The Learning With Errors Problem

From Periodic Hardness to Geometric Hardness

Shreya Sharma

Euler's Circle

July 2026

A Little History First

- **1976–1977:** Diffie–Hellman and RSA launch public-key cryptography, resting on two number-theoretic assumptions: factoring is hard, and discrete logarithms are hard.

A Little History First

- **1976–1977:** Diffie–Hellman and RSA launch public-key cryptography, resting on two number-theoretic assumptions: factoring is hard, and discrete logarithms are hard.
- **1994:** Peter Shor shows a quantum computer solves both in polynomial time.

A Little History First

- **1976–1977:** Diffie–Hellman and RSA launch public-key cryptography, resting on two number-theoretic assumptions: factoring is hard, and discrete logarithms are hard.
- **1994:** Peter Shor shows a quantum computer solves both in polynomial time.
- **1996:** Miklós Ajtai discovers the first worst-case-to-average-case reduction for a lattice problem: a random instance is, provably, at least as hard as the hardest instance. This is the theoretical basis of lattice-based cryptography.

A Little History First

- **1976–1977:** Diffie–Hellman and RSA launch public-key cryptography, resting on two number-theoretic assumptions: factoring is hard, and discrete logarithms are hard.
- **1994:** Peter Shor shows a quantum computer solves both in polynomial time.
- **1996:** Miklós Ajtai discovers the first worst-case-to-average-case reduction for a lattice problem: a random instance is, provably, at least as hard as the hardest instance. This is the theoretical basis of lattice-based cryptography.
- **2005:** Oded Regev introduces Learning With Errors (LWE), together with a quantum worst-case-to-average-case reduction from lattice problems, the object of this talk.

A Little History First

- **1976–1977:** Diffie–Hellman and RSA launch public-key cryptography, resting on two number-theoretic assumptions: factoring is hard, and discrete logarithms are hard.
- **1994:** Peter Shor shows a quantum computer solves both in polynomial time.
- **1996:** Miklós Ajtai discovers the first worst-case-to-average-case reduction for a lattice problem: a random instance is, provably, at least as hard as the hardest instance. This is the theoretical basis of lattice-based cryptography.
- **2005:** Oded Regev introduces Learning With Errors (LWE), together with a quantum worst-case-to-average-case reduction from lattice problems, the object of this talk.
- **2016–2024:** NIST runs an open, multi-round competition for post-quantum standards; in 2024, ML-KEM (built on a variant of LWE, formerly "Kyber") is standardized for public-key encryption.

A Little History First

- **1976–1977:** Diffie–Hellman and RSA launch public-key cryptography, resting on two number-theoretic assumptions: factoring is hard, and discrete logarithms are hard.
- **1994:** Peter Shor shows a quantum computer solves both in polynomial time.
- **1996:** Miklós Ajtai discovers the first worst-case-to-average-case reduction for a lattice problem: a random instance is, provably, at least as hard as the hardest instance. This is the theoretical basis of lattice-based cryptography.
- **2005:** Oded Regev introduces Learning With Errors (LWE), together with a quantum worst-case-to-average-case reduction from lattice problems, the object of this talk.
- **2016–2024:** NIST runs an open, multi-round competition for post-quantum standards; in 2024, ML-KEM (built on a variant of LWE, formerly "Kyber") is standardized for public-key encryption.

Where This Talk is Headed

We will watch a completely insecure linear-algebra problem become exponentially hard after one tiny perturbation, and see why that hardness survives a quantum adversary.

Public-Key Cryptography

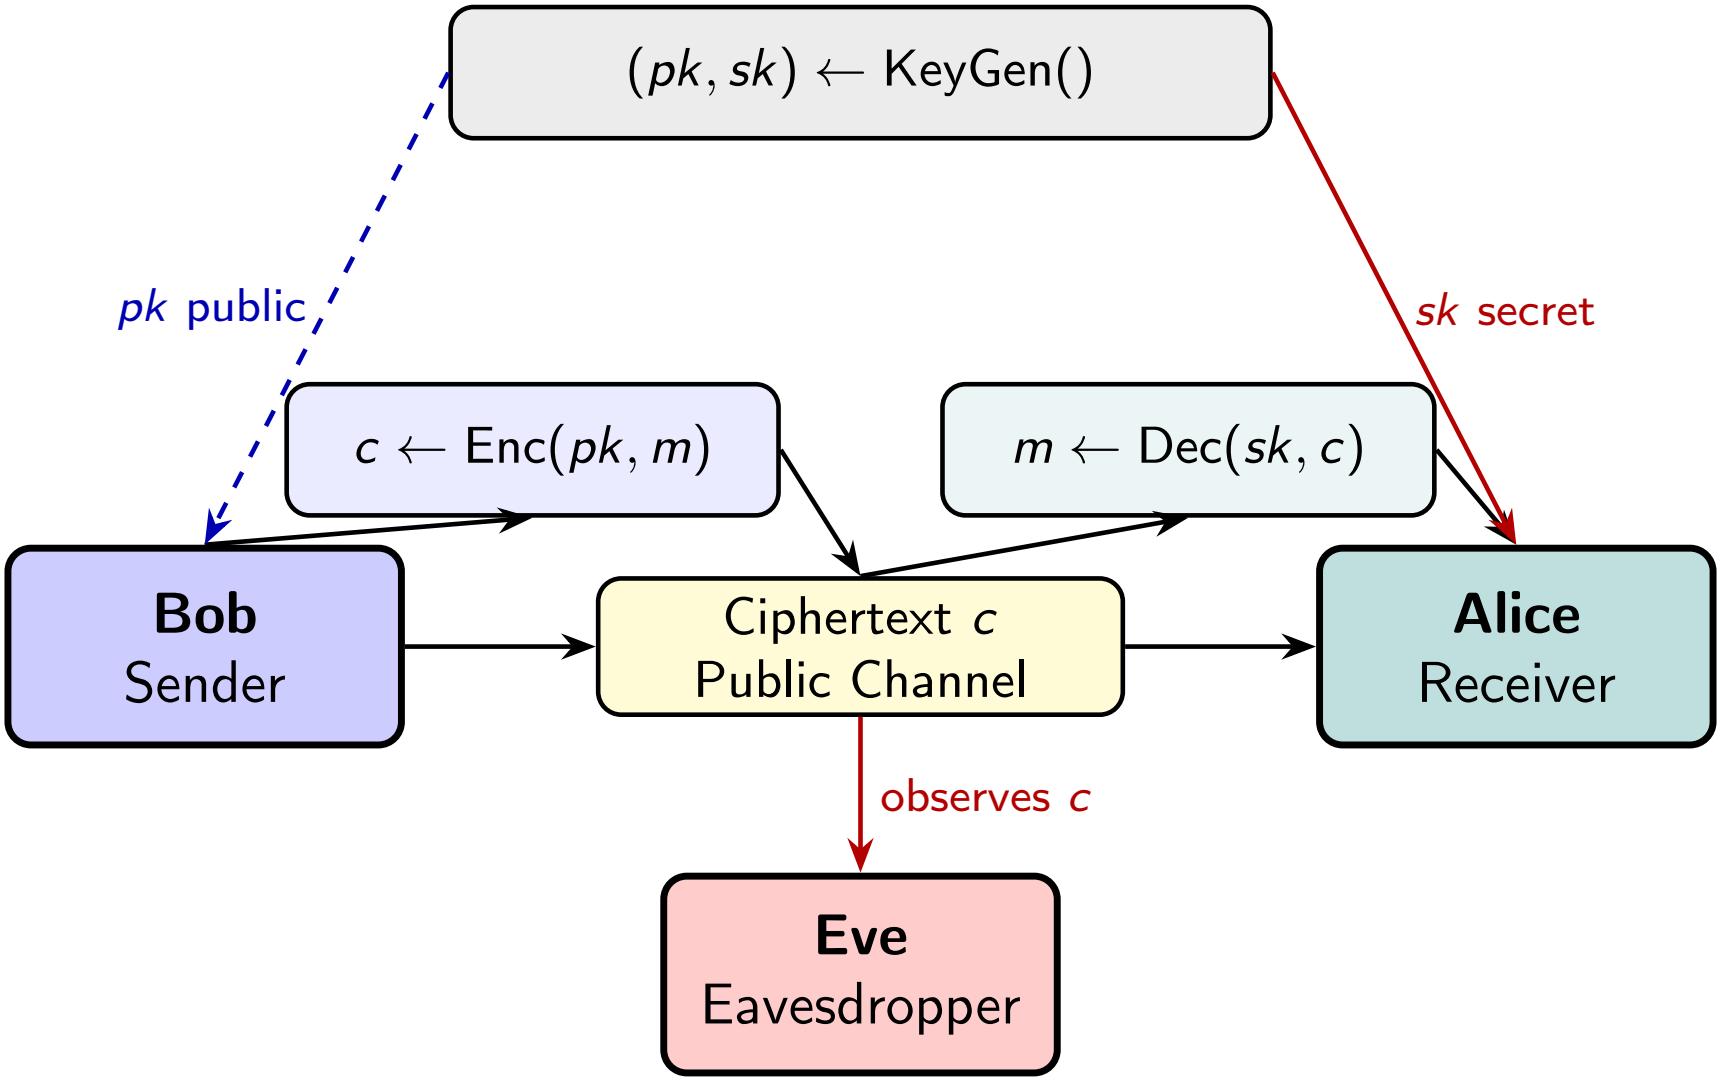

Correctness: $\mathrm{Dec}(sk, \mathrm{Enc}(pk, m)) = m$

Quantum Computation as Linear Algebra

A single qubit is a unit vector in the Hilbert space $\mathbb{C}^2$, written

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1.$$

Quantum Computation as Linear Algebra

A single qubit is a unit vector in the Hilbert space $\mathbb{C}^2$, written

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1.$$

For n qubits, $(\mathbb{C}^2)^{\otimes n} \cong \mathbb{C}^{2^n}$, and a general state is

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle, \qquad \sum_x |\alpha_x|^2 = 1.$$

Quantum Computation as Linear Algebra

A single qubit is a unit vector in the Hilbert space $\mathbb{C}^2$, written

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1.$$

For n qubits, $(\mathbb{C}^2)^{\otimes n} \cong \mathbb{C}^{2^n}$, and a general state is

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle, \qquad \sum_x |\alpha_x|^2 = 1.$$

Quantum computation consists of unitary operators $U \in U(2^n)$ and $UU^\dagger = I$, acting on this exponentially large vector space.

Quantum Computation as Linear Algebra

A single qubit is a unit vector in the Hilbert space $\mathbb{C}^2$, written

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1.$$

For n qubits, $(\mathbb{C}^2)^{\otimes n} \cong \mathbb{C}^{2^n}$, and a general state is

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle, \qquad \sum_x |\alpha_x|^2 = 1.$$

Quantum computation consists of unitary operators $U \in U(2^n)$ and $UU^\dagger = I$, acting on this exponentially large vector space. For example,

$$H|0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \qquad H|1\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}},$$

and therefore

$$H^{\otimes n}|0^n\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle.$$

Quantum Computation as Linear Algebra

A single qubit is a unit vector in the Hilbert space $\mathbb{C}^2$, written

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1.$$

For n qubits, $(\mathbb{C}^2)^{\otimes n} \cong \mathbb{C}^{2^n}$, and a general state is

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle, \qquad \sum_x |\alpha_x|^2 = 1.$$

Quantum computation consists of unitary operators $U \in U(2^n)$ and $UU^\dagger = I$, acting on this exponentially large vector space. For example,

$$H|0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \qquad H|1\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}},$$

and therefore

$$H^{\otimes n}|0^n\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle.$$

Quantum algorithms exploit interference among these amplitudes to solve certain algebraic problems exponentially faster than any known classical algorithm.

The Quantum Fourier Transform (QFT)

- The *Quantum Fourier Transform (QFT)* is the quantum analogue of the discrete Fourier transform. It changes the representation of a quantum state from the computational basis to the frequency basis.
- For an m-dimensional quantum system,

$$|x\rangle \longmapsto \frac{1}{\sqrt{m}} \sum_{k=0}^{m-1} e^{2\pi i x k/m} |k\rangle ,$$

  where the complex exponential introduces phase information that encodes frequency components.
- Unlike a classical Fourier transform, the QFT acts on the amplitudes of a quantum superposition, allowing interference between basis states to amplify some outcomes while suppressing others.
- The QFT can be implemented using only $O((\log m)^2)$ elementary quantum gates, making it exponentially more efficient than the classical discrete Fourier transform acting on a vector of length m.

Hidden Subgroup Problems

Fact

Both integer factorization and the discrete logarithm problem reduce to instances of the **Hidden Subgroup Problem** (HSP) over a finite abelian group G: given a function constant and injective on cosets of an unknown subgroup $H \leq G$, determine H. The abelian HSP admits a quantum polynomial-time algorithm via the QFT over G, for *every* finite abelian G.

- A QFT-based attack is not special to $\mathbb{Z}_N^\times$, but attacks any hardness assumption whose difficulty secretly reduces to locating the hidden period of an abelian group operation.

The Design Constraint This Imposes

A hardness assumption intended to resist quantum adversaries must not derive its difficulty from abelian periodic structure.

Our candidate: noisy linear algebra over a finite field, $As + e = b$ (Regev, 2005).

The Field $\mathbb{F}_q$, and the One Property We Need

- $\mathbb{Z}/q\mathbb{Z}$, the quotient of $\mathbb{Z}$ by the ideal $q\mathbb{Z}$, is always a commutative ring; it is a field iff q is prime. Indeed if $q = de$ nontrivially ($1 < d, e < q$), then $[d][e] = [q] = [0]$ while $[d], [e] \neq [0]$: these are zero divisors, and a field by definition has none, since $[d][e] = 0$ must force $[d] = 0$ or $[e] = 0$.
- Consequently every nonzero element of $\mathbb{F}_q$ is invertible, so Gaussian elimination (which requires dividing by pivots) runs to completion.
- $\mathbb{F}_q^n$ is an n-dimensional $\mathbb{F}_q$-vector space; by the Steinitz exchange lemma, any two bases have equal cardinality, so dimension is an invariant of the space.

The Property That Will Matter

$\mathbb{F}_q$ carries no compatible metric or ordering: as a field element, every nonzero residue is exactly as "far" from 0 as every other. There is no notion of a small field element. Hold onto this! It is why noise will later be invisible to finite-field arithmetic.

Three Statements Concerning $As = b$

Theorem (Elimination Complexity)

For $A \in \mathbb{F}_q^{m \times n}$, the system $As = b$ can be solved, or certified inconsistent, using $O(mn \cdot \min(m, n))$ field operations.

Theorem (Rank–Nullity)

If nonempty, the solution set of $As = b$ is precisely a coset $s_0 + \ker(A)$ of the linear subspace $\ker(A) = \{x : Ax = 0\}$, with $\dim \ker(A) = n - \operatorname{rank}(A)$; hence, the system has $q^{n - \operatorname{rank}(A)}$ solutions.

Proposition (Density of Invertibility)

For $A \in \mathbb{F}_q^{n \times n}$ drawn uniformly,

$$\Pr[A \text{ invertible}] = \prod_{k=1}^{n} \left(1 - q^{-k}\right) \geq 1 - \frac{1}{q-1},$$

obtained by choosing rows one at a time, each new row required only to avoid the span of the previous ones.

The Perturbation

Replace $As = b$ with

$$As + e = b \pmod{q}, \qquad e \equiv \tilde{e} \pmod{q}, \qquad \tilde{e} \in \mathbb{Z}^m,\ \|\tilde{e}\| \ll q.$$

- **The coset argument fails.** Now $b \in As + e$ rather than $As + \ker(A)$; these two cosets coincide only in the event $e \in \ker(A)$, which for A of rank n occurs with probability $q^{\operatorname{rank}(A)-m} \leq q^{-(m-n)}$.
- Row reduction over a field either finds an exact solution or correctly certifies inconsistency; there is no intermediate outcome. Fed $As + e = b$, it reports inconsistent, a true statement that discloses nothing useful to an attacker who already knows the system was generated from a genuine secret s.

Numerical illustration ($q = 97$, $s = (3, 61)$, $a_1 = (12, 40)$): the noiseless value is $\langle a_1, s \rangle = 2476 \equiv 51$; adding $e_1 = 2$ yields the published value 53. No finite sequence of the three elementary row operations (each one exact $\mathbb{F}_{97}$-arithmetic) can recover the split $51 + 2$ from the single residue class [53]: this was never encoded in the field structure to begin with.

Definition

A lattice of rank n is

$$\Lambda = \left\{ \sum_{i=1}^{n} z_i b_i : z_i \in \mathbb{Z} \right\}$$

for linearly independent $b_1, \ldots, b_n \in \mathbb{R}^n$. It is a discrete additive subgroup of $\mathbb{R}^n$: every bounded region contains only finitely many lattice points.

- $\det(\Lambda) := |\det B|$ is independent of the chosen basis (any two bases differ by $U \in GL_n(\mathbb{Z})$, $|\det U| = 1$), and equals the volume of the fundamental parallelepiped.
- $\lambda_1(\Lambda) := \min_{v \in \Lambda \setminus \{0\}} \|v\|$, the length of a shortest nonzero vector.

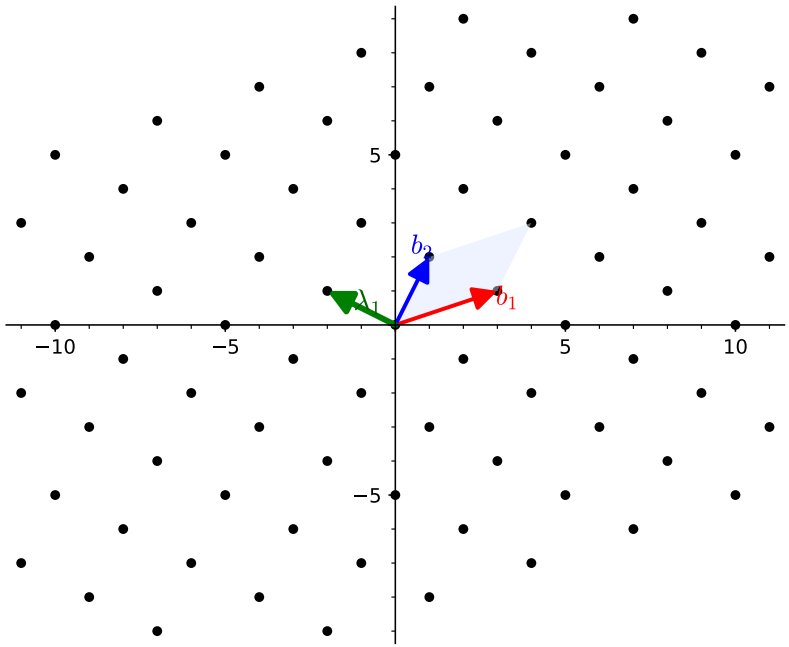

Figure: A two-dimensional lattice generated by basis vectors b_1 and b_2. The shortest nonzero lattice vector has length $\lambda_1(\Lambda)$.

Theorem (Minkowski)

$\lambda_1(\Lambda) \leq \sqrt{n} \det(\Lambda)^{1/n}$: an existence statement, with no efficient algorithm for showing this short vector.

SVP, GapSVP, and SIVP

- **SVP:** given a basis for Λ, locate a vector achieving $\lambda_1(\Lambda)$.
- GapSVP_γ: given (Λ, d) with the promise that either $\lambda_1(\Lambda) \leq d$ (YES) or $\lambda_1(\Lambda) > \gamma d$ (NO), decide which. For $\gamma = 1$ this is exact SVP; for $\gamma = \sqrt{n}$ the problem lies in NP $\cap$ coNP.
- SIVP_γ: output n linearly independent lattice vectors, each of length at most $\gamma\lambda_n(\Lambda)$, where $\lambda_n(\Lambda)$ is the n-th successive minimum (smallest radius containing n independent lattice vectors). A solution to SIVP_γ yields one to $\mathrm{GapSVP}_{\gamma\sqrt{n}}$ through Minkowski's inequality but not conversely.

State of the art

LLL and BKZ (combined with enumeration/sieving) solve these in $2^{\Theta(n)}$ time for polynomial γ. No subexponential quantum algorithm is known for either problem at polynomial approximation factors.

Formal Definitions: Discrete Gaussians and LWE

- For $s > 0$, the discrete Gaussian $D_{\Lambda,s}$ assigns $x \in \Lambda$ probability proportional to $\rho_s(x) = \exp(-\pi\|x\|^2/s^2)$. A sample satisfies $\|x\| \leq s\sqrt{n}$ with probability $1 - 2^{-\Omega(n)}$.
- Fix n, prime $q = q(n)$, noise rate $\alpha \in (0, 1)$, and $\chi = D_{\mathbb{Z},\alpha q}$ mod q.

Search-LWE$_{n,q,\chi}$

Given m samples $(a_i, \langle a_i, s\rangle + e_i)$ with a_i uniform in $\mathbb{F}_q^n$ and $e_i \sim \chi$ i.i.d., recover the fixed secret s.

Decision-LWE$_{n,q,\chi}$

Distinguish, with non-negligible advantage, m such samples from m pairs drawn uniformly from $\mathbb{F}_q^n \times \mathbb{F}_q$.

Stacked as rows this is $b = As + e$, the same equation as before, now reinterpreted as a BDD instance rather than an (insecure) linear system.

Bounded Distance Decoding

Figure: A target vector $t = v + e$ lies within a decoding radius $d < \lambda_1(\Lambda)/2$ of a unique lattice point v, making bounded distance decoding possible.

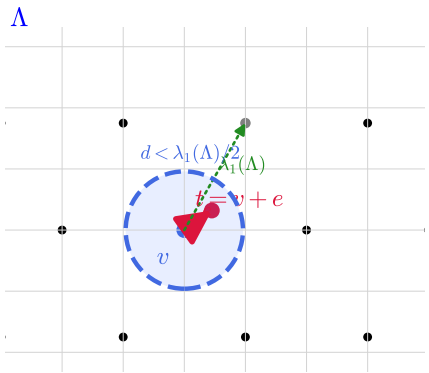

Definition (BDD)

Given Λ and $t = v + e$ with $v \in \Lambda$ and $\|e\| \leq d < \lambda_1(\Lambda)/2$, recover v. The radius bound guarantees uniqueness: any two lattice points within d of t would be within $2d < \lambda_1(\Lambda)$ of each other by the triangle inequality, contradicting the definition of λ_1.

Proposition

Let $\Lambda(A) = \{x \in \mathbb{Z}^m : x \equiv A^T y \pmod{q} \text{ for some } y \in \mathbb{Z}^n\}$, a sublattice of $\mathbb{Z}^m$ containing $q\mathbb{Z}^m$. Then $\text{Search-LWE}_{n,q,\chi}$ is $\text{BDD}_{\Lambda(A),d}$ with target b and $d = O(\alpha q\sqrt{m})$: the noiseless value As, lifted to integers, is a point of $\Lambda(A)$, and the noise e is the BDD displacement.

Regev's Theorem (2005): Worst-Case to Average-Case

Theorem

For noise rate α with $\alpha q \geq 2\sqrt{n}$, there is a probabilistic polynomial-time (quantum) reduction from worst-case lattice problems

$$\mathrm{GapSVP}_\gamma,\ \mathrm{SIVP}_\gamma \quad \text{on an arbitrary } n\text{-dimensional lattice } L$$

to

$$\text{Decision-LWE}_{n,q,\chi},$$

where $\gamma = \widetilde{O}(n/\alpha)$.

The QFT, Inverted: From Attack to Proof

- In Shor's algorithm, the QFT extracts a hidden period from the very function being attacked.
- In Regev's reduction, the QFT acts only on the auxiliary worst-case lattice L^*. It extracts no information from any particular LWE secret s.
- **Why no HSP attack applies here:** once $e \neq 0$, the function $a \mapsto \langle a, s \rangle + e$ is constant on no nontrivial coset of any subgroup of $\mathbb{F}_q^n$. The discrete Gaussian tail of e sees to that. There is, quite literally, no hidden subgroup for the QFT to find on the LWE side.

An Asymmetry Worth Noting

Peikert (2009) later gave a fully *classical* reduction for moduli q of special form, at the cost of a worse approximation factor γ. The fully general statement above, however, still requires the quantum step: LWE's strongest hardness proof is quantum, even though no quantum attack on LWE itself is known.

Our Takeaways?

- Classical public-key cryptography relies on algebraic problems that have become efficiently solvable by quantum algorithms.
- LWE replaces exact algebra with noisy linear systems whose security is based on lattices.
- Regev's worst-case-to-average-case reduction connects the security of random LWE instances to hard lattice problems such as GapSVP and SIVP.
- Thus, lattice-based cryptography is our leading foundation for post-quantum standards!

Thank you!

Questions?