Security considerations for LWE/RLWE

Kristin Lauter

WAM: Uhlenbeck Lectures #4

May 25, 2018

Learning With Errors:

It is hard to solve secret $\mathbf{s}$ from the linear system

$$\begin{cases} \langle \mathbf{a}_0, \mathbf{s} \rangle + e_0 = b_0 \pmod q \\ \langle \mathbf{a}_1, \mathbf{s} \rangle + e_1 = b_1 \pmod q \\ \langle \mathbf{a}_2, \mathbf{s} \rangle + e_2 = b_2 \pmod q \\ \qquad \vdots \qquad\qquad \vdots \\ \langle \mathbf{a}_{d-1}, \mathbf{s} \rangle + e_{d-1} = b_{d-1} \pmod q \end{cases}$$

unless e_j are known.

$q := 2^r$ an integer modulus (r not necessarily an integer)

n an integer, $\mathbf{s} \in \mathbb{Z}_q^n$ a secret vector chosen uniformly at random

$D_{\mathbb{Z},\sigma}$ (error distribution) the discrete Gaussian distribution centered at 0, with standard deviation σ

Definition 1 (LWE sample)

An LWE sample is a pair $(\mathbf{a}, t) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, where $\mathbf{a}$ is sampled uniformly at random from $\mathbb{Z}_q^n$, $e \leftarrow D_{\mathbb{Z},\sigma}$ and $t = \left[\langle \mathbf{a}, \mathbf{s} \rangle + e\right]_q = \langle \mathbf{a}, \mathbf{s} \rangle_q + e \in (-q/2, q/2)$.

Definition 2 (search-LWE$_{n,r,d,\sigma}$)

Given d LWE samples $(\mathbf{a}_i, t_i)$, the problem search-LWE$_{n,r,d,\sigma}$ is to recover the secret vector $\mathbf{s}$.

Let Λ be the $(n+d)$-dimensional lattice generated by the rows of the matrix

$$\begin{pmatrix} q & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ 0 & q & \cdots & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & \cdots & q & 0 & 0 & \cdots & 0 \\ \mathbf{a}_0[0] & \mathbf{a}_1[0] & \cdots & \mathbf{a}_{d-1}[0] & 1/2^{\ell-1} & 0 & \cdots & 0 \\ \mathbf{a}_0[1] & \mathbf{a}_1[1] & \cdots & \mathbf{a}_{d-1}[1] & 0 & 1/2^{\ell-1} & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots \\ \mathbf{a}_0[n-1] & \mathbf{a}_1[n-1] & \cdots & \mathbf{a}_{d-1}[n-1] & 0 & 0 & \cdots & 1/2^{\ell-1} \end{pmatrix}.$$

Easy to see:

$$\mathbf{v} = \left[\langle \mathbf{a}_0, \mathbf{s}\rangle_q, \langle \mathbf{a}_1, \mathbf{s}\rangle_q, \ldots, \langle \mathbf{a}_{d-1}, \mathbf{s}\rangle_q, \mathbf{s}[0]/2^{\ell-1}, \mathbf{s}[1]/2^{\ell-1}, \ldots, \mathbf{s}[n-1]/2^{\ell-1}\right] \in \Lambda$$

$$\mathbf{u} = \left[t_0, t_1, \ldots, t_{d-1}, 0, \ldots, 0\right] \notin \Lambda \text{ but is close to } \mathbf{v} \text{ if } \ell \text{ is big}$$

Lattice Basis Reduction: LLL

LLL polynomial time, exponentially bad approximation factor:

If λ = length of shortest vector, LLL finds a vector of length at most $\gamma \lambda$,

Where $\gamma < 2^{n/2}$

LLL runs in polynomial time: $O(n^5 \log(q)^3)$

To recover s:

1. Use LLL to find a reduced basis for Λ.
2. Use Babai's `NearestPlanes` algorithm to find a lattice point close to $\mathbf{u}$.
3. `NearestPlanes` will recover $\mathbf{w} \in \Lambda$ with

$$||\mathbf{w} - \mathbf{u}|| = 2^{\mu(n+d)}\,\mathrm{dist}(\Lambda, \mathbf{u})$$

   where $\mu \leq 1/4$.
4. But $\mathbf{v}$ is such a lattice point!

Theorem 6 (Laine-Lauter)

Any instance of LWE with $q > 2^{2n}$ *can be broken in polynomial-time using roughly* $2n$ *samples. In practice significantly smaller* q *are vulnerable.*

Examples of recovering the LWE secret: $(\sigma = 8/\sqrt{2\pi})$

n	Samples	$\log_2 q$	Time
80	255	16	10m
100	300	19	24m
120	335	22	61m
140	380	24	1.6h
160	420	27	2.9h
180	460	29	4.4h
200	500	32	7.2h
250	600	39	19h
300	705	45	1.8d
350	805	52	3.7d

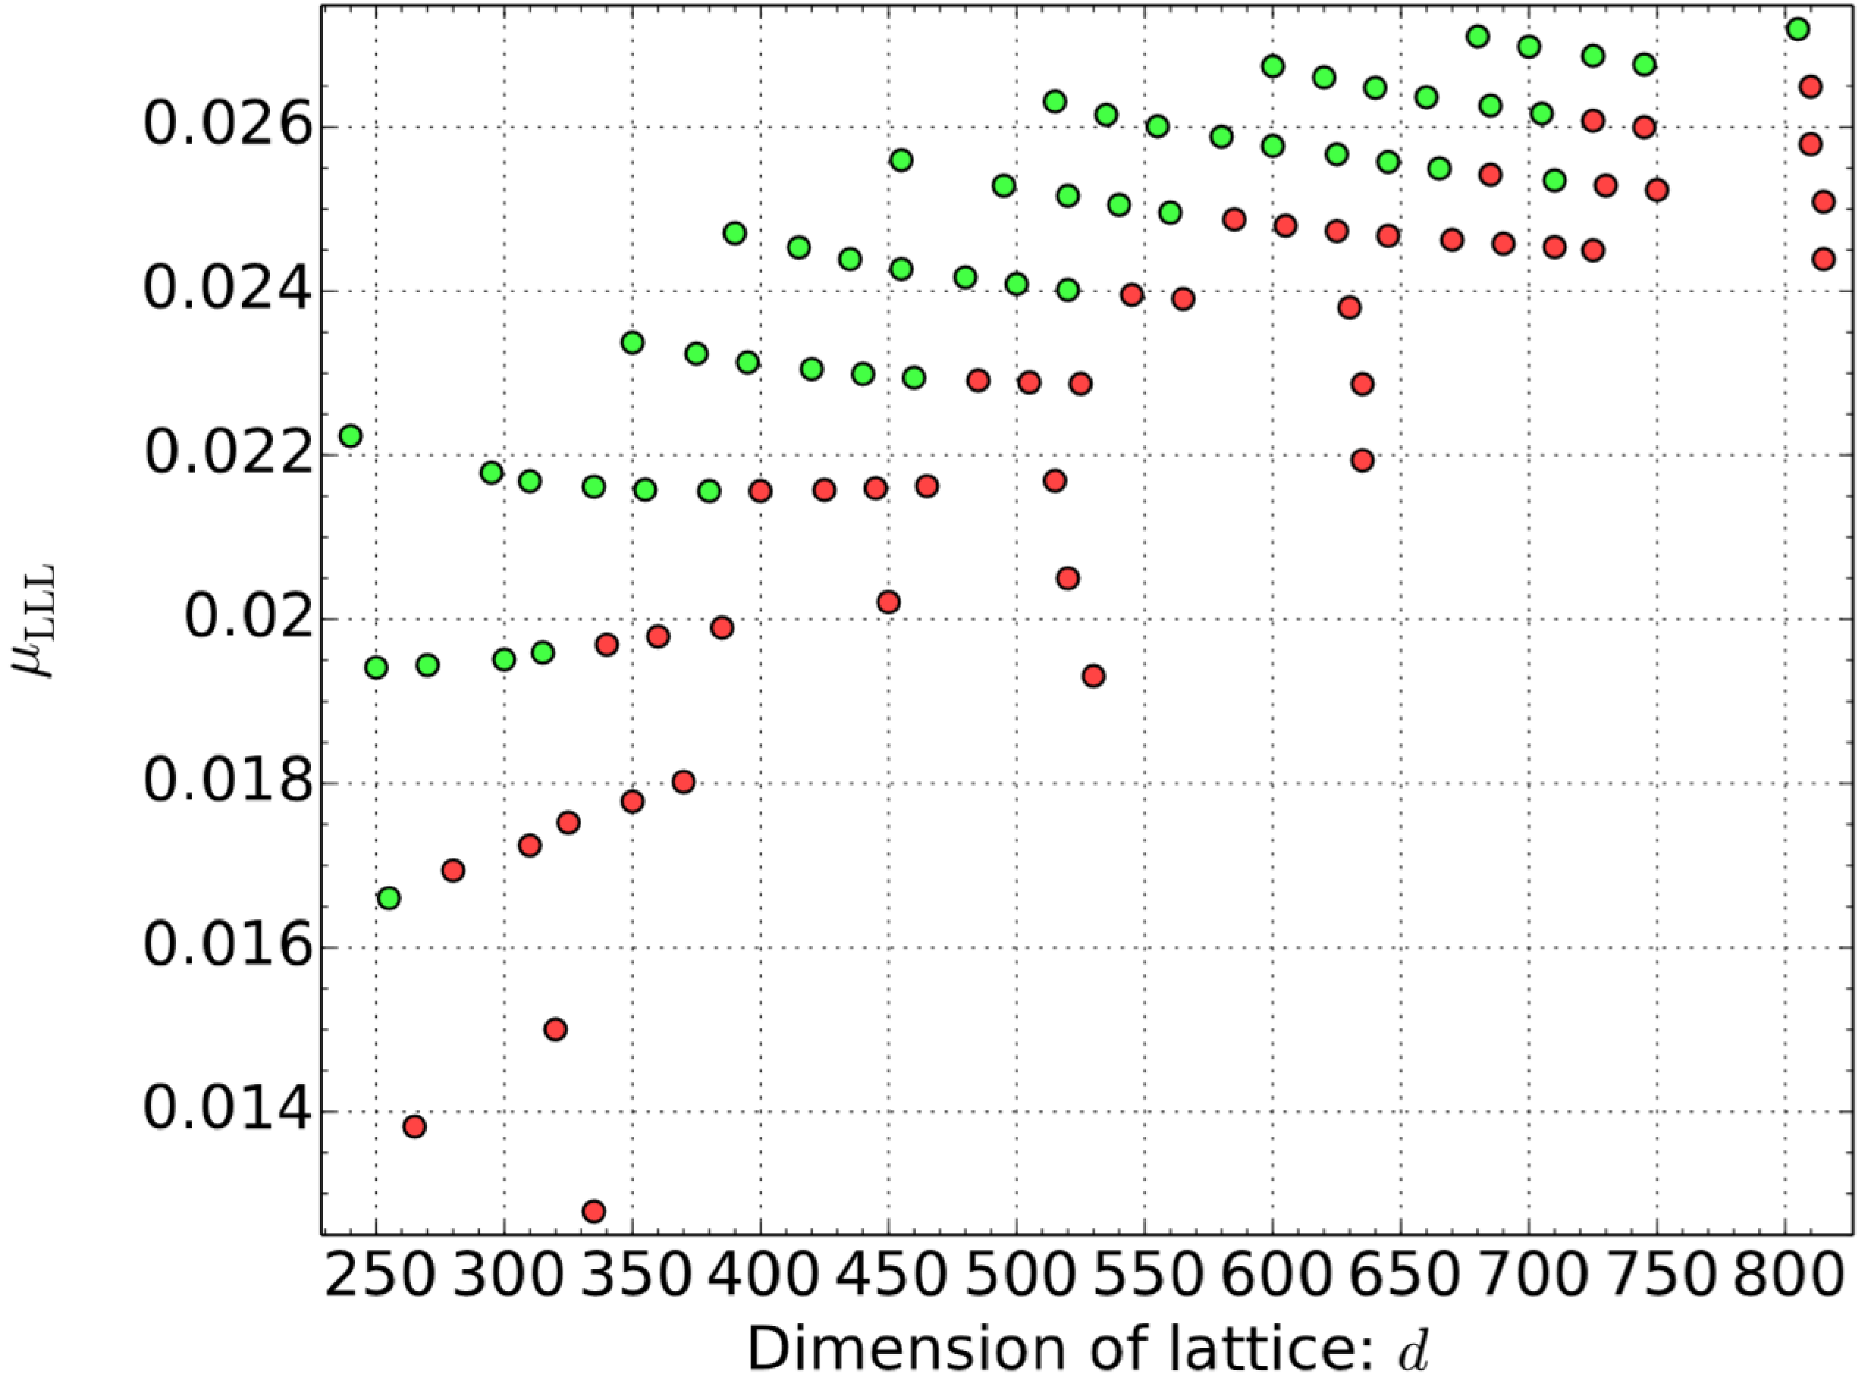

0.026
0.024
0.022
0.02
0.018
0.016
0.014
μ_{LLL}
250 300 350 400 450 500 550 600 650 700 750 800
Dimension of lattice: d

Parameter sizes

Secret picked from Uniform distribution

n	security level	log(q)	uSVP	dec	dual
1024	128	31	130.6	133.8	147.5
	192	22	203.6	211.2	231.8
	256	18	269.9	280.5	303.6
2048	128	59	129.5	129.7	139.2
	192	42	194.0	197.6	212.4
	256	33	263.8	270.7	289.9
4096	128	113	131.9	129.4	136.8
	192	80	192.7	193.2	203.2
	256	63	260.7	263.6	277.6
8192	128	222	132.9	128.9	134.9
	192	157	195.4	192.8	200.6
	256	124	257.0	256.8	266.7
16384	128	440	133.9	129.0	133.0
	192	310	196.4	192.4	198.7
	256	243	259.5	256.6	264.1
32768	128	880	134.3	129.1	131.6
	192	612	198.8	193.9	198.2
	256	480	261.6	257.6	263.6

Algorithm to select parameters ([BLN13])

Given a task:

determine the depth of the circuit required

determine bound on the potential plaintext growth

select plaintext modulus t to exceed this bound

now (n,q) selected to satisfy 2 conditions:

1. *q/t determines the error growth bound. Choose q large enough to allow for correct decryption after the circuit is evaluated (either with or without bootstrapping)*
2. *n must be chosen large enough to achieve 128-bit security with such a q*

Size of (n,q) and the size of the circuit determine the performance.

Ring-Learning With Errors:

It is hard to solve s from the polynomial system

$$\begin{cases} a_0(x)s(x) + e_0(x) = b_0(x) \\ a_1(x)s(x) + e_1(x) = b_1(x) \\ a_2(x)s(x) + e_2(x) = b_2(x) \\ \qquad \vdots \qquad\qquad \vdots \\ a_{d-1}(x)s(x) + e_{d-1}(x) = b_{d-1}(x) \end{cases}$$

unless $e_j(x)$ are known.

- $R = \mathbb{Z}[x]/(f)$, f monic irreducible over $\mathbb{Z}$
- $R_q = \mathbb{F}_q[x]/(f)$, q prime
- χ an error distribution on R_q
- Given a series of samples $(a, as + e) \in R_q^2$ where
  1. $a \in R$ uniformly,
  2. $e \in R$ according to χ,

  find s.

Decision Ring-LWE:

- Given samples (a, b), determine if they are LWE-samples or uniform $(a, b) \in R_q^2$.

Eisentraeger-Hallgren-Lauter attack:

Potential weakness: $f(1) \equiv 0$ **mod** q**.**

1. Ring homomorphism $R_q \to \mathbb{F}_q$ by evaluation at 1
2. Samples transported to $\mathbb{F}_q$:

$$(a(1), a(1)s(1) - e(1))$$

3. The error $e(1)$ is small if $e(x)$ has small coefficients.
4. Search for $s(1)$ exhaustively (try each, see if purported $e(1)$ is small).

Polynomial embedding: Think of R as a lattice via

$$R \hookrightarrow \mathbb{Z}^n \hookrightarrow \mathbb{R}^n, \quad a_n x^n + \ldots + a_0 \mapsto (a_n, \ldots, a_0).$$

Note: multiplication is ‘mixing’ on coefficients.
Actually work modulo q:

$$R_q \hookrightarrow \mathbb{F}_q^n, \quad a_n x^n + \ldots + a_0 \mapsto (a_n \bmod q, \ldots, a_0 \bmod q).$$

Naive sampling: Sample each coordinate as a one-dimensional discretized Gaussian. This leads to a discrete approximation to an n-dimensional Gaussian.

Minkowski embedding: A number field K of degree n can be embedded into $\mathbb{C}^n$ so that **multiplication and addition are componentwise**:

$$K \mapsto \mathbb{C}^n, \quad \alpha \mapsto (\alpha_1, \alpha_2, \ldots, \alpha_n)$$

where α_i are the n Galois conjugates of α. Massage into $\mathbb{R}^n$:

$$\phi : R \hookrightarrow \mathbb{R}^n, \quad (\underbrace{\alpha_1, \ldots, \alpha_r,}_{\text{real}} \underbrace{\Re(\alpha_{r+1}), \Im(\alpha_{r+1}), \ldots}_{\text{complex}}).$$

As usual, then we work modulo q (modulo prime above q).
Sampling: Discretize a Gaussian, spherical in $\mathbb{R}^n$ under the usual inner product.

WIN3 project: Elias-Lauter-Ozman-Stange attack [ELOS, Crypto15]

Suppose: CRT decomposition (f splits mod q):

$$R_q \cong \mathbb{F}_q^n$$

with n ring homomorphisms $\phi_i : R_q \to \mathbb{F}_q$,
Question: Given a distribution χ on R_q, when is the image distribution $\phi_i(\chi)$ distinguishable from uniform in $\mathbb{F}_q$?

- EHL: if ϕ_i takes $x \mapsto 1$, then it is distinguishable.
- Other cases with some hope for success on Poly-LWE:
  - $\phi_i(x)$ of small order (suggested by Eisenträger-Hallgren-Lauter)
  - $\phi_i(x)$ near 0.

- $\sigma =$ parameter for the Gaussian in Minkowski embedding
- $M =$ change of basis matrix from Minkowski embedding of R to its polynomial basis.

Theorem (Elias-Lauter-Ozman-Stange)

Let K be a number field with:

1. *ring of integers $\mathbb{Z}[\beta]$*
2. *q prime such that min poly of β has root 1 modulo q*
3. *spectral norm $\rho(M)$ satisfies*

$$\rho < \frac{q}{4\sqrt{2\pi}\sigma n}$$

Then Ring-LWE decision can be solved in time $\widetilde{O}(\ell q)$ with probability $1 - 2^{-\ell}$ using ℓ samples.

Search RLWE attacks: Chen-Lauter-Stange '15

Theorem (Elias-Lauter-Ozman-Stange)

Let $f = x^n + q - 1$ *be such that*

1. q *prime,* $q - 1$ *squarefree*
2. n *is a power of a prime* p
3. $\mathbf{p}^2 \nmid ((1-q)^n - (1-q))$
4. $\tau > 1$ *where*

$$\tau := \frac{q \det(M)^{1/n}}{4\sqrt{\pi}\sigma n(q-1)^{1/2-1/2n}}$$

Then Ring-LWE decision can be solved in time $\widetilde{O}(\ell q)$ *with probability* $1 - 2^{-\ell}$ *using* ℓ *samples.*

New questions in number theory

Are these problems hard for other number rings??

In general, NO: not for small error.

Eisentraeger-Hallgren-L (2014), Elias-L-Ozman-Stange (2015), Chen-L-Stange (2015)

Number Theory Questions:

distributions of elements of small order in finite fields,

relationship with Mahler measure,

construction of number rings with certain properties.

Course Goals:

- Introduce Post-Quantum Cryptography, overview of candidates
- familiarity with running time of algorithms and best attacks
- Introduce Supersingular Isogeny Graphs (SIG and SIKE)
- Introduce Lattice-based cryptography and applications

Thank you! To the participants, to IAS, to NSF, and to the Organizers!

Thanks to my coauthors Kim Laine and Kate Stange for slides in today's talk.

Joint work with:

LWE Attacks: Kim Laine

RLWE Attacks: Kirsten Eisentraeger, Sean Hallgren, Kate Stange, Ekin Ozman, Yara Elias, Hao Chen, SAC '14, Crypto '15, SAC'16, SIAGA