ produces two msgs m_0, m_1
- Run $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$ and give c to $\mathcal{A}$
- $\mathcal{A}$ outputs an output guess $b' \in \{0,1\}$

Since message is binary, might as well take to be 0,1

Our goal: bound $|\Pr[W_0(\lambda)] - \Pr[W_1(\lambda)]| \leq \epsilon(\lambda)$ for negligible ϵ

Lemma: Assuming decisional LWE, encryption scheme is CPA secure

Proof: Let $\mathcal{A}$ be a supposed adversary for the CPA-security of the encryption scheme

Define $W_b(\lambda)$ as the event that $\mathcal{A}$ outputs 1 in the following:

- Run $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Gen}(1^\lambda)$, give pk to $\mathcal{A}$
- Run $c \leftarrow \mathsf{Enc}(\mathsf{pk}, b)$ and give c to $\mathcal{A}$
- $\mathcal{A}$ outputs an output guess $b' \in \{0, 1\}$

Our goal: bound $|\Pr[W_0(\lambda)] - \Pr[W_1(\lambda)]| \leq \epsilon(\lambda)$ for negligible ϵ

Lemma: Assuming decisional LWE, encryption scheme is CPA secure

Proof: Let $\mathcal{A}$ be a supposed adversary for the CPA-security of the encryption scheme

Define $W_b(\lambda)$ as the event that $\mathcal{A}$ outputs 1 in the following:

- Give $\mathsf{pk} = (\mathbf{A}, \mathbf{u} = \mathbf{A}^T \cdot \mathbf{s} + \mathbf{e} \bmod q)$ to $\mathcal{A}$
- Give $(\mathbf{v}^T = \mathbf{r}^T \mathbf{A}^T \;,\; w = \mathbf{r}^T \mathbf{u} + b\lfloor q/2 \rceil \bmod q)$ to $\mathcal{A}$
- $\mathcal{A}$ outputs an output guess $b' \in \{0, 1\}$

Our goal: bound $|\Pr[W_0(\lambda)] - \Pr[W_1(\lambda)]| \leq \epsilon(\lambda)$ for negligible ϵ

Lemma: Assuming decisional LWE, encryption scheme is CPA secure

Proof:

Define $\boxed{V_b(\lambda)}$ as the event that $\mathcal{A}$ outputs 1 in the following:

- Give $\boxed{(\mathbf{A}, \mathbf{u} \text{ uniform in } \mathbb{Z}_q^m)}$ to $\mathcal{A}$
- Give $(\mathbf{v}^T = \mathbf{r}^T \mathbf{A}^T \ , \ w = \mathbf{r}^T \mathbf{u} + b\lfloor q/2 \rceil \bmod q)$ to $\mathcal{A}$
- $\mathcal{A}$ outputs an output guess $b' \in \{0, 1\}$

LWE → $|\Pr[W_b(\lambda)] - \Pr[V_b(\lambda)]|$ is negligible

Lemma: Assuming decisional LWE, encryption scheme is CPA secure

Proof: claim: $|\Pr[V_0(\lambda)] - \Pr[V_1(\lambda)]|$ is negligible

Recall: **Leftover Hash Lemma:** 2-universal hash functions are good randomness extractors

Since entropy of $\mathbf{r}$ is $m \gg (n+1)\log q$

➡ $\mathbf{r}^T\mathbf{A}^T, \mathbf{r}^T\mathbf{u}$ is statistically close to uniform in $\mathbb{Z}_q^{n+1}$ (even given $\mathbf{A}, \mathbf{u}$)

➡ $(\mathbf{v}^T = \mathbf{r}^T\mathbf{A}^T \ \ , \ \ w = \mathbf{r}^T\mathbf{u} + b\lfloor q/2 \rceil \bmod q)$ hides b

Why the LWE distribution

Simple algebraic structure is easy to work with

As hard as worst-case lattice problems

Search-to-decision reduction (decision is no easier than search)

In classical cryptography, used for tons of interesting applications that are not known from other tools

Hardness of LWE

For polynomial-time attacks, best algorithm is typically LLL or variants

Works when $m \geq \Omega(\sqrt{n \log q})$, $q/\sigma \geq 2^{\Omega(\sqrt{n \log q})}$

For typical parameter settings, best attacks run in time $2^{O(n)}$

Note that this is very slightly sub-exponential in the secret size $n \log q$

Next Wednesday (10/29): Quantum algorithms for lattice problems

Reminder: no class on Monday 10/27