Attacks on RLWE

Yara Elias, Kristin Lauter, Ekin Ozman, Katherine Stange
to appear at CRYPTO 2015

Simons Institute
July 7, 2015

Homomorphic Encryption

Practical Homomorphic Encryption schemes based on lattices, proposed in 2011 by Brakerski, Gentry, Vaikuntanathan in [BV], [BGV]. (also [GHS], [GSW], [SS], [BLLN], ...)

Applications to cloud storage and services:

*private cloud-based electronic medical records systems

*private predictive analysis

*machine learning on encrypted data

*genomic computation on encrypted data

Some History: Lattice-based Crypto

* Ajtai-Dwork public-key cryptosystem: based on the worst-case hardness of a variant of Shortest Vector Problem (SVP) [AD97]

* NTRU family of cryptosystems: defined in particularly efficient lattices connected to number fields [HPS98]

* NTRU standardized in IEEE P1363.1 Lattice-Based Public Key Cryptography standard [2008]

New Hardness Assumptions

*New assumption introduced, Learning-With-Errors (LWE) [Regev]

*Ring-Learning-With-Errors (RLWE) proposed [Lyubashevsky-Peikert-Regev]

* LWE/RLWE related via security reductions to hard lattice problems: (Gap-)SVP and Bounded Distance Decoding (BDD) [Regev, Lyubashevsky-Peikert-Regev, ...]

Ring-LWE distribution

$K =$ number field, $R = \mathcal{O}_K$,
$R^\vee = \{y \in R | Tr(xy) \in \mathbb{Z}$ for all $x \in R\}$.

$K_\mathbb{R} = K \otimes \mathbb{R}$ and $\mathbb{T} = K_\mathbb{R}/R^\vee$.

For $q \in \mathbb{Z}$, let $R_q := R/qR$.

Definition (Ring-LWE Distribution)

For $s \in R_q^\vee$ a secret, and an error distribution ψ over $K_\mathbb{R}$, the Ring-LWE distribution $A_{s,\psi}$ over $R_q \times \mathbb{T}$ consists of samples

$$(a, (a \cdot s)/q + e \bmod R^\vee)$$

$a \in R_q$ chosen uniformly at random, e chosen from the error distribution ψ.

Ring-LWE hardness assumptions

Definition (Ring-LWE Search Problem)

Let Ψ be a family of distributions over $K_{\mathbb{R}}$. The Ring-LWE Search problem ($RLWE_{q,\Psi}$), for some $s \in R_q^\vee$ and $\psi \in \Psi$, is to *find* s, given arbitrarily many independent samples from $A_{s,\psi}$.

Definition (Ring-LWE Average-Case Decision Problem)

Let Υ be a family of error distributions over $K_{\mathbb{R}}$. The Ring-LWE Average-Case Decision problem ($RDLWE_{q,\Upsilon}$) is to *distinguish* with non-negligible advantage between arbitrarily many independent samples from $A_{s,\psi}$, for a random choice of $s \in R_q^\vee$ and $\psi \in \Upsilon$, and the same number of samples chosen independently and uniformly at random from $R_q \times \mathbb{T}$.

Worst-case hardness of search version of ring-LWE

$K =$ cyclotomic number field of degree n, $R = \mathcal{O}_K$, $q =$ prime

$\Psi_\alpha =$ elliptical Gaussian of parameter α.

Theorem (Lyubashevsky, Peikert, Regev in LPR10)

Let $\alpha \in (0,1)$ be such that $\alpha \cdot q \geq \omega(\sqrt{\log n})$, then there is a probabilistic polynomial-time quantum reduction from the $\widetilde{O}(\sqrt{n}/\alpha)$-approximate SIVP problem on ideal lattices in K to $RLWE_{q,\Psi_{\zeta q}}$ given l samples where $\zeta = \alpha(ln\log(ln))^{\frac{1}{4}}$.

Search-to-Decision reductions: [LPR] for cyclotomics, [EHL] for Galois fields.

The PLWE problem

The PLWE problem was first defined in [LPR10] by Lyubashevsky, Peikert, Regev and in [BV11] by Brakerski and Vaikuntanathan.

For all $\kappa \in \mathbb{N}$, let $f(x) = f_\kappa(x)$ be a polynomial of degree $n = n(\kappa)$, and let $q = q(\kappa)$ be a prime integer. Let $R = \mathbb{Z}[x]/(f)$, let $R_q = R/qR$ and let χ denote a distribution over R.

Definition (The PLWE assumption)

The PLWE assumption $\mathrm{PLWE}_{f,q,\chi}$ states that for any $\ell = poly(\kappa)$ it holds that

$$\{a_i, a_i \cdot s + e_i\}_{i \in [\ell]}$$

is computationally *indistinguishable* from $\{a_i, u_i\}_{i \in [\ell]}$, where s is sampled from the noise distribution χ over R_q, the a_i are uniform in R_q, the e_i are sampled from χ and the ring elements u_i are uniformly random over R_q.

Attack on PLWE for some number fields ([EHL])

Let $K = \mathbb{Q}[x]/(f(x))$ be a number field such that $f(1) \equiv 0 \pmod{q}$, and such that q can be chosen large enough.

Let $R := \mathcal{O}_K$, and let $R_q := R/qR$.

Given samples, $(a_i, b_i) \in R_q \times R_q$, we have to decide whether the samples are uniform or come from a PLWE distribution.

To do this we take the representatives of a_i and b_i in R, call them a_i and b_i again, and evaluate them at 1.

The attack

This gives us elements $a_i(1), b_i(1) \in \mathbb{F}_q$.

If (a_i, b_i) are PLWE samples, then by definition,

$$b_i = a_i \cdot s + e_i,$$

and so

$$b_i(1) \equiv (a_i \cdot s)(1) + e_i(1) \pmod{q}.$$

Since $f(1) \equiv 0 \pmod{q}$, the Chinese Remainder Theorem gives us that

$$b_i(1) \equiv a_i(1) \cdot s(1) + e_i(1) \pmod{q}.$$

The attack

Now we can guess $s(1)$, and we have q choices.

For each of our guesses we compute $b_i(1) - a_i(1) \cdot s(1)$.

** If (a_i, b_i) are PLWE samples and our guess for $s(1)$ is correct, then $b_i(1) - a_i(1) \cdot s(1) = e_i(1)$, and we will detect that it is non-uniform, because e_i is taken from χ.

(For example, if e_i is taken from a Gaussian with small radius, then $e_i(1)$ will be "small" for all i and hence not uniform.)

**If (a_i, b_i) are uniform samples, then $b_i(1) - a_i(1) \cdot s(1)$ for any fixed choice of $s(1)$ will still be uniform, since $a_i(1)$, $b_i(1)$ are both uniform modulo q.

Overview of Eisentraeger-Hallgren-Lauter

$K = \mathbb{Q}(\beta) = \mathbb{Q}[x]/(f(x))$, $n =$ degree of K, $R = \mathcal{O}_K$, q prime

Consider the following properties:

1. (q) splits completely in K, and $q \nmid [R : \mathbb{Z}[\beta]]$;
2. K is Galois over $\mathbb{Q}$;
3. the ring of integers of K is generated over $\mathbb{Z}$ by β, $\mathcal{O}_K = \mathbb{Z}[\beta] = \mathbb{Z}[x]/(f(x))$ with $f'(\beta) \mod q$ "small" ;
4. the transformation between the Minkowski embedding of K and the power basis representation of K is given by a scaled orthogonal matrix;
5. $f(1) \equiv 0 \pmod{q}$;
6. q can be chosen suitably large.

Results: [Eisentraeger-Hallgren-Lauter 2014]

*For (K, q) satisfying conditions (1) and (2), we have a search-to-decision reduction from $RLWE_q$ to $RDLWE_q$.

*For (K, q) satisfying conditions (3) and (4), we have a reduction from $RDLWE_q$ to $PLWE_q$.

* For (K, q) satisfying conditions (5) and (6), we have an attack which breaks instances of the PLWE decision problem.

Consequence

For number fields K satisfying all 6 properties, we would have an attack on the RLWE problem!

However, this does not happen in general and we don't have any examples of number fields satisfying ***all 6 properties***.

For example, 2-power cyclotomic fields, which are used in practice, don't satisfy property (5).

Security Reductions

*The proof of the search-to-decision reduction from $RLWE_q$ to $RDLWE_q$ is a slight generalization of the proof given by Lyubashevsky, Peikert, Regev in [LPR] for the case of cyclotomic fields.

*The proof of the reduction from $RDLWE_q$ to $PLWE$ is a slightly more general restatement of the proof given by Ducas and Durmus in [DD] for the 2-power cyclotomic case.

We will not give details in this talk.

Extension of the attack on PLWE

... to a more general class of number fields:

Suppose that $f(x)$ has a root β modulo q which has small order in $(\mathbb{Z}/q\mathbb{Z})^*$.

If $f(\beta) \equiv 0 \mod q$, then the same attack above will work by evaluating samples at β, instead of at 1.

Now unfortunately, the value of the error polynomials $e_i(\beta)$ are harder to distinguish from random ones than in the case $\beta = 1$: although the $e_i(x)$ have small coefficients modulo q, the powers of β may grow large and also may wrap around modulo q.

Extension of the attack on PLWE...

However, if β has small order in $(\mathbb{Z}/q\mathbb{Z})^*$, then the set $\{\beta^i\}_{i=0,\ldots,n-1}$ takes on only a small number values, and this can be used to distinguish samples arising from $e_i(\beta)$ from random ones with non-negligible advantage.

Moving the attack to RLWE

Key point: hardness of RLWE is established when embedding R into $\mathbb{R}^n$ via the *canonical, i.e. Minkowski embedding*

PLWE uses a polynomial basis for the ring R.

Errors are generated coordinate-wise in the polynomial basis

In order to attack an RLWE instance, the error must not get too distorted when passing to the polynomial basis.

This distortion we will call the *spectral distortion* for $R = \mathbb{Z}[\beta]$.

Weak RLWE

A Ring-LWE instance is weak if the following three properties hold:

1. K is monogenic.
2. f satisfies $f(1) \equiv 0 \pmod{q}$.
3. ρ and σ are sufficiently small

where σ is the width of the error distribution and ρ is the spectral distortion.

Main Theorem

Theorem

Let K be a number field such that $K = \mathbb{Q}(\beta)$, $\mathcal{O}_K = \mathbb{Z}[\beta]$. Let f be the minimal polynomial of β, q a prime such that $f(1) \equiv 0 \pmod{q}$ and suppose that the spectral norm ρ satisfies

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

Then the non-dual Ring-LWE decision problem for K, q, σ can be solved in time $\widetilde{O}(\ell q)$ with probability $1 - 2^{-\ell}$, using a dataset of ℓ samples.

Weak Family

Consider the family of polynomials

$$f_{n,q}(x) = x^n + q - 1$$

for q a prime. These satisfy $f(1) \equiv 0 \pmod{q}$.

By the Eisenstein criterion, they are irreducible whenever $q - 1$ has a prime factor that appears to exponent 1.

Weak family

Theorem

Suppose q is prime, n is an integer and $f = f_{n,q}$ satisfies

1. *n is a power of the prime ℓ,*
2. *$q - 1$ is squarefree,*
3. *$\ell^2 \nmid ((1-q)^n - (1-q))$,*
4. *we have $\tau > 1$, where*

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

Then the non-dual Ring-LWE decision problem can be solved in time $\widetilde{O}(\ell q)$ with probability $1 - 2^{-\ell}$, using a dataset of ℓ samples.

Heuristics, examples, code

CRYPTO 2015 paper contains:

Examples of weak PLWE fields and weak RLWE fields

Weak PLWE cyclotomic fields with alternate polynomial basis

Code for attacks

Heuristics on spectral norms for general number fields

Questions in Number Theory

Parameter choices

Suggested parameter choices secure against the *distinguishing attack* by Micciancio and Regev [MR09] and the *decoding attack* by Lindner and Peikert [LP]

Concrete security estimates [LP] against these attacks lead to suggested parameters, at the "high security" level, of $n = 320$, $q \approx 2^{12}$, and $\sigma = 8$.

For those parameter choices, the distinguishing attack is estimated to run in time 2^{122} (seconds) to obtain a distinguishing advantage of 2^{-64}.

Parameter choices

The *decoding attack* in Lindner and Peikert [LP] recovers the secret. It requires a reduced basis, and the estimated time to compute the reduced basis when $n = 320$ and $q \approx 2^{12}$ is 2^{119} seconds for decoding probability 2^{-64}.

Our attack on PLWE for weak number fields runs in time $\widetilde{O}(q)$, so these parameters would not be safe against this attack.

Typically, Leveled and Practical Homomorphic Encryption schemes use much larger q, at least 2^{128}, and those parameters would be fine. ([GHS], [LNV], [GLN], [BLN])

Successfully coded attacks

Ring-LWE and Poly-LWE parameters attacked on a Thinkpad X220 laptop with Sage Mathematics Software

case	f	q	s	τ	samples per run	time per run
Poly-LWE	$x^{1024} + 2^{31} - 2$	$2^{31} - 1$	3.192	N/A	40	13.5 hrs
Ring-LWE	$x^{128} + 524288x + 524285$	524287	8.00	N/A	20	24 sec
Ring-LWE	$x^{192} + 4092$	4093	8.87	0.0136	20	25 sec
Ring-LWE	$x^{256} + 8189$	8190	8.35	0.0152	20	44 sec

Questions in Number theory

What are possible spectral distortions of algebraic numbers?

Are there fields of cryptographic size which are Galois and monogenic? (other than the cyclotomic number fields and their maximal real subfields?)

What is the distribution of elements of small order among residues modulo q?

What is the smallest residue modulo a prime q which has order exactly r ?