<u>Summary so far</u>: - The SIS problem can be used to realize many symmetric primitives such as OWFs, CRHFs, and signatures
- Useful trick: "Concealing" a trapdoor (e.g., short matrix/basis) within a random-looking one — common theme in lattice-based cryptography.

For public-key primitives, we will rely on a very similar assumption: learning with errors (LWE), which can also be viewed as a "dual" of SIS. We introduce the assumption below:

errors are typically much smaller than q

<u>Learning with Errors (LWE)</u>: The LWE problem is defined with respect to lattice parameters n, m, q, χ, where χ is an <u>error distribution</u> over $\mathbb{Z}_q$ (oftentimes, this is a discrete Gaussian distribution over $\mathbb{Z}_q$). The $\mathsf{LWE}_{n,m,q,\chi}$ assumption states that for a random choice $A \xleftarrow{R} \mathbb{Z}_q^{n \times m}$, $s \xleftarrow{R} \mathbb{Z}_q^n$, $e \leftarrow \chi^m$, the following two distributions are computationally indistinguishable:

$$(A, s^T A + e^T) \overset{c}{\approx} (A, r)$$

where $r \xleftarrow{R} \mathbb{Z}_q^m$.

In words, the LWE assumption says that <u>noisy</u> linear combinations of a secret vector over $\mathbb{Z}_q^n$ looks indistinguishable from random.

A few notes/observations on LWE:
- Typically, m is sufficiently large so that the LWE secret s is <u>uniquely</u> determined.
- Without the error terms, this problem is easy for $m > n$: simply use Gaussian elimination to solve for s
- Observe that if SIS is easy, then LWE is easy. Namely, if the adversary can find a short $u \in \mathbb{Z}_q^m$ such that $Au = 0$, then, the adversary can compute

$$(s^T A + e^T) u = s^T A u + e^T u = e^T u \implies \|e^T u\| \leq m \cdot \|e\| \cdot \|u\|$$

this is small (compared to q)

$r^T u$ will be uniform over $\mathbb{Z}_q$, are unlikely to be small

LWE in "normal form"

- We can also choose the LWE secret from the error distribution (so it is short) — can be useful for both efficiency and for functionality (this is at least as hard as LWE with secrets drawn from <u>any</u> distribution, including the uniform one)
- Can also consider search vs. decision versions of the problem (i.e., search LWE says given $(A, s^T A + e^T)$, find s). There are search-to-decision reductions for LWE

<u>LWE as a lattice problem</u>: The search version of LWE essentially asks one to find s given $s^T A + e^T$. This can be viewed as solving the "bounded-distance decoding" (BDD) problem on the q-ary lattice

$$\mathcal{L}(A^T) = \{ s \in \mathbb{Z}_q^n : A^T s \} + q\mathbb{Z}^n$$

i.e., given a point that is close to a lattice element $s \in \mathcal{L}(A^T)$, find the point s

<u>Connections to worst-case hardness</u>: Regev showed that for any $m = \mathrm{poly}(n)$ and modulus $q < 2^{\mathrm{poly}(n)}$ and for a discrete Gaussian noise distribution (with values bounded by β), solving $\mathsf{LWE}_{n,m,q,\chi}$ is as hard as <u>quantumly</u> solving GapSVP_γ on arbitrary n-dimensional lattices with approximation factor $\gamma = \tilde{O}(n \cdot q/\beta)$

↳ Long sequence of subsequent works have shown <u>classical</u> reductions to worst-case lattice problems (for suitable instantiations of the parameters)

Symmetric encryption from LWE (for binary-valued messages)

Setup(1^λ) : Sample $s \xleftarrow{R} \mathbb{Z}_q^n$.

Encrypt$(s, \underset{\in\{0,1\}}{\mu})$: Sample $a \xleftarrow{R} \mathbb{Z}_q^n$ and $e \leftarrow \chi$. Output $(a,\ s^T a + e + \mu \cdot \lfloor \frac{q}{2} \rfloor)$.

Decrypt(s, ct) : Output $\underbrace{\lfloor ct_2 - s^T ct_1 \rceil_2}_{\text{"rounding operation"}}$

$$\lfloor x \rceil_2 = \begin{cases} 0 & \text{if } -\frac{q}{4} \leq x < \frac{q}{4} \\ 1 & \text{otherwise} \end{cases}$$

take $x \in \mathbb{Z}_q$ to be representative between $-\frac{q}{2}$ and $\frac{q}{2}$

Visually:

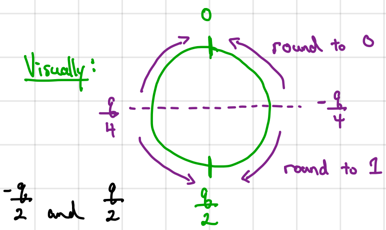

Correctness: $ct_2 - s^T ct_1 = s^T a + e + \mu \cdot \lfloor \frac{q}{2} \rfloor - s^T a$

$= \mu \cdot \lfloor \frac{q}{2} \rfloor + e$

if $|e| < \frac{q}{4}$, then decryption recovers the correct bit

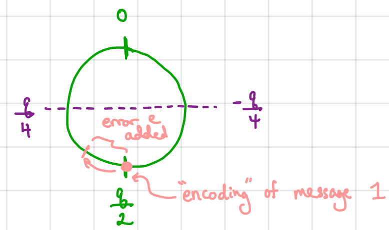

Security: By the LWE assumption, $(a, s^T a + e) \overset{c}{\approx} (a, r)$ where $r \xleftarrow{R} \mathbb{Z}_q$. Thus,

$$\underbrace{(a, s^T a + e)}_{\text{encryption of } 0} \underset{\text{LWE}}{\overset{c}{\approx}} \underbrace{(a, r) \equiv (a, r + \lfloor \tfrac{q}{2} \rfloor)}_{\text{since } r \text{ is uniform over } \mathbb{Z}_q} \underset{\text{LWE}}{\overset{c}{\approx}} \underbrace{(a, s^T a + e + \lfloor \tfrac{q}{2} \rfloor)}_{\text{encryption of } 1}$$

Observe: this encryption scheme is additively homomorphic (over $\mathbb{Z}_2$):

$$\begin{array}{l} (a_1,\ s^T a_1 + e_1 + \mu_1 \cdot \lfloor \frac{q}{2} \rfloor) \\ (a_2,\ s^T a_2 + e_2 + \mu_2 \cdot \lfloor \frac{q}{2} \rfloor) \end{array} \Rightarrow \left(a_1 + a_2,\ s^T(a_1 + a_2) + (e_1 + e_2) + (\mu_1 + \mu_2) \cdot \lfloor \tfrac{q}{2} \rfloor\right)$$

decryption then computes

$$(\mu_1 + \mu_2) \cdot \lfloor \tfrac{q}{2} \rfloor + e_1 + e_2$$

which when rounded yields $\mu_1 + \mu_2 \pmod 2$ provided that $|e_1 + e_2 + 1| < \frac{q}{4}$

Idea: We will rely on the LHL. We will include encryptions of 0 in the public key and refresh ciphertexts by taking a subset sum of encryptions of 0:

Regev's encryption scheme

Setup(1^λ): $A \xleftarrow{R} \mathbb{Z}_q^{n \times m}$ output $pk = (A, b^T)$

$s \xleftarrow{R} \mathbb{Z}_q^n$ $\quad b^T \leftarrow s^T A + e^T$ $\quad sk = s$

$e \leftarrow \chi^m$

can be viewed as m encryptions of 0 under the symmetric scheme with secret key s

Encrypt(pk, μ): sample $r \xleftarrow{R} \{0,1\}^m$

output $(Ar,\ b^T r + \mu \cdot \lfloor \frac{q}{2} \rfloor)$

Decrypt(sk, ct): output $\lfloor ct_2 - s^T ct_1 \rceil_2$

Correctness: $ct_2 - s^T ct_1 = b^T r + \mu \cdot \lfloor \frac{q}{2} \rfloor - s^T A r = s^T A r + e^T r + \mu \cdot \lfloor \frac{q}{2} \rfloor - s^T A r$

$= \mu \cdot \lfloor \frac{q}{2} \rfloor + e^T r$

if $|e^T r| < \frac{q}{4}$, then decryption succeeds (since e is small and r is binary, $e^T r$ is not large : $|e^T r| < m \|e\| \|r\| = m \|e\|$)

Security: Follows by LWE and LHL:

Hyb$_0$: Real public key

$\quad$ LWE

Hyb$_1$: Uniformly random public key (e.g. $b \xleftarrow{R} \mathbb{Z}_q^m$)

$\quad$ LHL: $(\bar{A}, \bar{A} r) \overset{s}{\approx} (\bar{A}, u)$ where $\bar{A} = \left[\frac{A}{b^T}\right] \xleftarrow{R} \mathbb{Z}_q^{(n+1) \times m}$, $r \leftarrow \{0,1\}^m$, and $u \xleftarrow{R} \{0,1\}$

Hyb$_2$: Uniformly random ciphertext (e.g., $ct = (u, t)$ where $u \xleftarrow{R} \mathbb{Z}_q^m$ and $t \xleftarrow{R} \{0,1\}$)