Proof Sketch: | one-wayness adversary B | challenger

$y^* \xleftarrow{R} y$

- will program y^* to i^* query to H
  (i^* is a random index)

assume A queries H on m before making signing query on m

signature adversary A

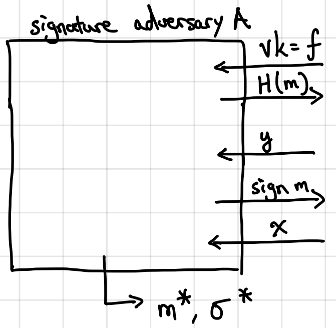

if this is query i^*: $y \leftarrow y^*$

else, $x \leftarrow D$, $y \leftarrow f(x)$, add $m \mapsto (x, y)$ to table

if $m \mapsto (x,y)$ is present in table, reply with x

otherwise abort

if m^* is query i^*, then output σ^*

otherwise abort

If A makes Q random oracle queries, B succeeds with probability $1/Q \cdot \mathrm{SigAdv}[A]$.

- All random oracle queries are properly distributed
  (since forward sampling and reverse sampling are statistically indistinguishable)
- All signature queries are properly distributed (as long as guess is correct)
- Guess is correct with prob. $1/Q$
- If guess is correct and A succeeds, then $f(\sigma^*) = H(m^*) = y^*$ so B succeeds.

Constructing preimage sampleable trapdoor functions from SIS.

$$f_A(x) := Ax \pmod q \qquad [A \in \mathbb{Z}_q^{n \times m},\ x \in \mathbb{Z}_q^m]$$

Goal: given a target vector $y \in \mathbb{Z}_q^n$, sample $x \in \mathbb{Z}^m$ such that $Ax = y$

Recall the SIS lattice

$$\mathcal{L}^\perp(A) = \{x \in \mathbb{Z}^m : Ax = 0 \pmod q\}$$

For a vector $u \in \mathbb{Z}_q^n$, recall that the coset $\mathcal{L}_u^\perp(A)$ is

$$\mathcal{L}_u^\perp(A) = c + \mathcal{L}^\perp(A) = \{x \in \mathbb{Z}^m : Ax = u\}$$

$c \in \mathbb{Z}_q^m$ is an arbitrary vector where $Ac = u$

} Equivalent formulation of objective: Sample from some "nice" distribution over $\mathcal{L}_u^\perp(A)$

Challenge: Defining a distribution over $\mathcal{L}_u^\perp(A)$ that is conducive for preimage sampling.

- Sampling preimages (given a trapdoor) must be efficient
- Samples must not leak trapdoor (can be simulated without knowledge of trapdoor)

The distribution is typically a discrete Gaussian distribution — shares many analytical properties with continuous Gaussians

Definition. For a parameter $s > 0$, we define the Gaussian function on $\mathbb{R}^n$ with width s as follows:

$$\rho_s(x) := \exp(-\pi \|x\|^2 / s^2)$$

$$\rho_{s,c}(x) := \exp(-\pi \|x - c\|^2 / s^2) \quad \text{Gaussian centered at } c \in \mathbb{R}^n$$

Discrete Gaussian centered at c:

$$D_{\mathcal{L},s,c}(x) \underbrace{\propto}_{\text{proportional to}} \begin{cases} \rho_{s,c}(x) & x \in \mathcal{L} \\ 0 & \text{otherwise} \end{cases} \quad \Bigg\} \quad D_{\mathcal{L},s,c}(x) = \frac{\rho_{s,c}(x)}{\underbrace{\sum_{x \in \mathcal{L}} \rho_{s,c}(x)}_{\text{normalization parameter}}} = \frac{\rho_{s,c}(x)}{\rho_{s,c}(\mathcal{L})}$$

Why discrete Gaussian distribution?

Nicely behaved:

- Rotationally invariant over $\mathbb{R}^n$: $\rho_s(\vec{x}) = \prod_{i \in [n]} \rho_s(x_i)$
  (density depends only on the norm $\|x\|$ of input)
- Sum of Gaussians is Gaussian
  (Gaussian convolution lemma)

Algorithms leverage these properties

Visually (2D setting) for $\mathbb{Z}^2$:

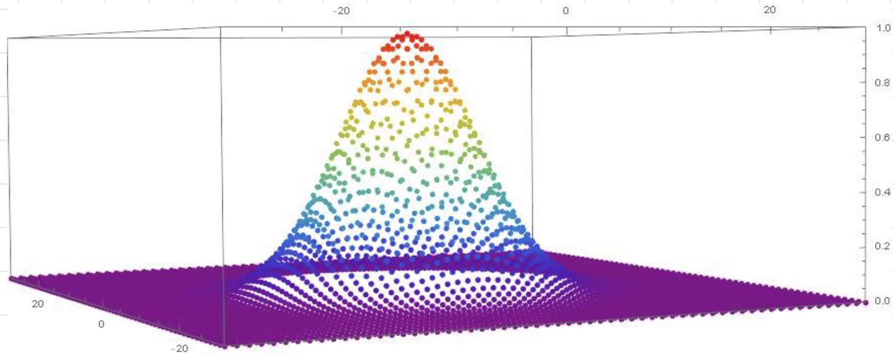

Here, we sketch one approach for discrete Gaussian sampling based on gadget trapdoors + perturbation [MP12, GPV08, Pei10]

We start by recalling the Gram-Schmidt orthogonalization method from linear algebra

- Let $b_1, \dots, b_n \in \mathbb{R}^d$ be a collection of vectors with $\mathrm{span}(b_1, \dots, b_n) = V$ [$b_1, \dots, b_n$ need not be linearly independent]
- The Gram-Schmidt orthogonalization process outputs $\tilde{b}_1, \dots, \tilde{b}_n \in \mathbb{R}^n$ where $\mathrm{span}(\tilde{b}_1, \dots, \tilde{b}_n) = V$ and $\tilde{b}_i^T \tilde{b}_j = 0$ for all $i \neq j$

  <u>Algorithm</u>: $\tilde{b}_1 \leftarrow b_1$

  for each $i = 2, \dots, n$:

$$\tilde{b}_i \leftarrow b_i - \sum_{j<i} \underbrace{\frac{b_i^T \tilde{b}_j}{\tilde{b}_j^T \tilde{b}_j} \cdot \tilde{b}_j}_{\text{projection of } b_i \text{ onto } \tilde{b}_j}$$

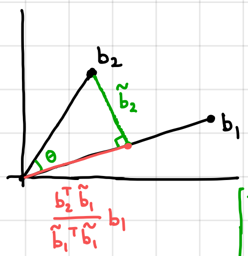

[In Euclidean space: $\frac{b_2^T \tilde{b}_1}{\tilde{b}_1^T \tilde{b}_1} b_1 = \frac{\|b_1\| \cdot \|b_2\| \cos \theta}{\|b_1\|^2} b_1 =$ projection of b_2 onto b_1]

  Not difficult to show that $\tilde{b}_1, \dots, \tilde{b}_n$ are pairwise orthogonal

- Let $\tilde{B} = [\tilde{b}_1 | \cdots | \tilde{b}_n]$ be the Gram-Schmidt basis.

<u>Important</u>: B and $\tilde{B}$ span the same vector space over $\mathbb{R}$, but do <u>not</u> necessarily generate the same lattice

($\tilde{b}_i$ is not necessarily an integer linear combination of $b_1, \dots, b_n$)

Simple counter-example: $b_1 = \begin{bmatrix} 2 \\ 0 \end{bmatrix}$ $b_2 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$

$\tilde{b}_1 = \begin{bmatrix} 2 \\ 0 \end{bmatrix}$ $\tilde{b}_2 = \begin{bmatrix} 0 \\ 2 \end{bmatrix}$ observe that $\tilde{b}_2 \notin \mathcal{L}(B)$

- The norm of the Gram-Schmidt vectors provides a bound on the minimum distance of a lattice:

$$\lambda_1(\mathcal{L}(B)) \geq \min_{i \in [n]} \|\tilde{b}_i\|$$

<u>Proof</u>. Take any lattice point $Bx \neq 0$ where $x \in \mathbb{Z}^n$. Let k be largest index where $x_k \neq 0$. Consider now the product

$$(Bx)^T \tilde{b}_k = \sum_{i \leq k} x_i b_i^T \tilde{b}_k = x_k b_k^T \tilde{b}_k$$ since $\tilde{b}_k$ is orthogonal to $b_1, \dots, b_{k-1}$ by construction

$$= x_k \|\tilde{b}_k\|^2$$ since $b_k^T \tilde{b}_k = \|\tilde{b}_k\|^2$ also by construction

By Cauchy-Schwarz ($|u^T v| \leq \|u\| \cdot \|v\|$), we now have

$$\|Bx\|_2 \cdot \|\tilde{b}_k\|_2 \geq |(Bx)^T \tilde{b}_k| = |x_k| \cdot \|\tilde{b}_k\|^2 \geq \|\tilde{b}_k\|^2$$ since $x_k \in \mathbb{Z}$ and $x_k \neq 0$

$$\geq \min_{i \in [n]} \|\tilde{b}_i\|.$$