



Post-quantum Cryptography



Data Privacy and Security

Prof. Daniele Venturi
Dipartimento di Informatica



SAPIENZA
UNIVERSITÀ DI ROMA

Academic Year 2024-2025

Modern Cryptography

- Cryptography is **everywhere**
 - Credit cards, electronic passports, electronic commerce, electronic voting, cryptocurrencies, ...
- **Provable** security: **Reductions** to solving **hard problems**, given an attacker breaking security of cryptographic primitives
 - Requires to **believe** $P \neq NP$ (and in fact, that OWFs exist)
 - Examples: factoring, discrete logarithm, bilinear maps...
- History of **success**
 - Secret-key cryptography, public-key cryptography, identity-based cryptography, attribute-based cryptography, program obfuscation, ...



The Quantum Threat

- An algorithm by Shor [Sho94] solves the factoring and discrete logarithm problems in **polynomial-time** on a **quantum** machine
 - The algorithm requires an **ideal** quantum Turing machine
 - Factoring a 1024-bit integer requires **2050** logical **qubits** and a quantum circuit with **billions** of quantum gates
 - Despite recent progress on quantum computation, current implementations can only factor **tiny numbers** (e.g., 15 and 21)
- Nevertheless, the NIST started in 2017 a process to solicit, evaluate, and standardize **quantum-resistant** cryptography
 - The selected algorithms were announced in 2022
 - Most of these algorithms are based on **lattices**

What's the Rush?

- Big quantum computers won't be available for **many years**
 - If **ever**...
 - Can't we just wait?
- Better safe than sorry
 - **Harvesting attacks**: Store today's keys/ciphertexts to break later
 - **Rewrite history**: Forge signatures for old keys
 - Deploying new cryptography **at scale** requires 10+ years




Lattices



What is a Lattice?

- Simply, a set of points in a **high-dimensional** space
 - Arranged **periodically**
- Formally, take n **linearly independent** vectors $(\vec{b}_1, \dots, \vec{b}_n)$ in \mathbb{R}^n and consider all **integer** combinations

$$\mathcal{L} = \{a_1 \vec{b}_1 + \dots + a_n \vec{b}_n : a_1, \dots, a_n \in \mathbb{Z}\}$$



- We call $(\vec{b}_1, \dots, \vec{b}_n)$ a **basis**
- The same lattice may have **different equivalent** basis
 - Even if base vectors are long, there are short vectors in the lattice

History


- **Geometric** objects with rich mathematical structure
- Considerable **mathematical interest** starting from Gauss (1801), Hermite (1850), and Minkowski (1896)



- Recently, many **interesting applications** (cryptanalysis, factoring rational polynomials, finding integer relations, ...)

Equivalent Bases

- Sometimes, we write $\mathcal{L}(B)$ where B is the matrix whose columns are $(\vec{b}_1, \dots, \vec{b}_n)$
 - One can also define a lattice as a **discrete additive subgroup** of \mathbb{R}^n




- **Equivalent** bases:

- Permute vectors (i.e., $\vec{b}_i \leftrightarrow \vec{b}_j$)
- Negate vectors (i.e., $\vec{b}_i \leftarrow -\vec{b}_i$)
- Add integer multiple of another vector (i.e., $\vec{b}_i \leftarrow \vec{b}_i + k \cdot \vec{b}_j, k \in \mathbb{Z}$)

- **Theorem:** Two bases B_1, B_2 are **equivalent** iff $B_2 = B_1 \cdot U$
 - U **unimodular** (i.e., integer matrix with $\det(U) = \pm 1$)

The Fundamental Region

- The **fundamental region** of a lattice corresponds to a **periodic tiling** of \mathbb{R}^n by copies of some body
 - For instance, $[0,1)$ is a fundamental region of the integer lattice \mathbb{Z} , as every $x \in \mathbb{R}$ is in the unique translate $[x] + [0,1)$



- A lattice base yields a fundamental region called the **fundamental parallelepiped**

$$\mathcal{P}(B) = B \cdot [0,1)^n = \left\{ \sum_{i=1}^n c_i \cdot \vec{b}_i : c_i \in [0,1) \right\}$$

- Useful for measuring arbitrary points **relative to a lattice**
 - Note $x \bmod \mathcal{P}(B) = (a_1 \bmod 1)\vec{b}_1 + \dots + (a_n \bmod 1)\vec{b}_n$
 - A point x is in a lattice iff $x \bmod \mathcal{P}(B) = (0, \dots, 0)$

Determinant


- The **determinant** of a lattice $\mathcal{L}(B)$ is $\det(\mathcal{L}) = |\det(B)|$
- Note that this is well defined, as for every **unilateral** U

$$|\det(B \cdot U)| = |\det(B) \cdot \det(U)| = \det(B)$$

- The determinant corresponds to the **volume** of the fundamental parallelepiped
 - The determinant is the **reciprocal** of the **density** (i.e., big determinant means sparse lattice)
 - Moreover, the volume is the **same** for **every** fundamental region

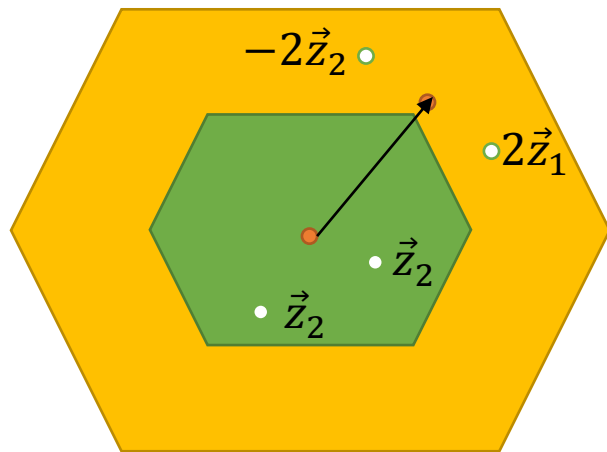
Successive Minima

- Let $\lambda_1(\mathcal{L})$ be the length of the **shortest non-zero** vector in a lattice \mathcal{L}
 - Usually, in terms of the **Euclidean** norm
 - The shortest vector is **never unique**, as for every $\vec{v} \in \mathcal{L}$ also $-\vec{v} \in \mathcal{L}$
- More generally, $\lambda_k(\mathcal{L})$ denotes the radius of the ball containing k **linearly independent** vectors
 - For $k = n$ the ball contains a basis of the entire space



Minkowski's Theorem

- **Lemma (Blichfeld)**: For any lattice \mathcal{L} and set \mathcal{S} with $\text{vol}(\mathcal{S}) > \det(\mathcal{L})$, there exist distinct $\vec{z}_1, \vec{z}_2 \in \mathcal{S}$ we have that $\vec{z}_1 - \vec{z}_2 \in \mathcal{L}$
 - The proof is simple and only requires volume arguments (exercise)
- **Theorem (Minkowski)**: For any lattice \mathcal{L} and **convex, zero-symmetric**, set \mathcal{S} with $\text{vol}(\mathcal{S}) > 2^n \det(\mathcal{L})$, there exists a **non-zero** lattice point in \mathcal{S}




- Let $\vec{z}_1, \vec{z}_2 \in \mathcal{S}/2$; by Blichfeld $\vec{z}_1 - \vec{z}_2 \in \mathcal{L}$
- Now, $2\vec{z}_1, -2\vec{z}_2 \in \mathcal{S}$
- So, their average $\vec{z}_1 - \vec{z}_2 \in \mathcal{S}$
- **Corollary (Minkowski)**: For every \mathcal{L} , we have that $\lambda_1(\mathcal{L}) \leq \sqrt{n} \cdot \det(\mathcal{L})^{1/n}$

Hard Problems

- **SVP $_{\gamma}$** : Given B , find a vector in $\mathcal{L}(B)$ with length $\leq \gamma \cdot \lambda_1(\mathcal{L}(B))$
- **GapSVP $_{\gamma}$** : Given B , **decide** if $\lambda_1(\mathcal{L}(B))$ is ≤ 1 or $\geq \gamma$
- **SIVP $_{\gamma}$** : Given B , find n **linearly independent** vectors in $\mathcal{L}(B)$ with length $\leq \gamma \cdot \lambda_n(\mathcal{L}(B))$
- **CVP $_{\gamma}$** : Given B and \vec{v} , find a lattice point that is at most γ times **farther** than the **closest** lattice point
 - It is known that **SVP $_{\gamma} \leq$ CVP $_{\gamma}$**
- **BDD**: Find **closest** lattice point, given that \vec{v} is **already close**

General Hardness Results



- Exact algorithms take time 2^n
- **Polynomial-time** algorithm for gap $\gamma = 2^n \log \log n / \log n$
- No better **quantum** algorithm known
- **NP hardness** for gap $\gamma = n^c / \log \log n$
 - For cryptographic applications, we need $\gamma = \Omega(n)$
 - Not believed to be *NP*-hard for $\gamma = \sqrt{n}$

Small Integer Solution Problem

- Fix **dimension** n , and **modulus** q (e.g., $q \approx n^2$)
- Given **random** vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{Z}_q^n$, find **non-zero small** $z_1, \dots, z_m \in \mathbb{Z}$ such that

$$z_1 \cdot \mathbf{a}_1 + z_2 \cdot \mathbf{a}_2 + \dots + z_m \cdot \mathbf{a}_m = \mathbf{0} \quad \text{in } \mathbb{Z}_q^n$$

- Observations:
 - Trivial if the size of the z_i 's is **not restricted** (Gaussian elimination)
 - Equivalently, find **non-zero short** $\mathbf{z} \in \mathbb{Z}^m$ s.t. $\mathbf{A} \cdot \mathbf{z} = \mathbf{0} \in \mathbb{Z}_q^n$

SIS as a Lattice Problem


- Matrix $A = (\mathbf{a}_1, \dots, \mathbf{a}_m) \in \mathbb{Z}_q^{n \times m}$
 $\mathcal{L}^\perp(A) = \{\mathbf{z} \in \mathbb{Z}^m : A \cdot \mathbf{z} = \mathbf{0}\}$

Find **short** ($\|\mathbf{z}\| \leq \beta \ll q$)
solutions for **random** A

- Theorem (Ajt96).** For **any** n -dimensional lattice, it holds that:

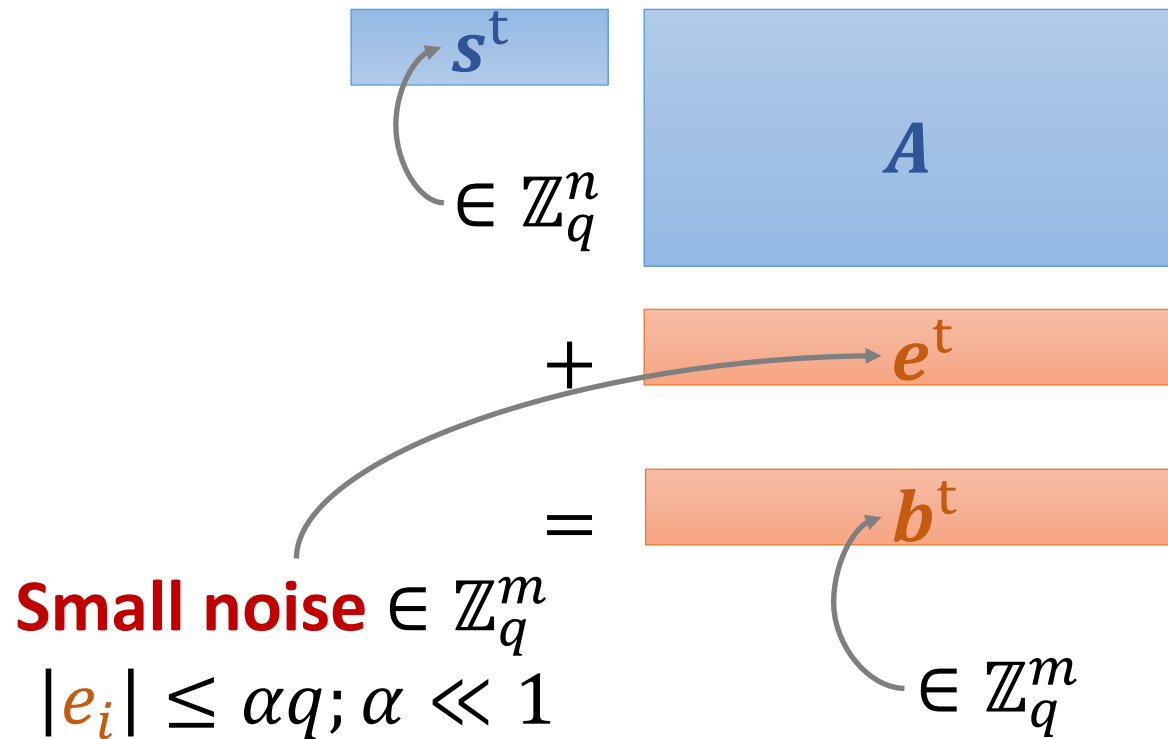
$$\text{GapSVP}_{\beta\sqrt{n}}, \text{SIVP}_{\beta\sqrt{n}} \leq \text{SIS}_\beta$$

- Also true for any lattice **coset** $\mathcal{L}_u^\perp(A) = \{\mathbf{z} \in \mathbb{Z}^m : A \cdot \mathbf{z} = \mathbf{u}\} = \mathbf{u} + \mathcal{L}^\perp(A)$ (i.e., **inhomogeneous** SIS)



Learning with Errors [Reg05]


- Dimension n , modulus $q > 2$, **noise** distribution χ
- **Find** $s \in \mathbb{Z}_q^n$ given m **noisy random** inner product equations



- Trivial **without** noise
- **Gaussian** distribution over \mathbb{Z} , with std deviation $\geq \sqrt{n}$ and $\ll q$
 - Rate parameter $\alpha \ll 1$
- Need $\alpha q > \sqrt{n}$ for **worst-case hardness** and because there is an $\exp((\alpha q)^2)$ -time attack

Decisional LWE

- **Distinguish** the matrix A and the vector b from random (A, b)
 - Decisional LWE is **equivalent** to Search LWE



LWE as a Lattice Problem

- Matrix $A = (\mathbf{a}_1, \dots, \mathbf{a}_m) \in \mathbb{Z}_q^{n \times m}$
 $\mathcal{L}(A) = \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{z}^t = \mathbf{s}^t \cdot A\}$


LWE is BDD on $\mathcal{L}(A)$: Given

$$\mathbf{b}^t \approx \mathbf{z}^t = \mathbf{s}^t \cdot A \text{ find } \mathbf{z}$$

- **Theorem (Reg05, Pei10).** For **any** n -dimensional lattice, it holds that:

$$\text{GapSVP}_{\alpha n}, \text{SIVP}_{\alpha n} \leq \text{LWE}$$

- **Quantum** reduction for **broad** parameters [Reg05]
- **Classical** reduction for **restricted** parameters (e.g., $q \approx 2^n$) [Pei10]



Hardness of LWE

- More formally define the **LWE distribution** as

$$\text{LWE}[n, m, q, \chi] = \left\{ (A, b) : \begin{array}{l} A \leftarrow \mathbb{Z}_q^{n \times m}; s \leftarrow \mathbb{Z}_q^n; \\ e \leftarrow \chi^m; b^t = [s^t \cdot A + e^t]_q \end{array} \right\}$$

- Parameters:
 - $\alpha = 1/\text{poly}(n)$ or $\alpha = 2^{-n^\epsilon}$ (**stronger** assumption as α **decreases**)
 - $m = \Theta(n \log q)$ or $m = \text{poly}(n)$ (**stronger** assumption as m **increases**)
 - $q = 2^{n^\epsilon}$ or $q = \text{poly}(n)$ (**stronger** assumption as q **increases**)
 - Noise distribution χ such that $\mathbb{P}[|e| > \alpha q : e \leftarrow \chi] \leq \text{negl}(n)$

Simple Properties

- Check a **candidate** solution $\mathbf{t} \in \mathbb{Z}_q^n$
 - Test if all $\mathbf{b} - \langle \mathbf{t}, \mathbf{a} \rangle$ are small
 - If $\mathbf{t} \neq \mathbf{s}$, then $\mathbf{b} - \langle \mathbf{t}, \mathbf{a} \rangle = \langle \mathbf{s} - \mathbf{t}, \mathbf{a} \rangle + \mathbf{e}$ is **well-spread** in \mathbb{Z}_q
- **Shift** the secret by any $\mathbf{r} \in \mathbb{Z}_q^n$
 - Given $(\mathbf{a}, \mathbf{b} = \langle \mathbf{s}, \mathbf{a} \rangle + \mathbf{e})$, output $(\mathbf{a}, \mathbf{b}' = \mathbf{b} + \langle \mathbf{r}, \mathbf{a} \rangle = \langle \mathbf{s} + \mathbf{r}, \mathbf{a} \rangle + \mathbf{e})$
 - Using **random** \mathbf{r} yields a random **self-reduction**
 - **Amplification** of success probabilities (i.e., **non-negligible** success probability for **random** $\mathbf{s} \in \mathbb{Z}_q^n$ implies **overwhelming** success probability for **every** $\mathbf{s} \in \mathbb{Z}_q^n$)
- **Multiple** secrets: $(\mathbf{a}, \mathbf{b}_1 = \langle \mathbf{s}_1, \mathbf{a} \rangle + \mathbf{e}_1, \dots, \langle \mathbf{s}_t, \mathbf{a} \rangle + \mathbf{e}_t)$ indistinguishable from **random** $(\mathbf{a}, \mathbf{b}_1, \dots, \mathbf{b}_t)$

Search/Decision Equivalence

- Suppose we are given an oracle that **perfectly distinguishes** pairs $(\mathbf{a}, b = \langle \mathbf{s}, \mathbf{a} \rangle + e)$ from random (\mathbf{a}, b)
- To find s_1 , it suffices to **test** if $s_1 = 0$
 - Because we can **shift** s_1 by $0, 1, \dots, q - 1$ (assuming $q = \text{poly}(n)$)
 - Then we can do the same for s_2, \dots, s_n
- The test: For each (\mathbf{a}, b) , choose **random** $r \in \mathbb{Z}_q$ and invoke the oracle on pairs $(\mathbf{a}' = \mathbf{a} - (r, 0, \dots, 0), b)$
- Note that $b = \langle \mathbf{s}, \mathbf{a}' \rangle + s_1 \cdot r + e$
 - If $s_1 = 0$, then $b = \langle \mathbf{s}, \mathbf{a}' \rangle + e$ and the oracle **accepts**
 - If $s_1 \neq 0$, then b is **uniform** (assuming q **prime**) and the oracle **rejects**

LWE with Short Secrets

- **Theorem [M01,ACPS09]:** LWE is **no easier** if the secret is drawn from the **error distribution** χ
 - Intuition: Finding **e equivalent** to finding s (i.e., $b^t - e^t = s^t \cdot A$)
- **Transformation** from secret $s \in \mathbb{Z}_q^n$ to secret $\bar{e} \leftarrow \chi^n$
 - Draw samples to get $(\bar{A}, \bar{b}^t = s^t \cdot \bar{A} + \bar{e}^t)$ for square, invertible, \bar{A}
 - Transform each **additional** sample $(a, b = \langle s, a \rangle + e)$ to

$$a' = -\bar{A}^{-1} \cdot a, b' = b + \langle \bar{b}, a' \rangle = \langle \bar{e}, a' \rangle + e$$


- This maps **uniform** (a, b) to **uniform** (a', b') , and thus works for **decision** LWE too

LWE vs SIS

- SIS has **many** valid solutions, whereas LWE only has **one**
- **LWE \leq SIS**
 - Given \mathbf{z} such that $\mathbf{A} \cdot \mathbf{z} = \mathbf{0}$ from an SIS oracle, compute $\mathbf{b}^t \cdot \mathbf{z}$
 - Now, $\mathbf{b}^t \cdot \mathbf{z} = \mathbf{e}^t \cdot \mathbf{z}$ is **small** in the LWE case, whereas $\mathbf{b}^t \cdot \mathbf{z}$ is **well-spread** in case \mathbf{b}^t is uniformly random
- What about the other direction?
 - Not known **in general**
 - True under **quantum reductions**

Efficiency of LWE/SIS

- Getting **one** random-looking scalar $b_i \in \mathbb{Z}_q$ requires an n -dimensional **inner product** mod q



- Can **amortize** each column a_i over **many secrets** s_j , but the latter still requires $\tilde{O}(n)$ work per scalar output
- Public keys are **rather large**, i.e. $> n^2$ time to encrypt/decrypt an n -bit message
- Can we do better?

Wishful Thinking...

$$\begin{array}{|c|} \hline s^t \\ \hline \end{array} \star \begin{array}{|c|} \hline a^t \\ \hline \end{array} + \begin{array}{|c|} \hline e^t \\ \hline \end{array} = \begin{array}{|c|} \hline b^t \\ \hline \end{array}$$

$\in \mathbb{Z}_q^d$

- Get d **pseudorandom** scalars from just one **cheap product** operation \star
- Replace $\mathbb{Z}_q^{d \times d}$ **chunks** with \mathbb{Z}_q^d

- **Main question:** How to define the product \star so that (a, b) is **pseudorandom**
 - Requires care: **coordinate-wise** product **insecure** for **small** errors
- **Answer:** Let \star be multiplication **in a polynomial ring**, e.g. $\mathbb{Z}_q^d[X]/(X^d + 1)$
 - **Fast** and **practical** with the FFT: $d \log d$ operations mod q
 - The same **ring structure** used in NTRU [HPS08]

LWE over Rings/Modules

- Let $R = \mathbb{Z}[X]/(X^d + 1)$ for d a power of 2 and $R_q = R/qR$
 - Elements of R_q are degree $< d$ **polynomials** with coefficients mod q
 - Operations over R_q are **very efficient** using FFT-like algorithms
- **Search LWE**: Find secret vector of **polynomials** \mathbf{s} in R_q^k given

$$\mathbf{s}^t \star \mathbf{a}_i^t + \mathbf{e}_i^t = \mathbf{b}_i^t$$

$\mathbf{b}_i^t \in \mathbb{Z}_q^d$

- Each equation is d **related equations** on a secret of dimension $n = d \cdot k$
 - LWE: $d = 1, k = n$
 - Ring-LWE: $d = n, k = 1$
 - Module-LWE: Interpolate
- **Decision LWE**: Distinguish $(\mathbf{a}_i, \mathbf{b}_i)$ from uniform $(\mathbf{a}_i, \mathbf{b}_i)$ in $R_q^k \times R_q$

Hardness of Ring/Module-LWE

- **Theorem [LPR10]**: For any $R = \mathcal{O}_K$

$$R^k\text{-GapSVP} \leq \text{search } R^k\text{-LWEdecision} \leq R^k\text{-LWE}$$

- Can we **dequantize** the worst-case/average-case reduction?
 - The **classical** GapSVP \leq LWE reduction is of little use: for the relevant factors, GapSVP for **ideals** (i.e., $k = 1$) is **easy**
- How **hard** (or not) is **GapSVP** on **ideal/module lattices**?
 - For **polynomial approximation** no significant improvement versus general lattices (even for ideals)
 - For **subexponential approximation** we have better **quantum** algorithms for **ideals**, but not for $k > 1$
- **Reverse** reductions? Seems not **without** increasing k ...

Why Lattice-based Cryptography?

- **Provable** security
 - If scheme is not secure, one can solve hard mathematical problems
 - Not always happens in current implementations (e.g., RSA)
- **Worst-case** security
 - If scheme not secure, one can break **every** instance of lattice problems
 - Factoring and discrete log only guarantee **average-case** security
- Still **unbroken** by quantum algorithms
 - No progress over the last 50 years
 - But we don't know: see <https://eprint.iacr.org/2024/555>
- Efficiency
 - Mainly additions/multiplications, no modular exponentiations


Basic Cryptographic Applications



One-Way Functions

- Parameters $m, n, q \in \mathbb{Z}$, key $A \in \mathbb{Z}_q^{n \times m}$
- Input $x \in \{0,1\}^m$, output $f_A(x) = A \cdot x$
- **Theorem [Ajt96]**: For $m > n \log q$, if **SIVP** is **hard** to approximate in the **worst-case**, then f_A is **one-way**
- Cryptanalysis: Given A, y , find x such that $y = A \cdot x$
 - **Easy** problem: find **arbitrary** u such that $y = A \cdot u$
 - All solutions $y = A \cdot x$ are of the form $t + \mathcal{L}^\perp(A)$
 - Requires to find **small** vector in $t + \mathcal{L}^\perp(A)$ or to find a lattice point $v \in \mathcal{L}^\perp(A)$ **close** to t (**average-case** instance of CVP w.r.t. $\mathcal{L}^\perp(A)$)

Collision-resistant Hash Functions



Collisions **exists**
inherently, but are
hard to find
efficiently

- Given $A = (\mathbf{a}_1, \dots, \mathbf{a}_m)$, define $h_A: \{0,1\}^m \rightarrow \mathbb{Z}_q^n$

$$h_A(z_1, \dots, z_m) = \mathbf{a}_1 \cdot z_1 + \dots + \mathbf{a}_m \cdot z_m$$

- Set $m > n \log q$ in order to get **compression**
- A collision $\mathbf{a}_1 \cdot z_1 + \dots + \mathbf{a}_m \cdot z_m = \mathbf{a}_1 \cdot z'_1 + \dots + \mathbf{a}_m \cdot z'_m$ yields $\mathbf{a}_1 \cdot (z_1 - z'_1) + \dots + \mathbf{a}_m \cdot (z_m - z'_m) = \mathbf{0}$, with $z_m - z'_m \in \{-1, 0, 1\}$

Commitments

- Analogy: **lock** message in a box, give the box, keep the key
 - Later give the key to **open** the box
- Implementation:
 - **Randomized** function $\mathbf{Com}(x; r)$, where x is the message and r is the randomness
 - To **open** a commitment simply reveal (x, r)
- Security properties
 - **Hiding:** $\mathbf{Com}(x; r)$ **reveals nothing** on x
 - **Binding:** **Can't open** $\mathbf{Com}(x; r)$ to $x' \neq x$

Commitments

- Take two **random** SIS matrices A_1, A_2
- The **message** is $x \in \{0,1\}^m$ and the **randomness** is $r \in \{0,1\}^m$
- Commitment: $\mathbf{Com}(x; r) = f_{A_1, A_2}(x, r) = A_1 \cdot x + A_2 \cdot r$
 - **Hiding:** $A_2 \cdot r = f_{A_2}(r)$ is **statistically** close to **uniform** over \mathbb{Z}_q^n , and thus x is information-theoretically **hidden**
 - **Binding:** Finding (x, r) and (x', r') such that $\mathbf{Com}(x; r) = \mathbf{Com}(x'; r')$ directly contradicts the **collision resistance** of f_{A_1, A_2}

Leftover Hash Lemma

- Let H be a family of **universal hash functions** with domain D and image I . Then, for $x \leftarrow_{\$} D$, $h \leftarrow_{\$} H$, and $u \leftarrow_{\$} I$:
$$\text{SD} \left((h, h(x)); (h, u) \right) \leq 1/2 \cdot \sqrt{|I|/|D|}$$
- Note that the function $h_A(\vec{r}) = [A \times \vec{r}]_q$ is **universal**
 - As $\forall \vec{r}_1 \neq \vec{r}_2: \mathbb{P}_A[h_A(\vec{r}_1) = h_A(\vec{r}_2)] = \mathbb{P}_A[A \times (\vec{r}_1 - \vec{r}_2) = \vec{0}] = q^{-n}$
- Hence, for $\vec{r} \leftarrow_{\$} \{0,1\}^m$, $A \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$, and $\vec{u} \leftarrow_{\$} \mathbb{Z}_q^n$, whenever $m = 2 + n \log q + 2n$

$$\text{SD} \left((A, [A \times \vec{r}]_q); (A, \vec{u}) \right) \leq 1/2 \cdot \sqrt{q^n / 2^m} \leq 2^{-n}$$

NIST Standards



Falcon



Lattice Trapdoors

- Recall: Lattice-based **one-way functions**

$$f_A(\mathbf{x}) = \mathbf{A} \cdot \mathbf{x} \bmod q \in \mathbb{Z}_q^n$$

(short \mathbf{x} , surjective)

$$f_A(\mathbf{s}, \mathbf{e}) = \mathbf{s}^t \cdot \mathbf{A} + \mathbf{e}^t \bmod q \in \mathbb{Z}_q^m$$

(short \mathbf{e} , injective)

- Task: **Invert** f_A
 - Find the **unique** \mathbf{s} (or \mathbf{e}) such that $f_A(\mathbf{s}, \mathbf{e}) = \mathbf{s}^t \cdot \mathbf{A} + \mathbf{e}^t \bmod q$
 - Given $\mathbf{u} = f_A(\mathbf{x}') = \mathbf{A} \cdot \mathbf{x}' \bmod q$, **sample random** $\mathbf{x} \leftarrow f_A^{-1}(\mathbf{u})$ with probability proportional to $\exp(-\|\mathbf{x}\|^2/s^2)$
- How? Via a **strong trapdoor** for \mathbf{A} (a **short basis** of $\mathcal{L}^\perp(\mathbf{A})$)
 - Deeply studied question [Babai86,Ajtai99,Klein01,GPV08,AP09,P10]

A Different Kind of Trapdoor [MP12]

- Drawbacks of previous solutions
 - Generating A with short basis is **complex** and **slow**
 - Inversion algorithms trade-off **quality** (i.e., length of basis vectors which depends on the Gaussian std parameter s) for **efficiency**
- Alternative: The trapdoor is **not a basis**
 - But just **as powerful**
 - **Simpler** and **faster**
- Overview of method
 - Start with **fixed, public**, lattice defined by **gadget matrix** G which admits very **fast**, and **parallel**, algorithms for f_G^{-1}
 - **Randomize** G into A via nice **unimodular** transform (the trapdoor)
 - **Reduce** f_A^{-1} to f_G^{-1} plus some pre/post-processing

Step 1: The Gadget Matrix

- Let $q = 2^k$ and take $\mathbf{g} = [1 \quad 2 \quad \dots \quad 2^{k-1}] \in \mathbb{Z}_q^{1 \times k}$
- To invert $f_{\mathbf{g}}: \mathbb{Z}_q \times \mathbb{Z}^k \rightarrow \mathbb{Z}_q^k$

$$f_{\mathbf{g}}(s, \mathbf{e}) = s \cdot \mathbf{g} + \mathbf{e} = [s + e_0 \quad 2s + e_1 \quad \dots \quad 2^{k-1}s + e_{k-1}] \bmod q$$

- Get lsb of s from $2^{k-1}s + e_{k-1}$, then repeat for the next bits of s
- Works when $e_{k-1} \in [-q/4, q/4)$
- To sample Gaussian preimage for $\mathbf{u} = f_{\mathbf{g}}(\mathbf{x}) = \langle \mathbf{g}, \mathbf{x} \rangle$
 - For $i \in [0, k-1]$, choose $x_i \leftarrow (2\mathbb{Z} + u)$ and let $u \leftarrow (u - x_i)/2 \in \mathbb{Z}$
 - E.g., $k = 2$: $x_0 \leftarrow (2z_0 + u)$, $u \leftarrow (u - 2z_0 - u)/2 = -z_0$, $x_1 \leftarrow (2z_1 - z_0)$, $\langle \mathbf{g}, \mathbf{x} \rangle = 2z_0 + u + 2(2z_1 - z_0) = u + 4z_1 = u \bmod 4$

Step 1: The Gadget Matrix G

- Alternative view: The **lattice** $\mathcal{L}^\perp(\mathbf{g})$ has **basis**

$$\mathbf{S} = \begin{bmatrix} 2 & & & & \\ -1 & 2 & & & \\ & -1 & \ddots & & \\ & & \ddots & 2 & \\ & & & -1 & 2 \end{bmatrix} \in \mathbb{Z}^{k \times k}, \text{ with } \tilde{\mathbf{S}} = 2 \cdot \mathbf{I}_k$$

- The above inversion algorithms are special cases of the randomized **nearest-plan algorithm** [Bab86,Kle01,GPV08]
- Define $\mathbf{G} = \mathbf{I}_n \otimes \mathbf{g} \in \mathbb{Z}^{n \times nk}$ (where \otimes is the **tensor** product)
 - Computing $f_{\mathbf{G}}^{-1}$ reduces to n **parallel calls** to $f_{\mathbf{g}}^{-1}$
 - Also applies to $\mathbf{H} \cdot \mathbf{G}$, for any **invertible** $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$

Step 2: Randomize G

- Define **semi-random** $[\bar{A}|G]$ for **uniform** $\bar{A} \in \mathbb{Z}_q^{n \times \bar{m}}$
 - It can be seen that inverting $f_{[\bar{A}|G]}^{-1}$ **reduces** to inverting f_G^{-1} [CHKP10]
- Choose a **short Gaussian** $R \in \mathbb{Z}^{\bar{m} \times n \log q}$ and let

$$A = [\bar{A}|G] \cdot \begin{bmatrix} I & R \\ & I \end{bmatrix} = [\bar{A}|G - \bar{A}R]$$

- A is **uniform** because, by the **leftover hash lemma**, $[\bar{A}|\bar{A}R]$ is **statistically close** to uniform when $\bar{m} \approx n \log q$
- Alternatively, $[I|\bar{A}] - \bar{A} \cdot R_1 + R_2$ is **pseudorandom** under the LWE assumption (in normal form)

A New Trapdoor Notion

- We constructed $A = [\bar{A} | G - \bar{A}R]$
- Say that R is a **trapdoor** for A with **tag** $H \in \mathbb{Z}_q^{n \times n}$ (invertible) if

$$A \cdot \begin{bmatrix} R \\ I \end{bmatrix} = H \cdot G$$

- The **quality** of R is $s_1(R) = \max_{u: \|u\|=1} \|R \cdot u\|$
- **Fact:** $s_1(R) \approx (\sqrt{\text{rows}} + \sqrt{\text{cols}}) \cdot r$ for Gaussian entries w/ std dev r
- Also R is a trapdoor for $A - [\mathbf{0} | H' \cdot G]$ with tag $H - H'$ [ABB10]
- Relating new and old trapdoors
 - Given basis S for $\mathcal{L}^\perp(G)$ and trapdoor R for A , one can **efficiently** construct **basis** S_A for $\mathcal{L}^\perp(G)$ where $\|\tilde{S}_A\| \leq (s_1(R) + 1) \cdot \|\tilde{S}\|$


Step 3: Reduce f_A^{-1} to f_G^{-1}

- Let R be a **trapdoor** for A with **tag** $H = I$: $A \cdot \begin{bmatrix} R \\ I \end{bmatrix} = G$
- Inverting LWE
 - Given $b^t = s^t \cdot A + e^t$, recover s from $b^t \cdot \begin{bmatrix} R \\ I \end{bmatrix} = s^t \cdot G + e^t \cdot \begin{bmatrix} R \\ I \end{bmatrix}$
 - Works if **each entry** of $e^t \cdot \begin{bmatrix} R \\ I \end{bmatrix} \in [-q/4, q/4)$
- Inverting SIS
 - Given u , sample $z \leftarrow f_G^{-1}(u)$ and output $x = \begin{bmatrix} R \\ I \end{bmatrix} \cdot z \in f_A^{-1}(u)$
 - Indeed, $A \cdot x = G \cdot z = u$

Leaks about R !

$$\Sigma = \mathbb{E}_x[x \cdot x^t] = \mathbb{E}_z[R \cdot z \cdot z^t \cdot R^t] \approx R \cdot R^t$$

Step 3: Perturbation Method [P10]



- To **fix** the covariance

$$u^t \cdot \Sigma_2 \cdot u = s^2 - u^t \cdot \Sigma_1 \cdot u > 0$$

- Generate **perturbation** vector p with covariance $s^2 \cdot I - R \cdot R^t$
- Sample **spherical** z such that $G \cdot z = u - A \cdot p$
- Output $x = p + \begin{bmatrix} R \\ I \end{bmatrix} \cdot z$

$$A \cdot x = A \cdot p + A \cdot \begin{bmatrix} R \\ I \end{bmatrix} \cdot z = A \cdot p + G \cdot z = u$$


Falcon: Digital Signatures from SIS

- Generate **uniform** $vk = A$ with **trapdoor** $sk = T$
- To sign μ , use T to **sample** $\sigma = x \in \mathbb{Z}^m$ such that $A \cdot x = H(\mu)$, where H is a **public** hash function
 - Recall that x is drawn from a **Gaussian distribution**, which **reveals nothing** about the trapdoor T
- To verify $(\mu, \sigma = x)$ under $vk = A$ simply check $A \cdot x = H(\mu)$ and that x is **sufficiently short**
- Security: **Forging** a signature for a new message μ^* requires finding a **short** x^* such that $A \cdot x^* = H(\mu^*)$
 - This is **equivalent** to solving the SIS problem
 - Signatures queries **do not help** because they **reveal nothing** about the trapdoor T

Crystals-Dilithium

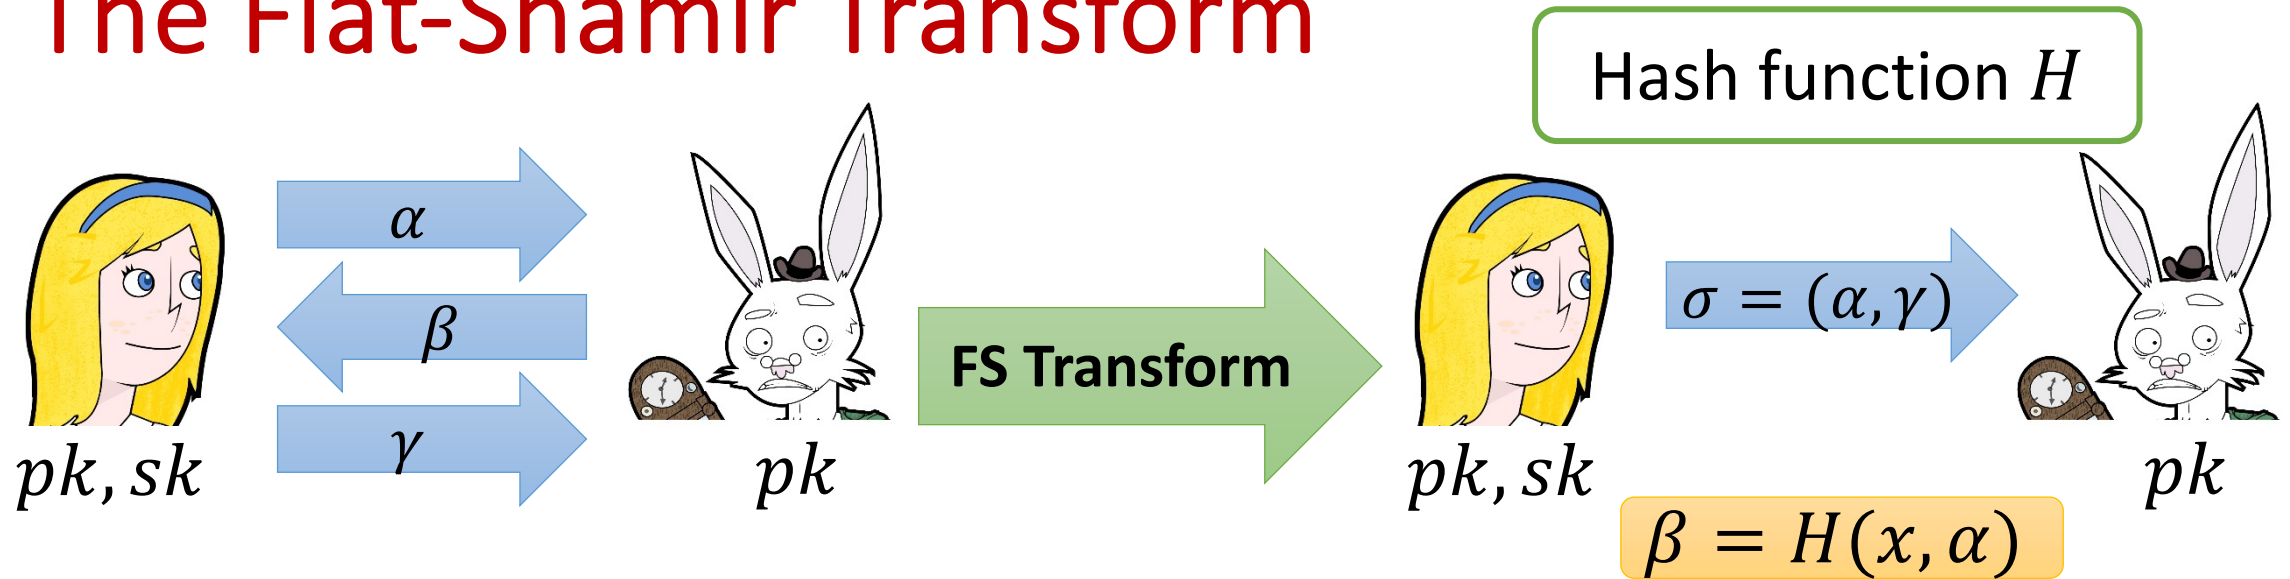


Canonical Identification Schemes



- **Completeness**: The **honest** prover convinces the **honest** verifier (with all but a negligible probability)
- **Passive Security**: No (**efficient**) **malicious** prover knowing only pk can convince the **honest** verifier
 - Even in case the attacker knows many **accepting transcripts** corresponding to **honest** protocol executions

The Fiat-Shamir Transform




- Given a **canonical** ID scheme, we can derive a **signature scheme** as follows:
 - Alice obtains $\sigma = (\alpha, \gamma)$ from the **prover**, using the **secret key** sk and choosing $\beta = H(x, \alpha)$
 - Bob checks that (α, β, γ) is a **valid transcript**, with $\beta = H(x, \alpha)$

The Fiat-Shamir Transform

Theorem [FS86]. If the ID scheme is **passively** secure, the signature derived via the **Fiat-Shamir** transform is **UF-CMA**

- **Remark**: The original proof requires to model H as an **ideal** hash function (**random oracle**)
 - It is **debatable** in the community what such a proof means in **practice**
- Can we prove security in the **plain model** (i.e., no random oracles)?
 - Many **impossibility** results for **general** ID schemes [???]
 - **Possible** for **some** classes of ID schemes assuming so-called **correlation intractability** [???]

Sufficient Criteria for Passive Security



- One can show the following criteria are **sufficient** for achieving **passive security**:
 - **Special soundness**: Given any pk and two **accepting** transcripts (α, β, γ) and $(\alpha, \beta', \gamma')$ for pk with $\beta \neq \beta'$, there is a polynomial-time algorithm **outputting** sk
 - **HVZK**: **Honest** proofs **reveal nothing** about the secret key sk

Proofs of Knowledge

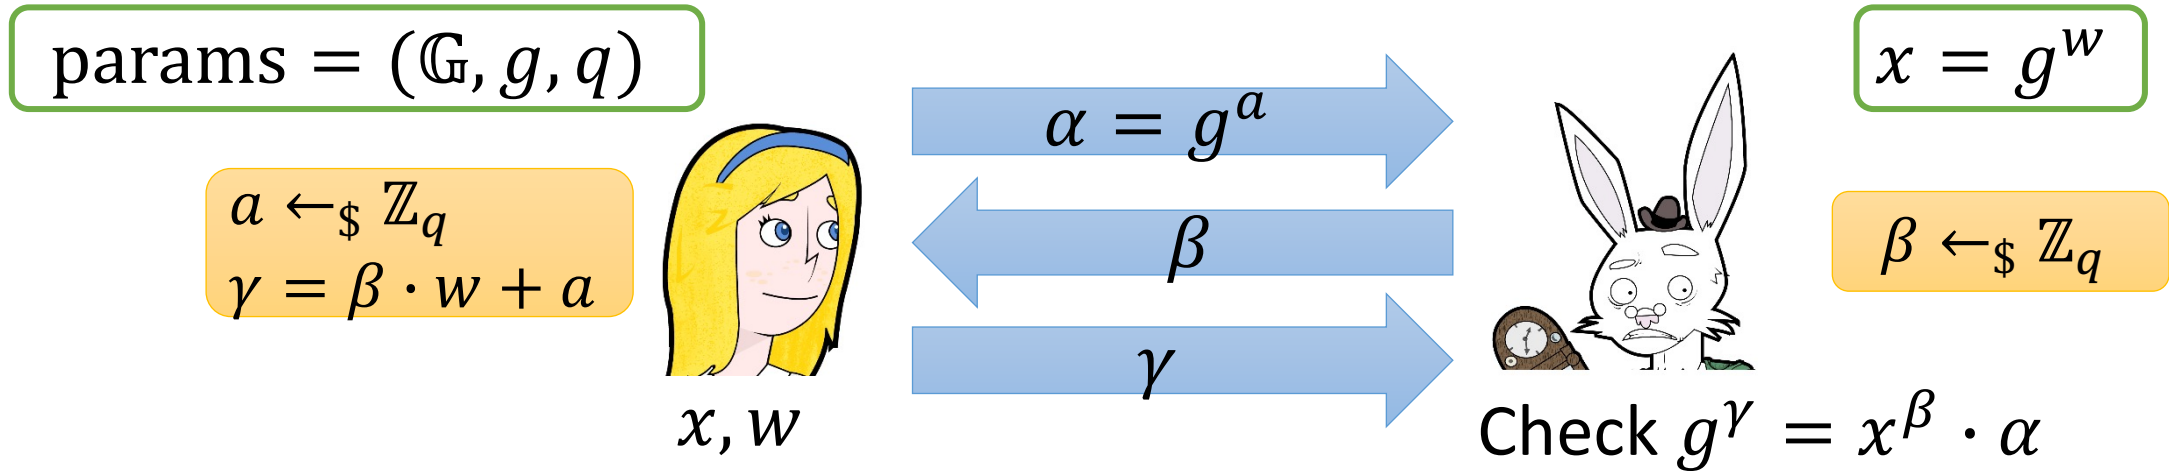
- The **special soundness** property implies that any successful prover must essentially **know the secret key**
- In fact, any such prover can be used to **extract** the secret key:
 - Run the prover upon input pk in order to obtain a transcript (α, β, γ)
 - **Rewind** the prover after it already sent α and forward it **another random challenge** β' , which yields a transcript $(\alpha, \beta', \gamma')$
 - As long as $\beta \neq \beta'$, **special soundness** allows us to obtain sk
- The above can be formalized, but the proof requires **some care**
 - Because the transcripts (α, β, γ) and $(\alpha, \beta', \gamma')$ are **correlated**

Honest-Verifier Zero-Knowledge

- How do we formalize that a transcript **reveals nothing** on sk ?
 - This is tricky: transcripts shall not reveal even **one bit** of sk
- Require that honest transcripts can be **efficiently simulated** given just pk (but not sk)
 - Whatever the verifier could compute via the protocol, he could have computed by **talking to himself** (i.e., by running the simulator)
- A canonical ID scheme is **perfect honest-verifier zero-knowledge** (HVZK) if \exists PPT \mathcal{S} such that:


$$(pk, sk, \mathcal{S}(pk)) \equiv (pk, sk, \langle \mathcal{P}(pk, sk), \mathcal{V}(pk) \rangle)$$

Canonical ID Scheme from Discrete Log



- **Special HVZK:** Upon input $pk = x$, **simulator** \mathcal{S} outputs (α, β, γ) such that $\alpha = g^\gamma / x^\beta$ and $\beta, \gamma \leftarrow_{\$} \mathbb{Z}_q$
- **Special soundness:** Assume we are given two accepting transcripts (α, β, γ) and $(\alpha, \beta', \gamma')$ for $pk = x$, with $\beta \neq \beta'$
 - This implies $g^{\gamma - \gamma'} = x^{\beta - \beta'}$
 - Thus, $w = (\gamma - \gamma') \cdot (\beta - \beta')^{-1}$ is the **discrete logarithm** of x

Let's Try the Same Idea using Lattices




- **HVZK:** Upon input $pk = (A, t)$, **simulator** \mathcal{S} outputs (α, β, γ) such that $\alpha = A \cdot \gamma - \beta \cdot t$ and $\beta \leftarrow_{\$} \mathbb{Z}_q, \gamma \leftarrow_{\$} \mathbb{Z}_q^m$
- **Special soundness:** Assume we are given two accepting transcripts (α, β, γ) and $(\alpha, \beta', \gamma')$ for $pk = (A, t)$, with $\beta \neq \beta'$
 - This implies $A \cdot (\gamma - \gamma') = (\beta - \beta') \cdot t$
 - Thus, $s = (\gamma - \gamma') \cdot (\beta - \beta')^{-1}$ is the **solution** for $A \cdot s = t$

Many Problems...

- The challenge space is **small**
 - $q \approx 2^{12}$ for **encryption**
 - $q \approx 2^{30}$ for **signatures**
 - $q \approx 2^{32}$ for **advanced applications**
- This means that a **successful prover** can just **guess** β
- The vector s we extract is **not guaranteed to be small**
 - Recall that **removing** the requirement of s being **small** makes lattice problems **trivial**
- **Solution:** Choose **small** u, β and **repeat** the protocol in **parallel**

Modified Protocol (Take 1)



- The verifier checks the above $\forall j = 1, \dots, k$ and that the coefficients of each γ_j are **small** (i.e., in $\{0,1,2\}$)
- **Special soundness:** Given $A \cdot \gamma_j = \beta_j \cdot t + \alpha_j$ and $A \cdot \gamma'_j = \beta'_j \cdot t + \alpha_j$ with $\beta_j \neq \beta'_j$, extract $s = (\gamma_j - \gamma'_j) \cdot (\beta_j - \beta'_j)^{-1}$
 - The elements of $\gamma_j - \gamma'_j$ are in $\{-2, -1, 0, 1, 2\}$, and $\beta_j - \beta'_j$ is in $\{-1, 1\}$, so s also lies in $\{-2, -1, 0, 1, 2\}$

Insecurity of the Protocol

- There are some **caveats**:
 - We **extracted** a **slightly bigger** secret
 - We need to **repeat** for $k = 128$ or $k = 256$ times
- Even worse, the protocol **does not** satisfy **HVZK**
 - Suppose that the challenge is $\beta = 1$

0 ? 1 ? 1 0 ? 0 ? ?

+

0 ? 1 ? 1 0 ? 0 ? ?

=

0 1 2 1 2 0 ? 0 ? ?

$\beta \cdot s = s$ has coefficients in $\{0,1\}$

u has coefficients in $\{0,1\}$

γ coefficients

Possible Fix?

- Maybe we can sample u from a **larger domain**?
 - Suppose that the challenge is $\beta = 1$

$$\begin{array}{cccccccccc} 0 & ? & ? & ? & 1 & ? & 0 & ? & ? & ? \\ + & & & & & & & & & \\ 0 & ? & ? & ? & 5 & ? & 0 & ? & ? & ? \\ = & & & & & & & & & \\ 0 & 4 & 2 & 3 & 6 & 5 & 0 & 2 & 4 & 1 \end{array}$$

$\beta \cdot s = s$ has coefficients in $\{0,1\}$

u has coefficients in $\{0,1,2,3,4,5\}$

γ coefficients

- Whenever a γ coefficient is 0 or 6 we know that s is 0 or 1, but the other coefficients are **hidden** (i.e., they could be **equally** 0 or 1)
- So, s **only** effects the probability that a γ coefficient is 0 or 6

Possible Fix?

- Maybe we can sample u from a **larger domain**?
 - Suppose that the challenge is $\beta = 1$

$$\begin{array}{cccccccccc} 0 & ? & ? & ? & 1 & ? & 0 & ? & ? & ? \\ + & & & & & & & & & \\ 0 & ? & ? & ? & 5 & ? & 0 & ? & ? & ? \\ = & & & & & & & & & \\ 0 & 4 & 2 & 3 & 6 & 5 & 0 & 2 & 4 & 1 \end{array}$$

$\beta \cdot s = s$ has coefficients in $\{0,1\}$

u has coefficients in $\{0,1,2,3,4,5\}$

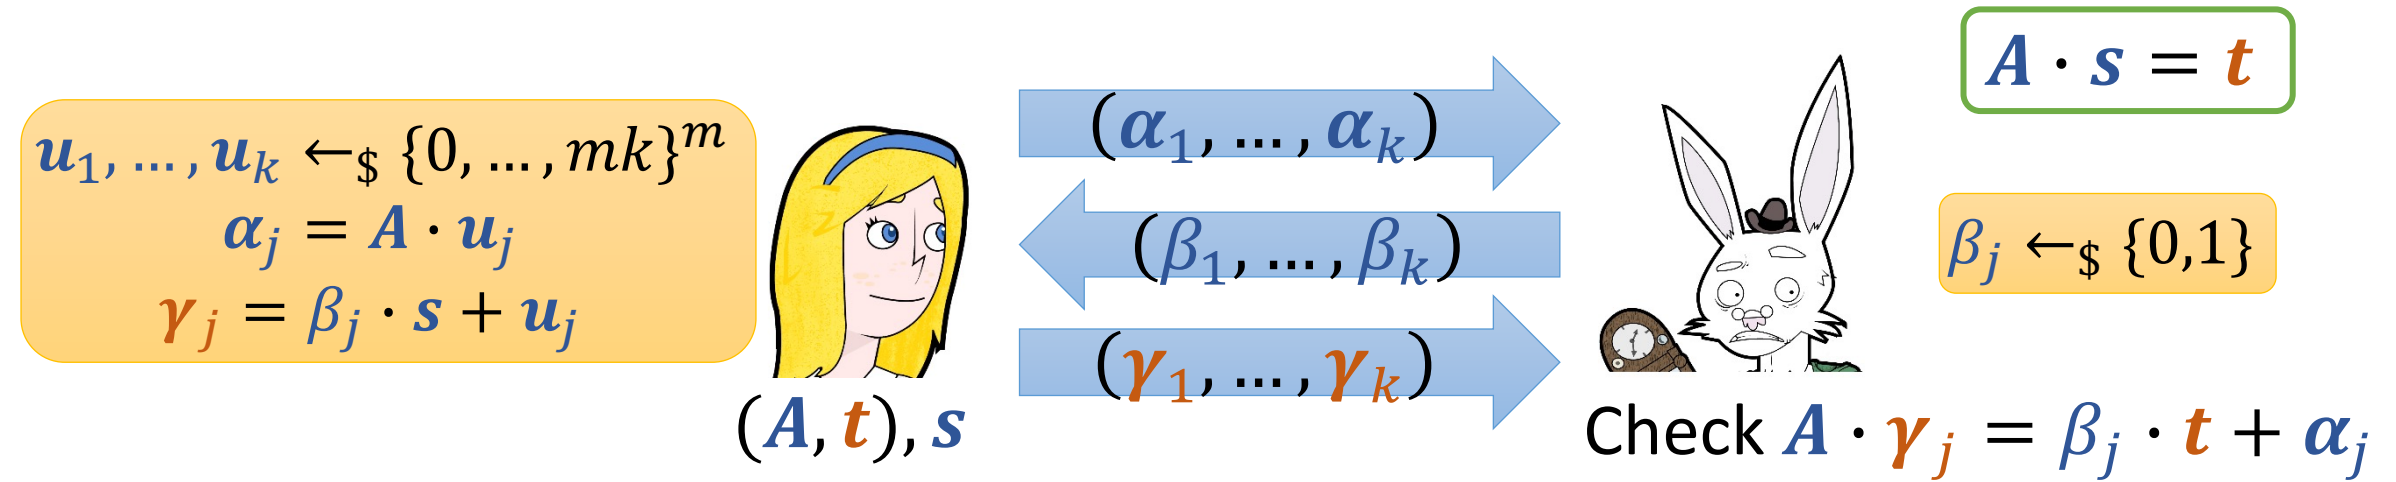
γ coefficients

- In other words, the coefficients 1,2,3,4,5 are **equally likely** to appear **regardless** of the **secret key**
- Natural idea: Send γ only when **all the coefficients** are **in this range**

In General...

- Suppose \mathbf{s} has coefficients in $\{0, 1, \dots, a\}$ and that \mathbf{u} has coefficients in $\{0, 1, \dots, b - 1\}$
 - Here, $b > a$
- Then, for all $a \leq i < b$, we have $\mathbb{P}[\mathbf{s} + \mathbf{u} = \mathbf{i}] = 1/b$
 - Moreover, there are $b - a$ such \mathbf{j} 's and thus $1 - a/b$ probability of keeping the value \mathbf{s} **secret**
- The probability that a γ coefficient is in $\{1, \dots, b - 1\}$ is $1 - 1/b$
 - The probability that they **all are** is $(1 - 1/b)^m$
 - The probability that they **all are for all** $\gamma_1, \dots, \gamma_k$ is $(1 - 1/b)^{mk}$
 - By setting $b = mk$, we get $(1 - 1/b)^{mk} \approx 1/e$

Modified Protocol (Take 2)




- The prover checks whether **any** of the coefficients contained in γ_j is 0 or $mk + 1$
 - If it is, **abort** and **restart** the protocol
- The verifier checks the above $\forall j = 1, \dots, k$ and that the coefficients of each γ_j are **small** (i.e., in $\{0, \dots, mk\}$)

Modified Protocol (Take 2)

$$u_1, \dots, u_k \leftarrow_{\$} \{0, \dots, mk\}^m$$

$$\alpha_j = A \cdot u_j$$

$$\gamma_j = \beta_j \cdot s + u_j$$




$(A, t), s$

$$(\alpha_1, \dots, \alpha_k)$$

$$(\beta_1, \dots, \beta_k)$$

$$(\gamma_1, \dots, \gamma_k)$$



Check $A \cdot \gamma_j = \beta_j \cdot t + \alpha_j$

$$A \cdot s = t$$

$$\beta_j \leftarrow_{\$} \{0, 1\}$$

- **Special soundness:** Given $A \cdot \gamma_j = \beta_j \cdot t + \alpha_j$ and $A \cdot \gamma'_j = \beta'_j \cdot t + \alpha_j$ with $\beta_j \neq \beta'_j$, extract $s = (\gamma_j - \gamma'_j) \cdot (\beta_j - \beta'_j)^{-1}$
 - The elements of $\gamma_j - \gamma'_j$ are in $\{-mk, \dots, mk\}$, and $\beta_j - \beta'_j$ is in $\{-1, 1\}$, so s also lies in $\{-mk, \dots, mk\}$
- **HVZK:** Yes, as now γ_j **never depends** on s
 - **Caveat:** What is α_j in case of **abort**?

Modified Protocol (Take 3)


$$\begin{aligned} u_1, \dots, u_k &\leftarrow_{\$} \{0, \dots, mk\}^m \\ \alpha_j &= A \cdot u_j \\ \gamma_j &= \beta_j \cdot s + u_j \end{aligned}$$



$$\alpha = H(\alpha_1, \dots, \alpha_k)$$

$$(\beta_1, \dots, \beta_k)$$

$$(\gamma_1, \dots, \gamma_k)$$



$$A \cdot s = t$$

$$\beta_j \leftarrow_{\$} \{0,1\}$$

$$\text{Check } A \cdot \gamma_j = \beta_j \cdot t + \alpha_j$$

- The verifier checks the above $\forall j = 1, \dots, k$ and that the coefficients of each γ_j are **small** (i.e., in $\{0, \dots, mk\}$)
- But now it also **additionally checks** that

$$\alpha = H(A \cdot \gamma_1 - \beta_1 \cdot t, \dots, A \cdot \gamma_k - \beta_k \cdot t)$$

- In case of **abort**, the HVZK simulator can still send a **random** α

In Practice

- The previous protocol still needs to be **repeated in parallel** $k = 128$ or 256 times
 - And this is the best one can get for **arbitrary** lattices
- However:
 - The proof size for **one equation** is roughly the same as the proof size for **many equations** (amortization with **logarithmic** growth)
 - Working with **polynomial rings** instead of \mathbb{Z}_q allows for **one-shot approximate** proofs (i.e., the coefficients of \mathbf{s} are **small**)
 - Using more **complex techniques**, one obtains **almost one-shot exact** proofs (i.e., the coefficients of \mathbf{s} are in $\{0,1\}$)

Crystals-Kyber



Regev PKE [Reg05]

- **Key Generation:** $pk = (\mathbf{A}, \mathbf{b})$ and $sk = \mathbf{s}$, where $\mathbf{b}^t = \mathbf{s}^t \cdot \mathbf{A} + \mathbf{e}^t$ and $\mathbf{s} \in \mathbb{Z}_q^n, \mathbf{A} \in \mathbb{Z}_q^{n \times m}$
- **Encryption:** The encryption of x w.r.t. pk is made of two parts
 - Ciphertext preamble $\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{r}$ for random $\mathbf{r} \in \{0,1\}^m$
 - Ciphertext payload $\mathbf{c}_1 = \mathbf{b}^t \cdot \mathbf{r} + x \cdot q/2$
 - Bob outputs $\mathbf{c}_1 - \mathbf{s}^t \cdot \mathbf{c}_0 \approx x \cdot q/2$
- **Security:** By LWE we can switch (\mathbf{A}, \mathbf{b}) with (\mathbf{A}, \mathbf{b}) for uniformly random \mathbf{b}^t
 - By the **leftover hash lemma**, we can finally replace \mathbf{c}_0 with uniformly random \mathbf{c}_0 , so that \mathbf{c}_1 hides x **information theoretically**

Dual Regev [GPV08]

- **Key Generation:** $pk = (A, u)$ and $sk = r$, where $u = A \cdot r$ and $r \in \{0,1\}^m$, $A \in \mathbb{Z}_q^{n \times m}$
- **Encryption:** The encryption of x w.r.t. pk is made of two parts
 - Ciphertext preamble $c_0 = b^t = s^t \cdot A + e^t$ for random $s \in \mathbb{Z}_q^n$
 - Ciphertext payload $c_1 = s^t \cdot u + e' + x \cdot q/2$
 - Bob outputs $c_1 - c_0 \cdot r \approx x \cdot q/2$
- **Security:** By the leftover hash lemma, we can switch u with **uniformly random u**
 - By LWE we can switch (c_0, c_1) with **uniformly random (c_0, c_1)**

Primal versus Dual

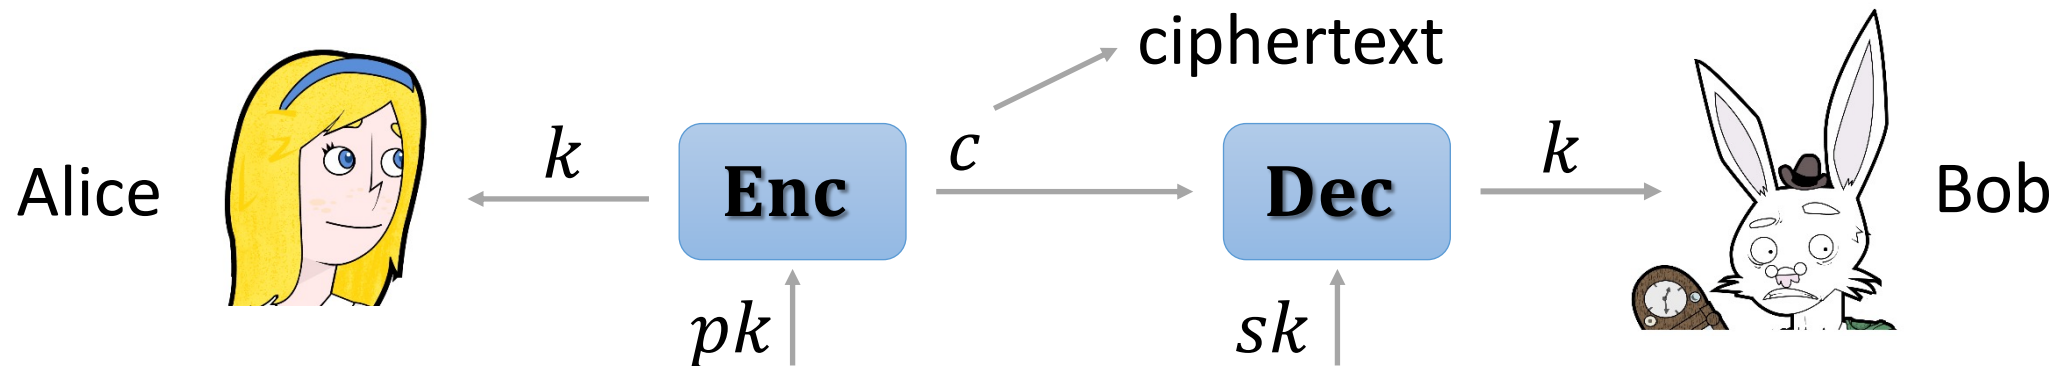
- Public key
 - Primal: pk is **pseudorandom** with **unique** sk
 - Dual: pk is **statistically random** with **many possible** sk
- Ciphertext
 - Primal: A fresh LWE sample with **many possible** coins
 - Dual: Multiple LWE samples with **unique** coins
- Security
 - Primal: Encrypting with **uniform** pk induces **random** ciphertext
 - Dual: By LWE can switch the ciphertext to **random**
- Efficiency: The matrix A can be **shared** by different users

Most Efficient [LP11]

- **Key Generation:** $pk = (A, u)$ and $sk = s$, where $u^t = s^t \cdot A + e^t$ and $s \in \chi^n, A \in \mathbb{Z}_q^{n \times n}$
- **Encryption:** The encryption of x w.r.t. pk is made of two parts
 - Ciphertext preamble $c_0 = A \cdot r + e'$ for $r \in \chi^n$
 - Ciphertext payload $c_1 = u^t \cdot r + e' + x \cdot q/2$
 - Bob outputs $c_1 - s^t \cdot c_0 \approx x \cdot q/2$
- **Security:** By LWE we can switch (A, u) with (A, u) for **uniformly random u**
 - This requires LWE with secrets from the **error distribution**
 - Next, we can replace (c_0, c_1) with **uniformly random (c_0, c_1)**

Fujisaki-Okamoto Transform

- The **FO transform** [FO99,FO13] turns **passively (IND-CPA)** secure PKE schemes into **actively (IND-CCA)** secure ones
 - The transformation requires two **hash functions** (random oracles)
 - The obtained scheme is better understood as a **key encapsulation mechanism (KEM)**

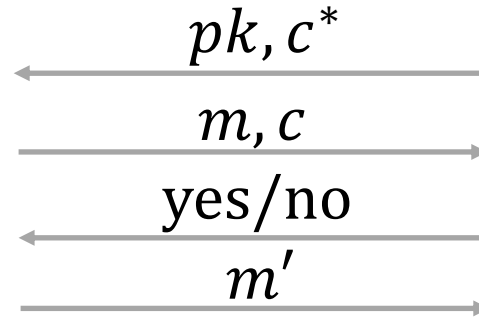


- We can combine a **KEM** with an **SKE** scheme to get a **PKE** scheme

One-Wayness of PKE



Eve



Challenger




pk, sk

$m^* \leftarrow \mathcal{M}$

$c^* \leftarrow \mathbf{Enc}(pk, m^*)$

- **OW-CPA**: PKE makes it **hard to guess** the message
 - The message is **uniformly random** and **unknown** to the attacker
- **OW-PCA**: As before but now the attacker can query a **plaintext-checking oracle** which allows to check if $\mathbf{Dec}(sk, c) = m$

Modularization of the FO Transform




- We can view FO as the **concatenation** of **two transforms** $U \circ T$
 - The first transformation takes care of **derandomization** and allows to go from **IND-CPA** to **OW-PCA**
 - The second transformation takes care of **hashing** and allows to go from **OW-PCA** to **IND-CCA**

Transformation **T**: From IND-CPA to OW-PCA



- Encryption becomes **deterministic** (the **randomness** is $\mathbf{G}(m)$)
- Decryption **re-encrypts** m' using randomness $\mathbf{G}(m')$ and outputs m' if and only if it obtains c
- **Theorem [HKK17]:** Assuming $(\mathbf{Enc}, \mathbf{Dec})$ is **IND-CPA** (**OW-CPA**), $(\mathbf{Enc}', \mathbf{Dec}')$ is **OW-PCA**

Transformation **U**: From OW-PCA to IND-CCA



- Encapsulation outputs $k = \mathbf{H}(c, m)$ and c
- Decapsulation obtains $m' = \mathbf{Dec}(sk, c)$ and outputs m'
 - Here, m' could be \perp (**explicit rejection**)
- **Theorem [HKK17]:** Assuming $(\mathbf{Enc}', \mathbf{Dec}')$ is **OW-PCA**, $(\mathbf{Encaps}, \mathbf{Decaps})$ is **IND-CCA**


Advanced Cryptographic Applications



Computing over Encrypted Data

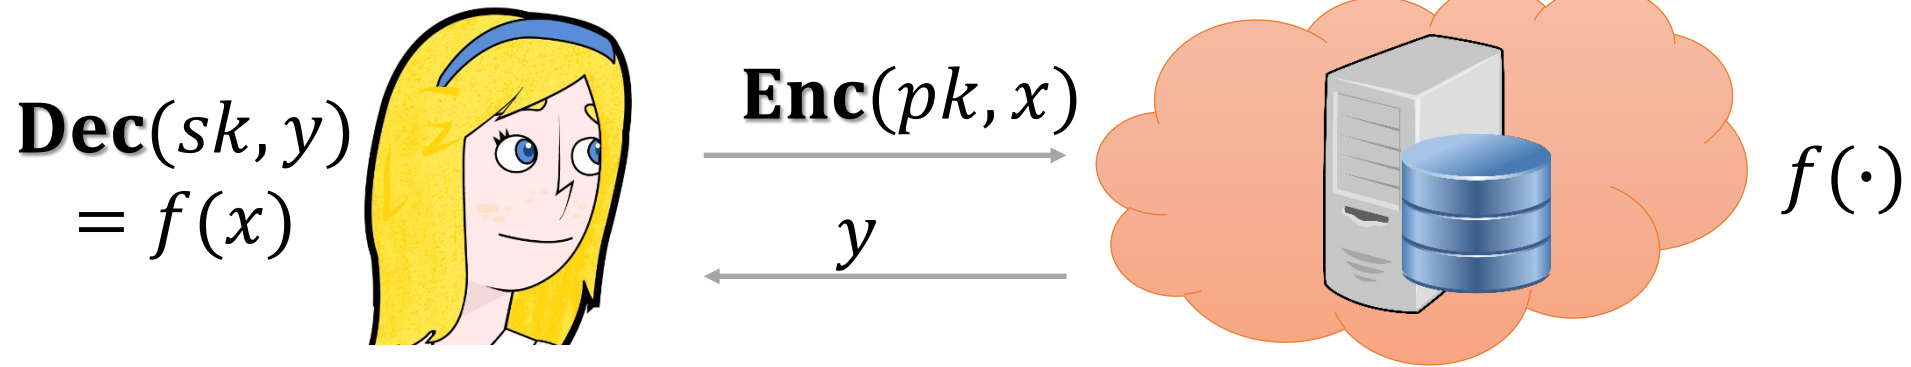
- Can we have a (public-key) encryption scheme which allows to run **computations** over **encrypted data**?
- Question dating back to the late 70s
 - Ron Rivest and "privacy homomorphisms"
- Partial solutions known
 - E.g., RSA and Elgamal enjoy limited forms of homomorphism
- First solution by Craig Gentry after 30 years
 - The "Swiss Army knife of cryptography"

Motivation: Outsourcing of Computation




- Email, web search, navigation, social networking, ...
- What about **private** x ?

Outsourcing of Computation - Privately




Wish: Homomorphic **evaluation** function:
Eval: $pk, f, \text{Enc}(pk, x) \rightarrow \text{Enc}(pk, f(x))$

Fully-Homomorphic Encryption (FHE)



$$c = \mathbf{Enc}(pk, x)$$
$$y = \mathbf{Eval}(pk, f, c)$$



Correctness:

$$\mathbf{Dec}(sk, y) = f(x)$$

Privacy:

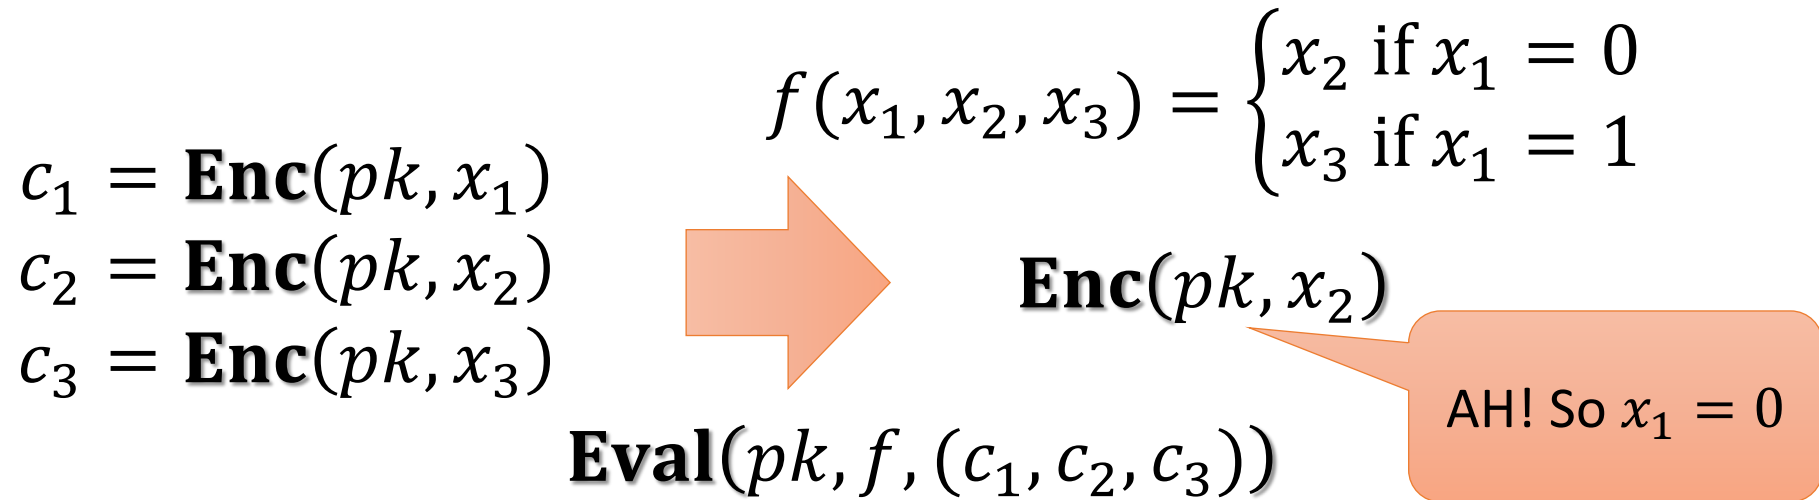
$$\mathbf{Enc}(pk, x) \approx \mathbf{Enc}(pk, 0^{|x|})$$

FHE = Correctness \forall efficient f = Correctness for universal set

Levelled FHE: Bounded depth f

- NAND
- $(+, \times)$ over a ring

A Paradox (and its Resolution)



- But remember that encryption is **randomized!**
- Output of **Eval** will look as a **fresh and random** ciphertext

Trivial FHE?

- Let $(\mathbf{KGen}, \mathbf{Enc}, \mathbf{Dec})$ be **any PKE** scheme
- Define the following **fully-homomorphic** PKE $(\mathbf{KGen}, \mathbf{Enc}, \mathbf{Eval}', \mathbf{Dec}')$:
 - $\mathbf{Eval}'(pk, \Gamma, c) = (\Gamma, c)$
 - $\mathbf{Dec}'(sk, c) = \Gamma(\mathbf{Dec}(sk, c))$

Wish: Complexity of decryption **much less** than running the circuit from scratch

The Gentry-Sahai-Waters FHE Scheme

- In what follows we will present the FHE scheme due to:
 - C. Gentry, A. Sahai, B. Waters: "Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based." CRYPTO 2013
- Based on the **Learning with Errors (LWE)** assumption
- Only achieves **levelled homomorphism**
 - But can be **bootstrapped** to **full homomorphism** using a trick by Gentry (under additional assumptions)
- Plaintext space will be $\mathbb{Z}_q = [-q/2, q/2)$, for a large prime q
 - For simplicity let us write $[a]_q$ for $a \bmod q$

Eigenvectors Method (Basic Idea)

- Let C_1 and C_2 be matrices for **eigenvector** \vec{s} , and **eigenvalues** x_1, x_2 (i.e., $\vec{s} \times C_i = x_i \cdot \vec{s}$)
 - $C_1 + C_2$ has eigenvalue $x_1 + x_2$ w.r.t. \vec{s}
 - $C_1 \times C_2$ has eigenvalue $x_1 \cdot x_2$ w.r.t. \vec{s}
- Idea: Let C be the ciphertext, \vec{s} be the secret key and x be the plaintext (say over \mathbb{Z}_q)
 - Homomorphism for **addition/multiplication**
 - But **insecure**: Easy to compute eigenvalues

Approximate Eigenvectors (1/2)

- Approximate variant: $\vec{s} \times C = x \cdot \vec{s} + \vec{e} \approx x \cdot \vec{s}$
 - Decryption **works** as long as $\|\vec{e}\|_\infty \ll q$

$$\begin{aligned} \vec{s} \times C_1 &= x_1 \cdot \vec{s} + \vec{e}_1 & \vec{s} \times C_2 &= x_2 \cdot \vec{s} + \vec{e}_2 \\ \|\vec{e}_1\|_\infty &\ll q & \|\vec{e}_2\|_\infty &\ll q \end{aligned}$$

- Goal: Define **homomorphic** operations

$$C_{\text{add}} = C_1 + C_2:$$

$$\begin{aligned} \vec{s} \times (C_1 + C_2) &= \vec{s} \times C_1 + \vec{s} \times C_2 \\ &= x_1 \cdot \vec{s} + \vec{e}_1 + x_2 \cdot \vec{s} + \vec{e}_2 \\ &= (x_1 + x_2) \cdot \vec{s} + (\vec{e}_1 + \vec{e}_2) \end{aligned}$$

Noise **grows** a little!

Approximate Eigenvectors (2/2)

- Approximate variant: $\vec{s} \times C = x \cdot \vec{s} + \vec{e} \approx x \cdot \vec{s}$
 - Decryption **works** as long as $\|\vec{e}\|_\infty \ll q$

$$\begin{aligned} \vec{s} \times C_1 &= x_1 \cdot \vec{s} + \vec{e}_1 & \vec{s} \times C_2 &= x_2 \cdot \vec{s} + \vec{e}_2 \\ \|\vec{e}_1\|_\infty &\ll q & \|\vec{e}_2\|_\infty &\ll q \end{aligned}$$

- Goal: Define **homomorphic** operations

$$\begin{aligned} C_{\text{mult}} &= C_1 \times C_2: \\ \vec{s} \times (C_1 \times C_2) &= (x_1 \cdot \vec{s} + \vec{e}_1) \times C_2 \\ &= x_1 \cdot (x_2 \cdot \vec{s} + \vec{e}_2) + \vec{e}_1 \times C_2 \\ &= x_1 \cdot x_2 \cdot \vec{s} + (x_1 \cdot \vec{e}_2 + \vec{e}_1 \times C_2) \end{aligned}$$

Noise **grows!**
Needs to be
small!

Shrinking Gadgets

- Write entries in C using **binary decomposition**; e.g.

$$C = \begin{bmatrix} 3 & 5 \\ 1 & 4 \end{bmatrix} \pmod{8} \xrightarrow{\text{yields}} \text{bits}(C) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \\ 1 & 1 \\ 0 & 1 \\ 0 & 0 \\ 1 & 0 \end{bmatrix} \pmod{8}$$

- **Reverse** operation:


$$C = G \times G^{-1}(C) = \begin{bmatrix} 2^{N-1} & \dots & 2 & 1 & 0 & \dots & 0 & 0 \\ 0 & \dots & 0 & 0 & 2^{N-1} & \dots & 2 & 1 \end{bmatrix} \times \text{bits}(C)$$

$\xleftarrow{k \cdot N = k \lceil \log q \rceil}$

$\Rightarrow \vec{s} \times C = \vec{s} \times G \times G^{-1}(C)$

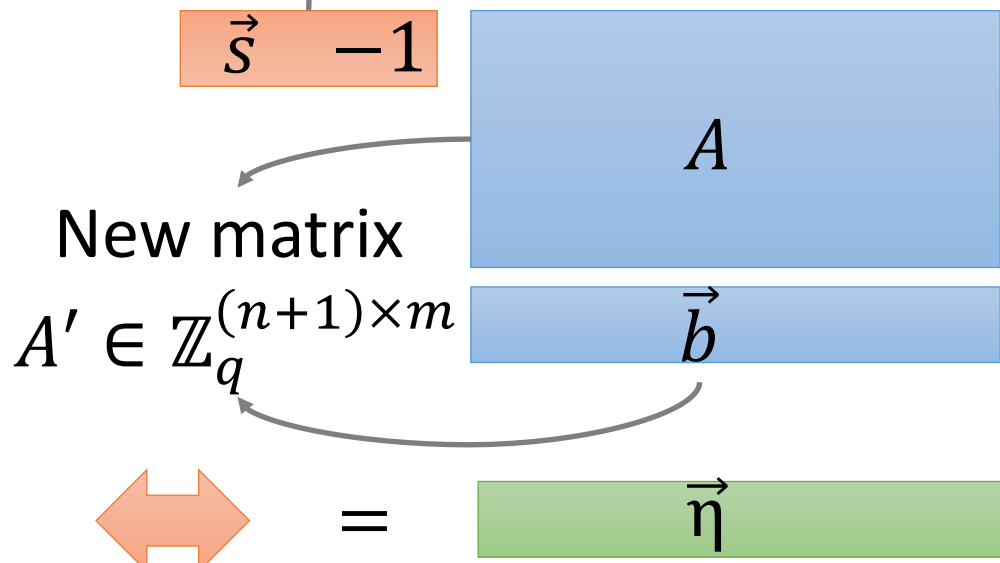
LWE – Rearranging Notation

$$\vec{b} = \vec{s} \times A + \vec{\eta}$$



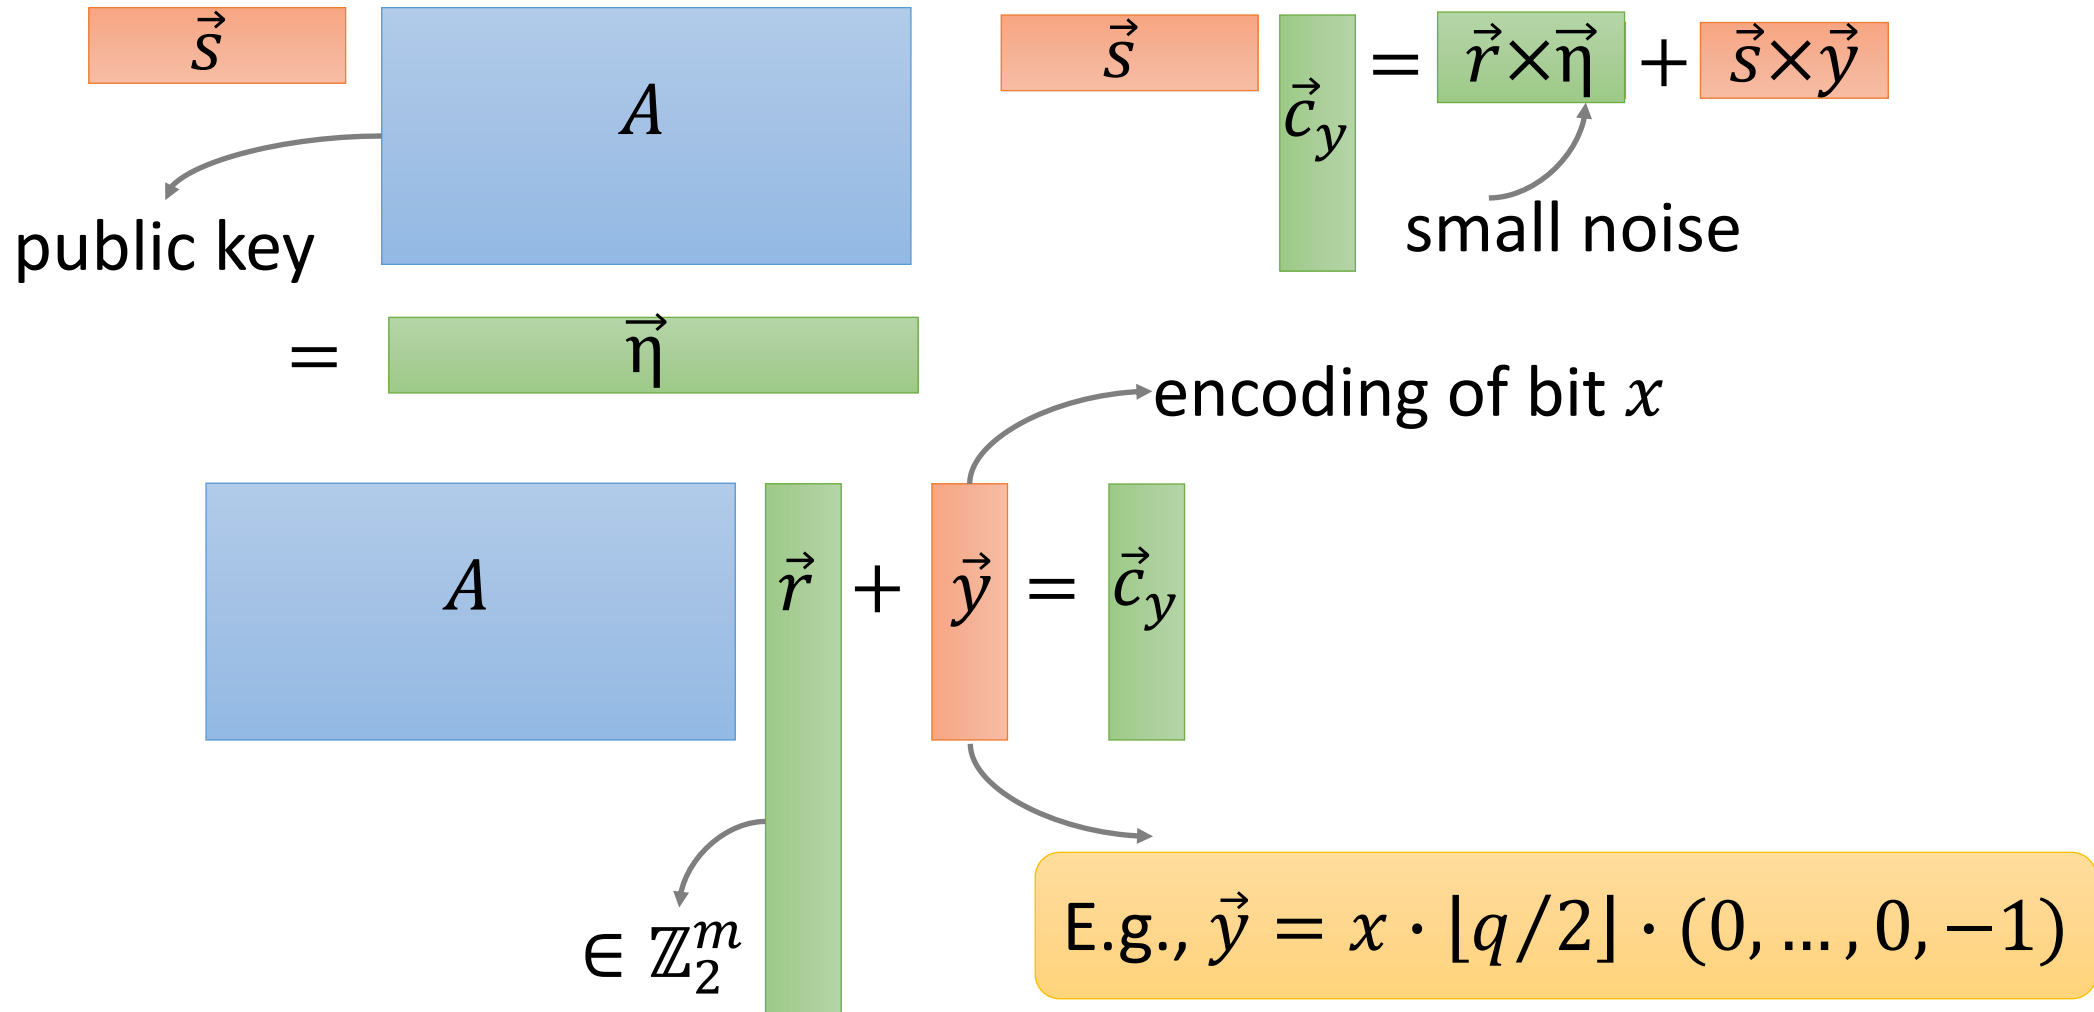
Small noise $\in \mathbb{Z}_q^m$
 $|\eta_i| \leq \alpha q; \alpha \ll 1$

New secret $\vec{s} \in \mathbb{Z}_q^{n+1}$

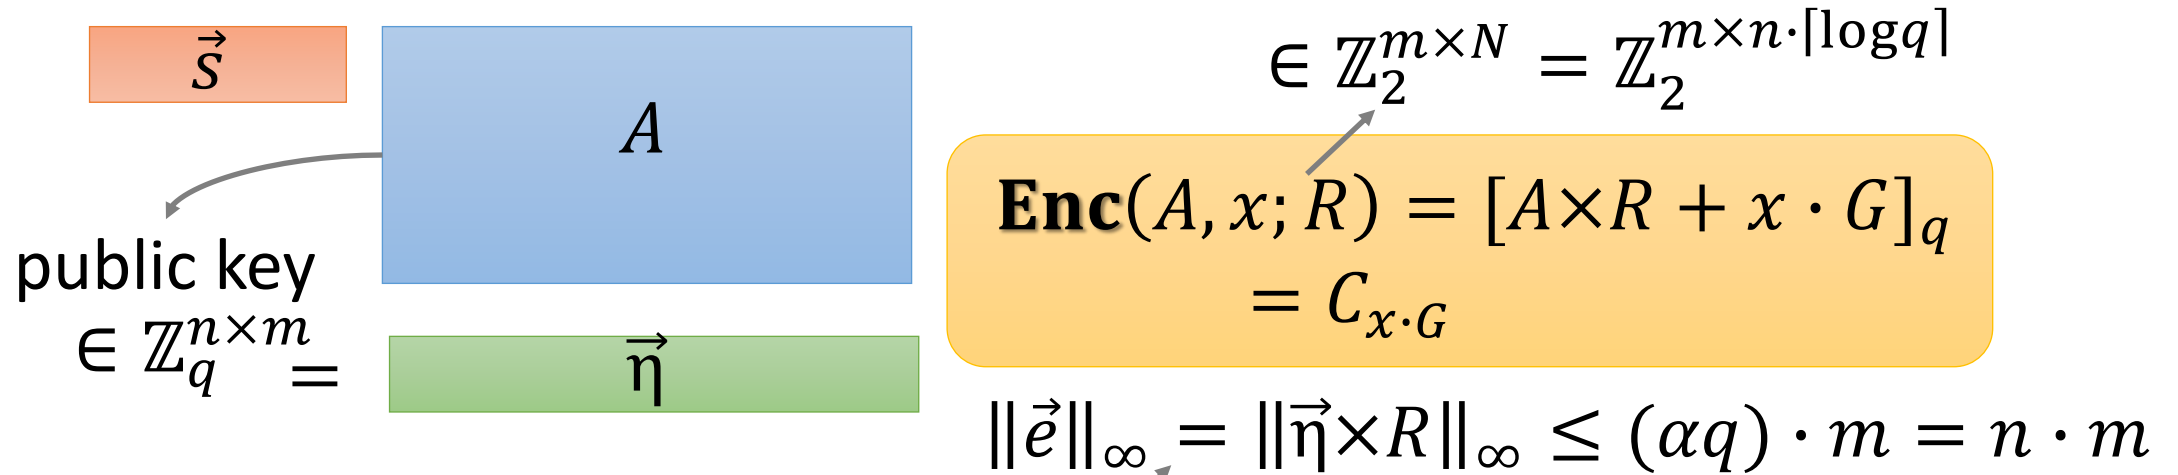


$$\text{LWE: } A' = (A || \vec{b}) \approx_c \mathbf{U}_q^{(n+1) \times m}$$

Regev PKE – Pictorially



The GSW Scheme



Invariant: $\vec{s} \times C = \vec{e} + x \cdot \vec{s} \times G$

$$\begin{aligned} \text{Dec}(\vec{s}, C) &= \vec{s} \times C \times G^{-1}((0, \dots, 0, -\lfloor q/2 \rfloor)) \\ &= \vec{e} \times G^{-1}(\dots) + x \cdot \vec{s} \times G \times G^{-1}((0, \dots, 0, -\lfloor q/2 \rfloor)) \\ &= \vec{e} \times G^{-1}(\dots) + \lfloor q/2 \rfloor \cdot x = z \end{aligned}$$



Output: $0 \Leftrightarrow |z| < q/4$

The GSW Scheme – Homomorphism


$$\text{Invariant: } \vec{s} \times C = \vec{e} + x \cdot \vec{s} \times G$$

$$C_{\text{mult}} = C_1 \times G^{-1}(C_2)$$

$$\begin{aligned} \vec{s} \times C_1 \times G^{-1}(C_2) &= (\vec{e}_1 + x_1 \cdot \vec{s} \times G) \cdot G^{-1}(C_2) \\ &= \vec{e}_1 \times G^{-1}(C_2) + x_1 \cdot \vec{s} \times G \times G^{-1}(C_2) \\ &= \vec{e}_1 \times G^{-1}(C_2) + x_1 \cdot \vec{s} \times C_2 \\ &= \vec{e}_1 \times G^{-1}(C_2) + x_1 \cdot (\vec{e}_2 + x_2 \cdot \vec{s} \times G) \\ &= (\vec{e}_1 \times G^{-1}(C_2) + x_1 \cdot \vec{e}_2) + x_1 x_2 \cdot \vec{s} \times G \\ &= \vec{e}_{\text{mult}} + x_1 x_2 \cdot \vec{s} \times G \end{aligned}$$

$$\|\vec{e}_{\text{mult}}\|_{\infty} \leq N \cdot \|\vec{e}_1\|_{\infty} + \|\vec{e}_2\|_{\infty} \leq (N + 1) \cdot \max\{\|\vec{e}_1\|, \|\vec{e}_2\|\}$$

The GSW Scheme – Correctness



$$\|\vec{e}_{\text{out}}\|_{\infty} \leq (N + 1)^{\tau+1} m \cdot \alpha q$$

Correctness:

$$n \cdot m \cdot (N + 1)^{\tau+1} < q/4$$

$$\|\vec{e}_{i+1}\|_{\infty} \leq (N + 1) \|\vec{e}_i\|_{\infty}$$

$$\|\vec{e}_{\text{in}}\|_{\infty} \leq m \cdot n = m \cdot \alpha q$$

The GSW Scheme – Semantic Security


- Similar as in the proof of Regev PKE
- Using LWE we move to a **mental experiment** with $A \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$
- Hence, by the **leftover hash lemma**, with $m = \Theta(n \log q)$, the statistical distance between $(A, A \times \vec{r})$ and uniform is negligible
 - By a **hybrid argument** over the columns of R , it follows that the statistical distance between $(A, A \times R)$ and uniform is also negligible
 - Thus, the ciphertext **statistically hides** the plaintext

The GSW Scheme – Parameters

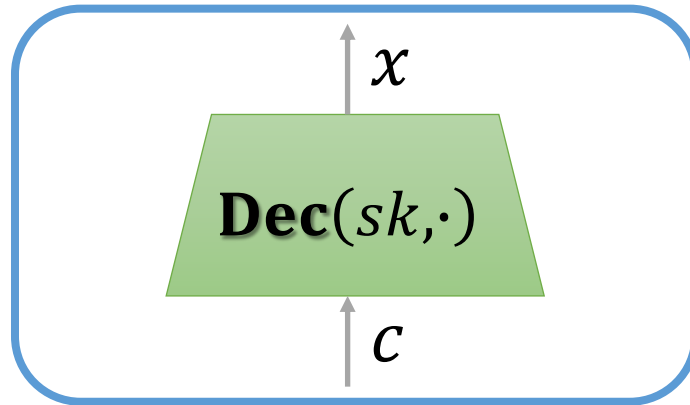
- **Correctness** requires $n \cdot m \cdot (N + 1)^{\tau+1} < q/4$
- **Security** requires $m = \Theta(n \log q)$, e.g. $m \geq 1 + 2n(2 + \log q)$
- **Hardness** of LWE requires $q \leq 2^{n^\epsilon}$ for $\epsilon < 1$
 - Substituting we get $q > (2n \log q)^{\tau+3}$
 - And thus $n^\epsilon > (\tau + 3)(\log n + \log \log q + 1)$ which for large τ, n yields $n^\epsilon > 2\tau \log n$
 - So we set $n = \max(\lambda, \lceil 4\tau/\epsilon \log \tau^{1/\epsilon} \rceil)$, $q = \lceil 2^{n^\epsilon} \rceil$, $m = O(n^{1+\epsilon})$, and $\alpha = n/q = n \cdot 2^{-n^\epsilon}$
- Hence, the size of ciphertexts is polynomial in λ, τ thus yielding a **weakly-compact** FHE

Increasing the Homomorphic Capacity

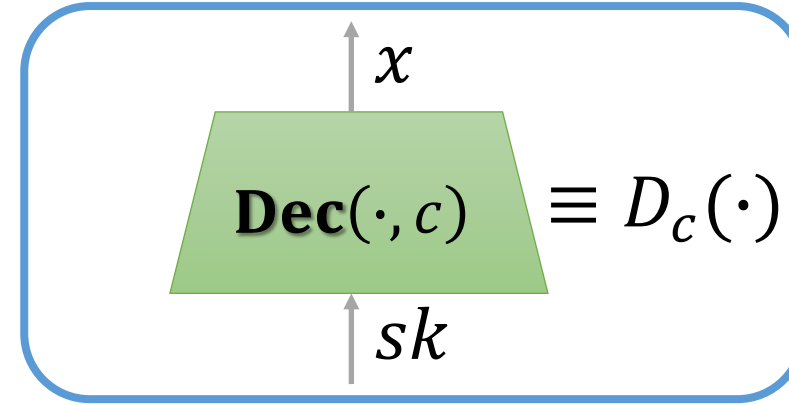
- The only way to increase the homomorphic capacity of GSW is to pick **larger parameters**
- This dependence can be **broken** using a trick by Gentry
- Main idea: Do a few operations, then **switch keys**



How to Switch Keys



Decryption circuit



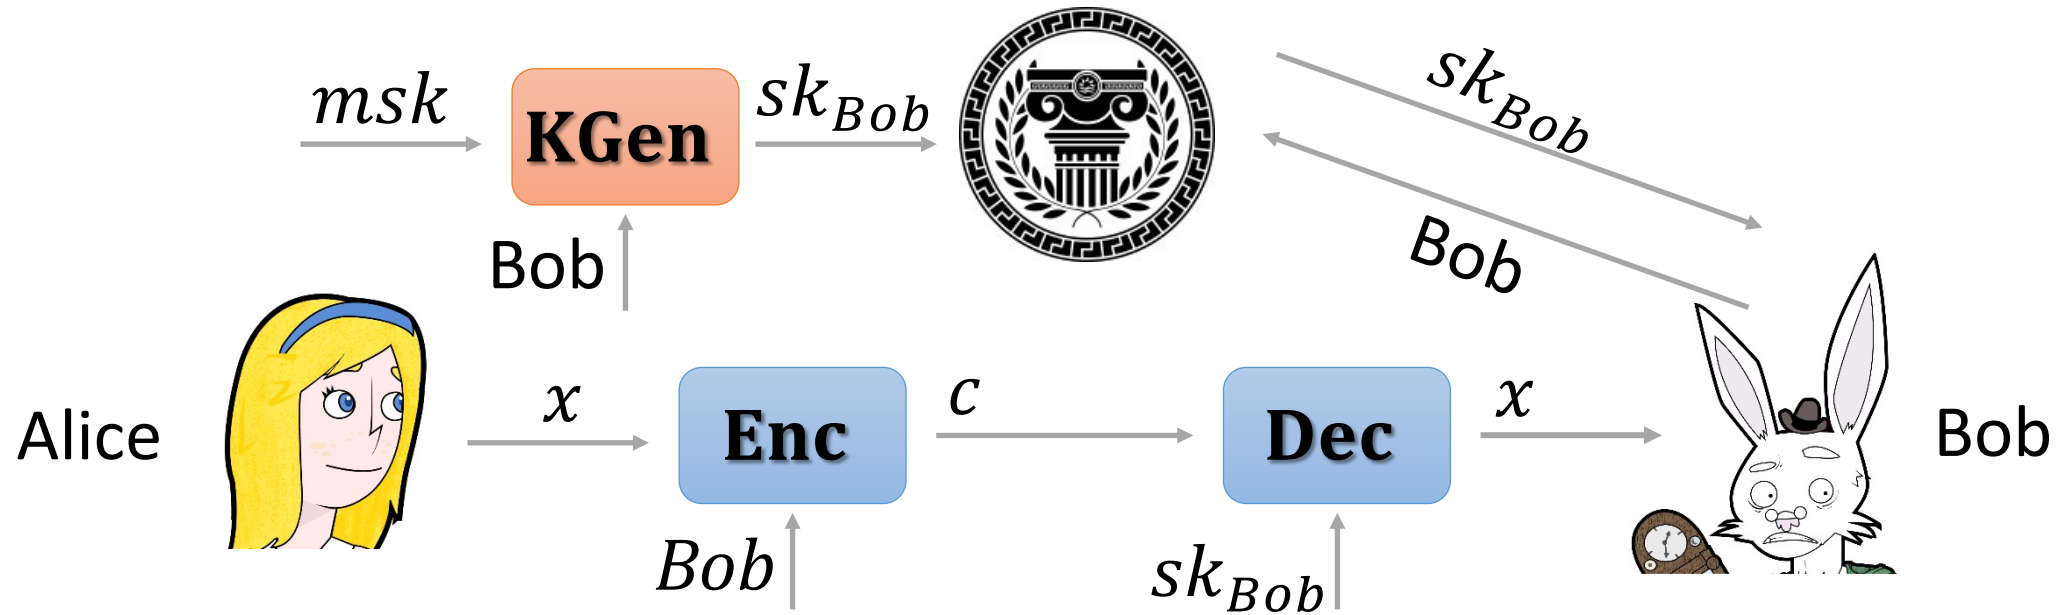
Dual view

$$\begin{aligned} \mathbf{Eval}_{pk'}(D_c, aux) &= \mathbf{Eval}_{pk'}(D_c, \mathbf{Enc}_{pk'}(sk)) \\ &= \mathbf{Enc}_{pk'}(D_c(sk)) \\ &= \mathbf{Enc}_{pk'}(x) \end{aligned}$$

Circular Security

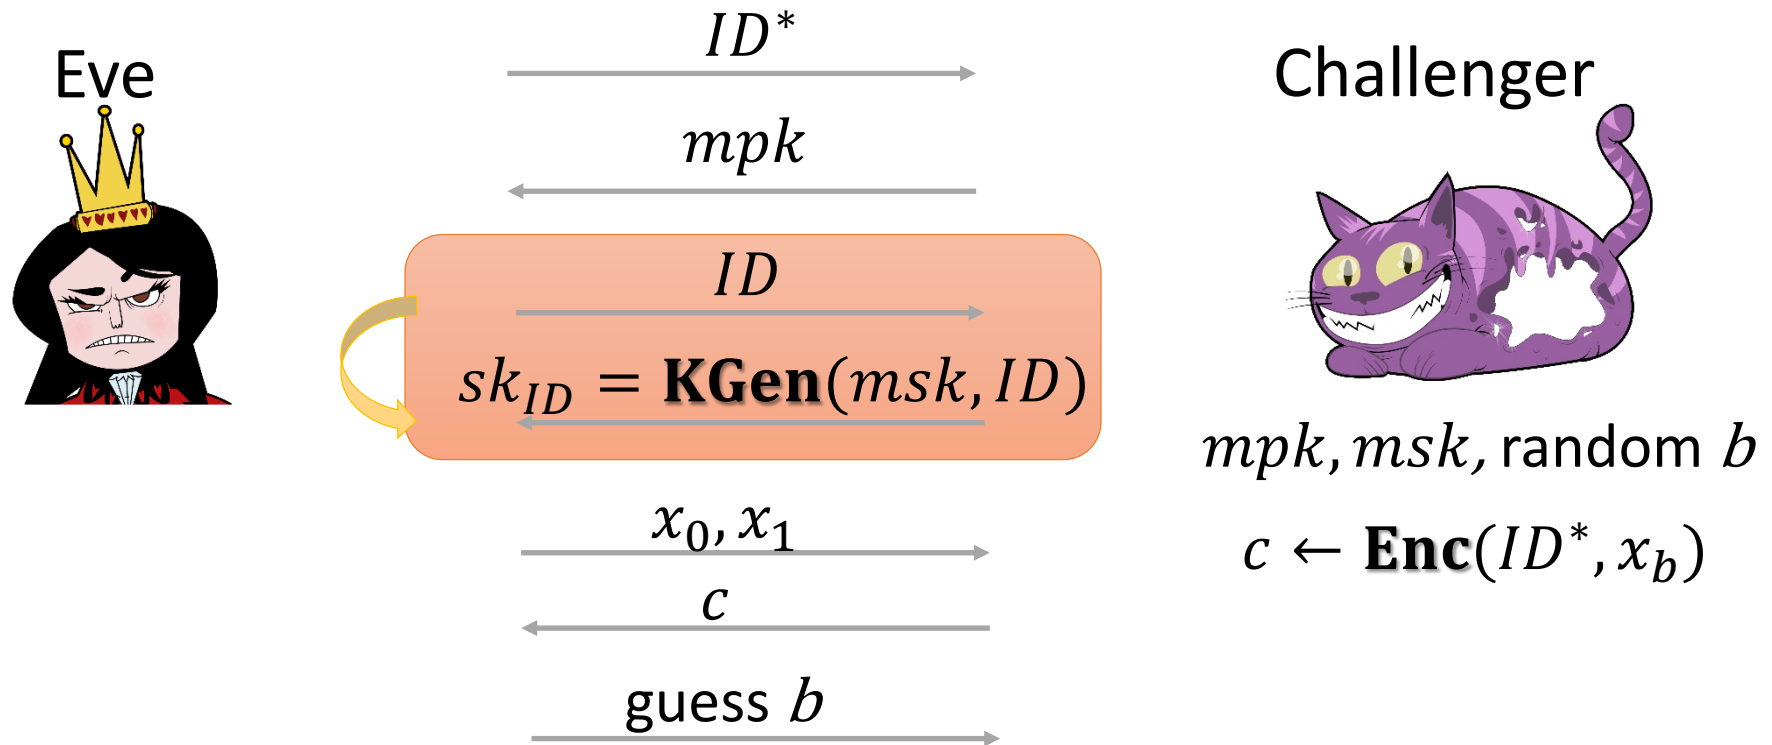
- The above scheme is **compact**, but **not fully homomorphic**, as we need a pair of keys **for each level** in the circuit
- A natural idea is to use a **single pair** (pk, sk) and include in pk' a ciphertext $\vec{c}^* \leftarrow_{\$} \mathbf{Enc}(pk, sk)$
 - Correctness still holds for this variant, but the reduction to **semantic security breaks**
- Workaround: Assume **circular security**
 - I.e., $\mathbf{Enc}(pk, 0) \approx_c \mathbf{Enc}(pk, 1)$ even given $\vec{c}^* \leftarrow_{\$} \mathbf{Enc}(pk, sk)$
 - GSW is **conjectured** to have this property, but no proof of this fact is currently known

Identity-Based Encryption



- **Postulated** by Shamir in 1984 [Sha84]
 - Avoids the need of **certificates**
 - Introduces the so-called **key escrow** problem
- First **realization** by Boneh and Franklin in 2001 [BF01]

Selective Security of IBE



- Every **selectively** secure IBE is also **fully** secure with an **exponential** loss in the parameters
 - Also, general **transformations** are known

Warm-up Construction [CHKP10]

- **Public parameters:** $mpk = (A_0, A_1^0, A_1^1, A_2^0, A_2^1, u)$
 - Assume, for simplicity, $|ID| = 2$
- **Master secret key:** Trapdoor for A_0
 - Secret key for identity $ID = 01$: **Short vector** s s.t. $F_{01} \cdot s = u \pmod q$, where $F_{01} = [A_0 | A_1^0 | A_1^1]$
 - Note: A trapdoor for A_0 **implies** a trapdoor for F_{01}
- **Encryption: Dual** Regev encryption of x w.r.t. matrix F_{01}
 - The ciphertext is $c_0^t = r^t \cdot F_{01} + e^t$ and $c_1 = r^t \cdot u + e' + x \cdot q/2$
 - Bob outputs $c_1 - c_0^t \cdot s \approx x \cdot q/2$

Simulation

- Assume the **challenge** identity is $ID^* = 11$
 - The reduction **can't know** the secret key for ID^*
- Choose A_0, A_1^1, A_2^1 uniformly at **random**, but sample A_1^0, A_2^0 with the corresponding **trapdoors**
- The reduction can derive trapdoors for $F_{00} = [A_0 | A_1^0 | A_2^0]$, $F_{01} = [A_0 | A_1^0 | A_2^1]$, and $F_{10} = [A_0 | A_1^1 | A_2^0]$ but not for $F_{11} = [A_0 | A_1^1 | A_2^1]$
 - This allows the reduction to **simulate** key extraction queries while **embedding** the LWE challenge in the simulation

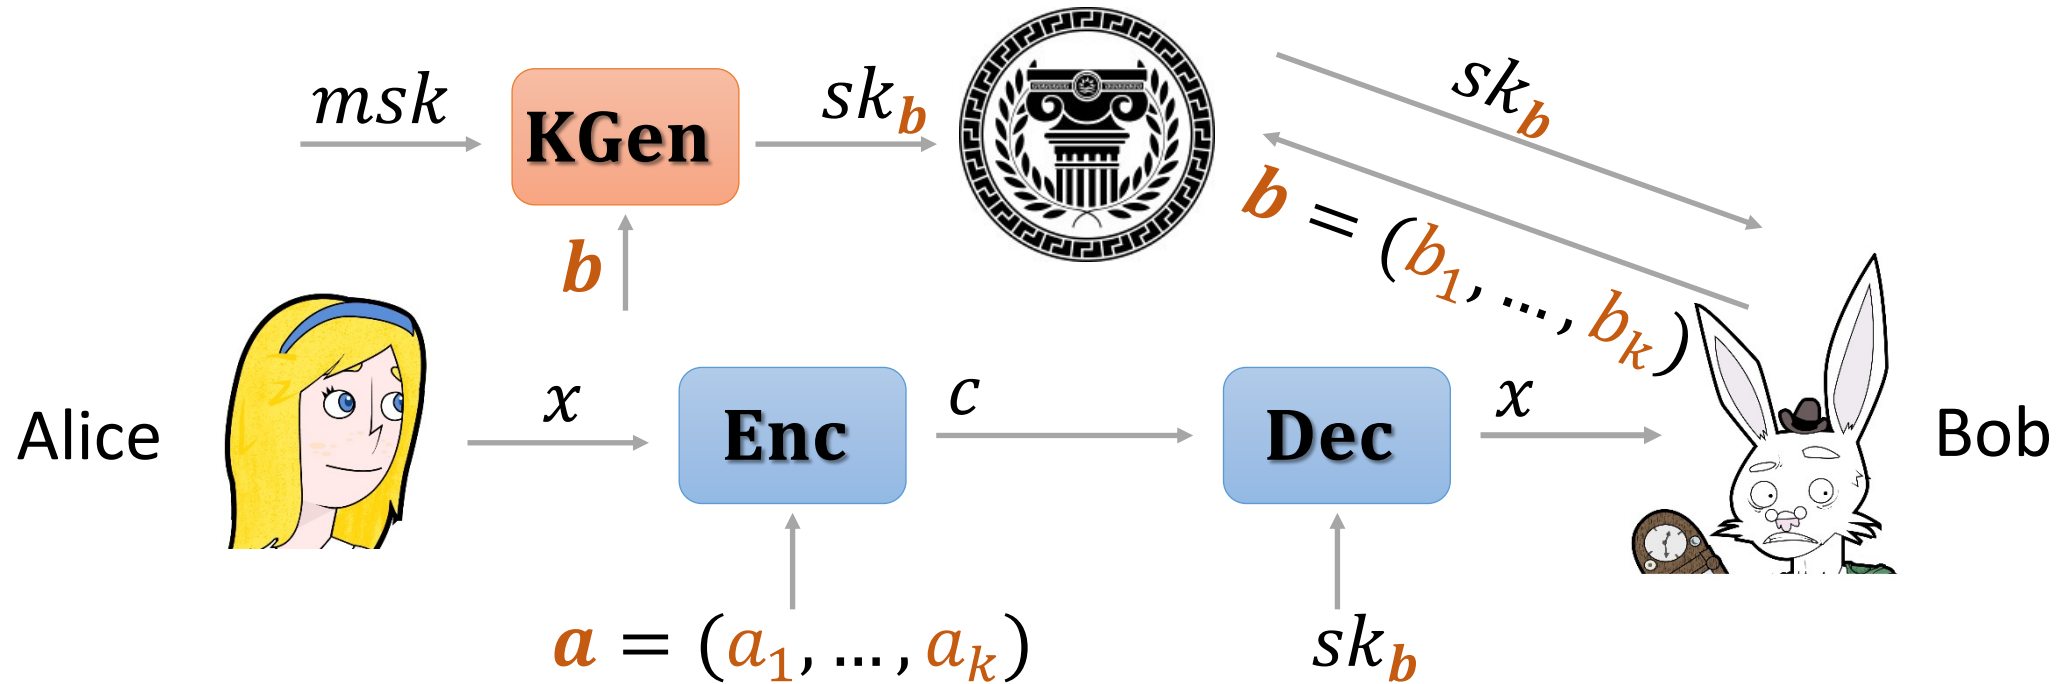
A More Efficient Construction [ABB10]

- **Public parameters:** $mpk = (A_0, A_1, G, u)$
- **Master secret key:** Trapdoor for A_0
 - Secret key for identity ID : **Short vector** s s.t. $F_{ID} \cdot s = u \pmod q$, where $F_{ID} = [A_0 | A_1 + ID \cdot G]$
 - As before, a trapdoor for A_0 **implies** a trapdoor for F_{ID}
- **Encryption: Dual** Regev encryption of x w.r.t. matrix F_{ID}
 - The ciphertext is $c_0^t = r^t \cdot F_{ID} + e^t$ and $c_1 = r^t \cdot u + e' + x \cdot q/2$
 - Bob outputs $c_1 - c_0^t \cdot s = r^t \cdot u + e' + x \cdot q/2 - r^t \cdot F_{ID} \cdot s + e^t \cdot s = r^t \cdot u + e' + x \cdot q/2 - r^t \cdot u + e^t \cdot s \approx x \cdot q/2$

Simulation Revisited

- Assume the **challenge** identity is ID^*
 - The reduction **can't know** the secret key for ID^*
- The reduction does **not** know a trapdoor for A_0 , but it knows a trapdoor for the gadget matrix G
- Let $A_1 = [A_0 \cdot R - ID^* \cdot G]$, where R is **random** and **low-norm**
 - This is **indistinguishable** from the real A_1
- Note that $F_{ID} = [A_0 | A_0 \cdot R + (ID - ID^*) \cdot G]$
 - Using the technique of [MP12], we can **derive** a trapdoor for F_{ID} given a trapdoor for A_0
 - This allows to **simulate** key extraction queries for all $ID \neq ID^*$
 - The LWE challenge can be **embedded** as before

Inner-product Encryption [KSW08]

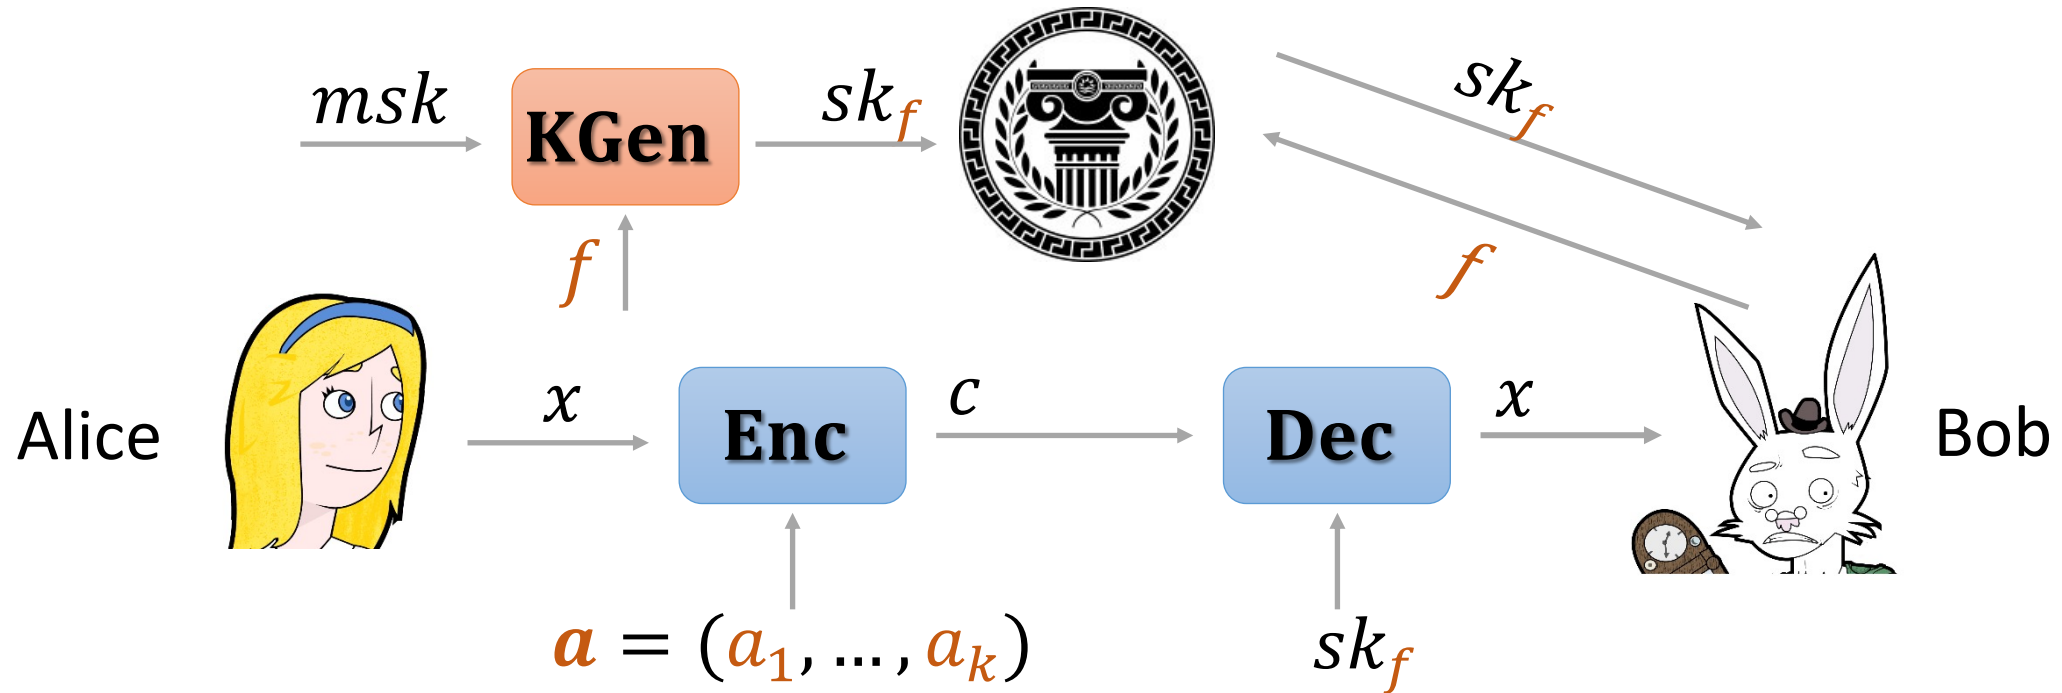


- Decryption reveals x **if and only if** $\langle a, b \rangle = 0$
 - Here, we can also be interested in **attributes privacy**
- Can be used to obtain **predicate encryption** for polynomial evaluation, CNFs/DNFs of bounded degree, and **fuzzy** IBE

Generalizing to Inner Products [AFV11]

- **Public parameters:** $mpk = (A, A_1, \dots, A_k, G, u)$
- **Master secret key:** Trapdoor for A
 - Secret key for b : **Short vector** s_b s.t. $F_b \cdot s_b = u \pmod q$, where $F_b = [A \mid \sum_i b_i \cdot A_i]$
- **Encryption: Dual** Regev encryption of x w.r.t. matrix A
 - The ciphertext is $c_0^t = r^t \cdot A + e^t$, $c' = r^t \cdot u + e' + x \cdot q/2$, and $c_i^t = r^t \cdot (A_i + a_i \cdot G) + e_i^t$ (so it indeed hides a)
 - Bob sets $c_b = \sum_i b_i \cdot c_i = r^t \cdot (\sum_i b_i \cdot A_i + \sum_i a_i \cdot b_i \cdot G) + \sum_i b_i \cdot e_i$ which equals $r^t \cdot \sum_i b_i \cdot A_i + \sum_i b_i \cdot e_i$
 - Hence, $[c_0 \mid c_b] \approx r^t \cdot [A \mid \sum_i b_i \cdot A_i]$ is a dual Regev ciphertext
 - Bob outputs $c' - c_0^t \cdot s_b - c_b^t \cdot s_b \approx x \cdot q/2$

Attribute-based Encryption [SW04]



- Decryption reveals x **if and only if** $f(a) = 0$
 - Here, we are not interested in **attributes privacy**
- Plenty of applications for **privacy-preserving data mining** and in cryptography for **big data**

Handling Multiplications [BGG+14]

- Let $\mathbf{c}_1^t = \mathbf{r}^t \cdot (\mathbf{A}_1 + a_1 \cdot \mathbf{G}) + \mathbf{e}_1^t$ and $\mathbf{c}_2^t = \mathbf{r}^t \cdot (\mathbf{A}_2 + a_2 \cdot \mathbf{G}) + \mathbf{e}_2^t$
- Want: $\mathbf{c}_{12}^t = \mathbf{r}^t \cdot (\mathbf{A}_{12} + a_1 \cdot a_2 \cdot \mathbf{G}) + \mathbf{e}_{12}^t$
 - Compute $(\mathbf{A}_1 + a_1 \cdot \mathbf{G}) \cdot \mathbf{G}^{-1}(-\mathbf{A}_2) = \mathbf{A}_1 \cdot \mathbf{G}^{-1}(-\mathbf{A}_2) - a_1 \cdot \mathbf{A}_2$
 - Compute $(\mathbf{A}_2 + a_2 \cdot \mathbf{G}) \cdot a_1 = a_1 \cdot \mathbf{A}_2 + a_1 \cdot a_2 \cdot \mathbf{G}$
 - The **difference** is $\mathbf{A}_{12} + a_1 \cdot a_2 \cdot \mathbf{G}$
- So, we let $\mathbf{c}_{12}^t = \mathbf{c}_1^t \cdot \mathbf{G}^{-1}(-\mathbf{A}_2) + \mathbf{c}_2^t \cdot a_1$
 - $\mathbf{G}^{-1}(-\mathbf{A}_2)$ and a_1 are **small** and **do not effect noise**
 - As usual, additionally let $\mathbf{c}_0^t = \mathbf{r}^t \cdot \mathbf{A} + \mathbf{e}^t$, $\mathbf{c}' = \mathbf{r}^t \cdot \mathbf{u} + \mathbf{e}' + x \cdot q/2$
 - If $a_1 \cdot a_2 = 0$, then $[\mathbf{c}_0 | \mathbf{c}_{12}] \approx \mathbf{r}^t \cdot [\mathbf{A} | \mathbf{A}_{12}]$
 - The secret key is a **short vector** \mathbf{s}_{12} s.t. $[\mathbf{A} | \mathbf{A}_{12}] \cdot \mathbf{s}_{12} = \mathbf{u} \bmod q$
 - Bob outputs $\mathbf{c}' - \mathbf{c}_0^t \cdot \mathbf{s}_{12} - \mathbf{c}_{12}^t \cdot \mathbf{s}_{12} \approx x \cdot q/2$