How (Not) to Instantiate Ring-LWE

Chris Peikert
University of Michigan

Security and Cryptography for Networks
1 September 2016

Conclusions

Conclusions

1. Prior insecure Ring-LWE instantiations turn out to use quite narrow error distributions that are incongruous to the ring geometry. This explains their vulnerability to attacks.

Conclusions

1. Prior insecure Ring-LWE instantiations turn out to use quite narrow error distributions that are incongruous to the ring geometry. This explains their vulnerability to attacks.

2. 'Peculiar' aspects of the Ring-LWE definition and worst-case hardness theorems—adopted for generality and tightness—also yield provable immunity to the attacks (and generalizations).

Conclusions

1. Prior insecure Ring-LWE instantiations turn out to use quite narrow error distributions that are incongruous to the ring geometry.
   This explains their vulnerability to attacks.

2. ‘Peculiar’ aspects of the Ring-LWE definition and worst-case hardness theorems—adopted for generality and tightness—also yield provable immunity to the attacks (and generalizations).

3. For Ring-LWE security, proper choice of error distribution is essential: error should be ‘well spread’ relative to the ring and its small-norm ideals.

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$
\begin{array}{lcl}
\mathbf{a}_1 \leftarrow \mathbb{Z}_q^n & , & b_1 \approx \langle \mathbf{a}_1 , \mathbf{s} \rangle \bmod q \\
\mathbf{a}_2 \leftarrow \mathbb{Z}_q^n & , & b_2 \approx \langle \mathbf{a}_2 , \mathbf{s} \rangle \bmod q \\
 & \vdots &
\end{array}
$$

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$\begin{array}{lcl} \mathbf{a}_1 \leftarrow \mathbb{Z}_q^n & , & b_1 = \langle \mathbf{a}_1 \,,\, \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q \\ \mathbf{a}_2 \leftarrow \mathbb{Z}_q^n & , & b_2 = \langle \mathbf{a}_2 \,,\, \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q \\ & \vdots & \end{array}$$

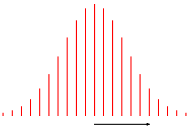

$\sqrt{n} \leq \text{error} \ll q$

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$
\begin{array}{lcl}
\mathbf{a}_1 \leftarrow \mathbb{Z}_q^n & , & b_1 = \langle \mathbf{a}_1 \, , \, \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q \\
\mathbf{a}_2 \leftarrow \mathbb{Z}_q^n & , & b_2 = \langle \mathbf{a}_2 \, , \, \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q \\
 & \vdots &
\end{array}
$$

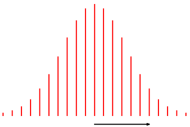

$\sqrt{n} \leq \text{error} \ll q$

- **Decision:** distinguish $(a_i \, , \, b_i)$ from uniform $(a_i \, , \, b_i)$

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$\mathbf{a}_1 \leftarrow \mathbb{Z}_q^n \quad , \quad b_1 = \langle \mathbf{a}_1 \, , \, \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q$$

$$\mathbf{a}_2 \leftarrow \mathbb{Z}_q^n \quad , \quad b_2 = \langle \mathbf{a}_2 \, , \, \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q$$

$$\vdots$$

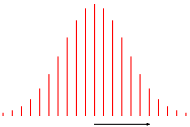

$\sqrt{n} \leq \text{error} \ll q$

- **Decision:** distinguish $(a_i \, , \, b_i)$ from uniform $(a_i \, , \, b_i)$

LWE is Versatile and Hard (...maybe even for quantum!)

worst case lattice problems $\leq$ search-LWE $\leq$ decision-LWE $\leq$ much crypto

(quantum [R'05]) [BFKL'93,R'05,...]

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$\begin{aligned} \mathbf{a}_1 \leftarrow \mathbb{Z}_q^n \quad &, \quad b_1 = \langle \mathbf{a}_1 \,,\, \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q \\ \mathbf{a}_2 \leftarrow \mathbb{Z}_q^n \quad &, \quad b_2 = \langle \mathbf{a}_2 \,,\, \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q \\ &\vdots \end{aligned}$$

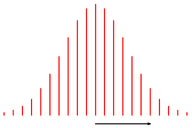

$\sqrt{n} \le \text{error} \ll q$

- **Decision:** distinguish $(a_i \,,\, b_i)$ from uniform $(a_i \,,\, b_i)$

LWE is Versatile and Hard (... maybe even for quantum!)

worst case lattice problems $\le$ search-LWE $\le$ decision-LWE $\le$ much crypto

(quantum [R'05]) [BFKL'93,R'05,...]

- Also a *classical* reduction for search-LWE [P'09,BLPRS'13]

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$\begin{array}{lcl} \mathbf{a}_1 \leftarrow \mathbb{Z}_q^n & , & b_1 = \langle \mathbf{a}_1 , \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q \\ \mathbf{a}_2 \leftarrow \mathbb{Z}_q^n & , & b_2 = \langle \mathbf{a}_2 , \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q \\ & \vdots & \end{array}$$

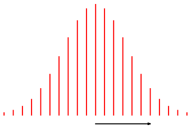

$\sqrt{n} \leq \text{error} \ll q$

- **Decision:** distinguish (a_i , b_i) from uniform (a_i , b_i)

LWE is (sort of) Efficient

- Getting one pseudorandom $\mathbb{Z}_q$-scalar requires an n-dim inner product.

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \text{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$
\begin{aligned}
\mathbf{a}_1 \leftarrow \mathbb{Z}_q^n \quad &, \quad b_1 = \langle \mathbf{a}_1 , \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q \\
\mathbf{a}_2 \leftarrow \mathbb{Z}_q^n \quad &, \quad b_2 = \langle \mathbf{a}_2 , \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q \\
&\vdots
\end{aligned}
$$

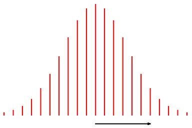

$\sqrt{n} \leq \text{error} \ll q$

- **Decision:** distinguish (a_i , b_i) from uniform (a_i , b_i)

LWE is (sort of) Efficient

- Getting one pseudorandom $\mathbb{Z}_q$-scalar requires an n-dim inner product.
- Cryptosystems have large keys: $\Omega(n^2 \log^2 q)$ bits.

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus $q = \mathrm{poly}(n)$ (usually)
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

$$\mathbf{a}_1 \leftarrow \mathbb{Z}_q^n \quad , \quad b_1 = \langle \mathbf{a}_1 \, , \, \mathbf{s} \rangle + e_1 \in \mathbb{Z}_q$$

$$\mathbf{a}_2 \leftarrow \mathbb{Z}_q^n \quad , \quad b_2 = \langle \mathbf{a}_2 \, , \, \mathbf{s} \rangle + e_2 \in \mathbb{Z}_q$$

$$\vdots$$

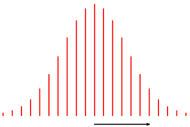

$\sqrt{n} \leq \text{error} \ll q$

- **Decision:** distinguish $(a_i \, , \, b_i)$ from uniform $(a_i \, , \, b_i)$

LWE is (sort of) Efficient

- Getting one pseudorandom $\mathbb{Z}_q$-scalar requires an n-dim inner product.
- Cryptosystems have large keys: $\Omega(n^2 \log^2 q)$ bits.
- Inspired by NTRU [HPS'96], for efficiency we go to the ring setting...

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over R (usually Gaussian in 'canonical' geometry)

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over R (usually Gaussian in 'canonical' geometry)
- Modulus $q \geq 2$ defining $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over R (usually Gaussian in 'canonical' geometry)
- Modulus $q \geq 2$ defining $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$

Search: find secret ring element $s \in R_q$, given independent samples

$$\begin{array}{lcl} a_1 \leftarrow R_q & , & b_1 = a_1 \cdot s + e_1 \in R_q \\ a_2 \leftarrow R_q & , & b_2 = a_2 \cdot s + e_2 \in R_q \qquad\qquad (e_i \leftarrow \chi) \\ & \vdots & \end{array}$$

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over R (usually Gaussian in 'canonical' geometry)
- Modulus $q \geq 2$ defining $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$

Search: find secret ring element $s \in R_q$, given independent samples

$$\begin{array}{lcl} a_1 \leftarrow R_q & , & b_1 = a_1 \cdot s + e_1 \in R_q \\ a_2 \leftarrow R_q & , & b_2 = a_2 \cdot s + e_2 \in R_q \qquad\qquad (e_i \leftarrow \chi) \\ & \vdots & \end{array}$$

Decision: distinguish $(a_i \, , \, b_i)$ from uniform $(a_i \, , \, b_i) \in R_q \times R_q$

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over $R^\vee$ (usually Gaussian in 'canonical' geometry)
- Modulus $q \geq 2$ defining $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$

Search: find secret ring element $s \in R_q^\vee$, given independent samples

$$\begin{aligned} a_1 \leftarrow R_q \quad &, \quad b_1 = a_1 \cdot s + e_1 \in R_q^\vee \\ a_2 \leftarrow R_q \quad &, \quad b_2 = a_2 \cdot s + e_2 \in R_q^\vee \qquad (e_i \leftarrow \chi) \\ &\vdots \end{aligned}$$

Decision: distinguish $(a_i \, , \, b_i)$ from uniform $(a_i \, , \, b_i) \in R_q \times R_q^\vee$

!!! [LPR'10] actually defines R-LWE using 'dual' ideal $R^\vee = t^{-1}R$.

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over $R^\vee$ (usually Gaussian in 'canonical' geometry)
- Modulus $q \geq 2$ defining $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$

Search: find secret ring element $s \in R_q^\vee$, given independent samples

$$\begin{array}{lcl} a_1 \leftarrow R_q & , & b_1 = a_1 \cdot s + e_1 \in R_q^\vee \\ a_2 \leftarrow R_q & , & b_2 = a_2 \cdot s + e_2 \in R_q^\vee \qquad (e_i \leftarrow \chi) \\ & \vdots & \end{array}$$

Decision: distinguish $(a_i\ ,\ b_i)$ from uniform $(a_i\ ,\ b_i) \in R_q \times R_q^\vee$

!!! [LPR'10] actually defines R-LWE using 'dual' ideal $R^\vee = t^{-1}R$.
'(Non-)Dual' forms are equivalent up to χ, via a 'tweak:' [AP'13]

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
- Error distribution χ over $R^\vee$ (usually Gaussian in 'canonical' geometry)
- Modulus $q \geq 2$ defining $R_q := R/qR = \mathbb{Z}_q[X]/(f(X))$

Search: find secret ring element $s \in R_q^\vee$, given independent samples

$$\begin{array}{lll} a_1 \leftarrow R_q, & b_1 = a_1 \cdot s + e_1 \in R_q^\vee & \\ a_2 \leftarrow R_q, & b_2 = a_2 \cdot s + e_2 \in R_q^\vee & (e_i \leftarrow \chi) \\ & \vdots & \end{array}$$

Decision: distinguish $(a_i\ ,\ b_i)$ from uniform $(a_i\ ,\ b_i) \in R_q \times R_q^\vee$

!!! [LPR'10] actually defines R-LWE using 'dual' ideal $R^\vee = t^{-1}R$.

'(Non-)Dual' forms are equivalent up to χ, via a 'tweak:' [AP'13]

$$b \leftrightarrow t \cdot b \quad \text{induces} \quad s \leftrightarrow t \cdot s,\ e \leftrightarrow t \cdot e.$$

Tweak may dramatically change width and shape of χ!

Ring-LWE Instantiations, Hard and Easy

- 'Dual' R-LWE with wide enough (near-)spherical error is hard:

worst-case approx-SVP on *ideal* lattices in R $\le$ search R-LWE $\le$ decision R-LWE

(quantum, any $R = \mathcal{O}_K$)

(classical, any Galois R)

Ring-LWE Instantiations, Hard and Easy

- 'Dual' R-LWE with wide enough (near-)spherical error is hard:

worst-case approx-SVP on *ideal* lattices in R $\leq$ search R-LWE $\leq$ decision R-LWE

(quantum, any $R = \mathcal{O}_K$) (classical, any Galois R)

- But some other R-LWE instantiations are insecure:

Ring-LWE Instantiations, Hard and Easy

- 'Dual' R-LWE with wide enough (near-)spherical error is hard:

worst-case approx-SVP on *ideal* lattices in R $\leq$ search R-LWE $\leq$ decision R-LWE

(quantum, any $R = \mathcal{O}_K$) (classical, any Galois R)

- But some other R-LWE instantiations are insecure:

[EHL'14] Solves decision-"Poly-LWE" for rings w/ certain properties

Ring-LWE Instantiations, Hard and Easy

- 'Dual' R-LWE with wide enough (near-)spherical error is hard:

> worst-case approx-SVP on *ideal* lattices in R $\leq$ search R-LWE $\leq$ decision R-LWE
>
> (quantum, any $R = \mathcal{O}_K$) (classical, any Galois R)

- But some other R-LWE instantiations are insecure:

[EHL'14] Solves decision-"Poly-LWE" for rings w/ certain properties

[ELOS'15] Solves decision for non-dual, spherical error in certain $R = \mathbb{Z}[X]/(X^n + aX + b)$

Ring-LWE Instantiations, Hard and Easy

- 'Dual' R-LWE with wide enough (near-)spherical error is hard:

> worst-case approx-SVP on *ideal* lattices in R $\leq$ search R-LWE $\leq$ decision R-LWE
>
> (quantum, any $R = \mathcal{O}_K$) (classical, any Galois R)

- But some other R-LWE instantiations are insecure:

[EHL'14] Solves decision-"Poly-LWE" for rings w/ certain properties

[ELOS'15] Solves decision for non-dual, spherical error in certain $R = \mathbb{Z}[X]/(X^n + aX + b)$

[CIV'16] Solves search for [ELOS'15] instantiations, via errorless LWE

Ring-LWE Instantiations, Hard and Easy

- 'Dual' R-LWE with wide enough (near-)spherical error is hard:

> worst-case approx-SVP on *ideal* lattices in R $\leq$ search R-LWE $\leq$ decision R-LWE
>
> (quantum, any $R = \mathcal{O}_K$) (classical, any Galois R)

- But some other R-LWE instantiations are insecure:

[EHL'14] Solves decision-"Poly-LWE" for rings w/ certain properties

[ELOS'15] Solves decision for non-dual, spherical error in certain $R = \mathbb{Z}[X]/(X^n + aX + b)$

[CIV'16] Solves search for [ELOS'15] instantiations, via errorless LWE

[CLS'15,'16] Solves search (via decision) for non-dual, spherical error in certain Galois fields. (Not solvable via errorless LWE.)

What to Make of All This?

Glib answer:

> *The insecure instantiations aren't covered by the worst-case hardness theorems, so all bets are off.*

What to Make of All This?

Glib answer:

> *The insecure instantiations aren't covered by the worst-case hardness theorems, so all bets are off.*

But in practice people often don't use provably hard instantiations; e.g., narrower and/or non-Gaussian error.

What to Make of All This?

Glib answer:

> *The insecure instantiations aren't covered by the worst-case hardness theorems, so all bets are off.*

But in practice people often don't use provably hard instantiations; e.g., narrower and/or non-Gaussian error.

- How "close" are the insecure instantiations to worst-case-hard ones, or those used in practice?

What to Make of All This?

Glib answer:

> *The insecure instantiations aren't covered by the worst-case hardness theorems, so all bets are off.*

But in practice people often don't use provably hard instantiations; e.g., narrower and/or non-Gaussian error.

- How "close" are the insecure instantiations to worst-case-hard ones, or those used in practice?
- Are some kinds of rings inherently less secure for Ring-LWE?

What to Make of All This?

Glib answer:

> *The insecure instantiations aren't covered by the worst-case hardness theorems, so all bets are off.*

But in practice people often don't use provably hard instantiations; e.g., narrower and/or non-Gaussian error.

- How "close" are the insecure instantiations to worst-case-hard ones, or those used in practice?
- Are some kinds of rings inherently less secure for Ring-LWE?
- How can we evaluate the security of Ring-LWE instantiations that aren't supported by hardness theorems?

Contributions and Findings

1. A comprehensive review of prior attacks and insecure instantiations.

Contributions and Findings

1. A comprehensive review of prior attacks and insecure instantiations.
   - New, unified exposition in terms of short elements in dual ideals, and formal analysis that explains prior experimental results.

Contributions and Findings

1. A comprehensive review of prior attacks and insecure instantiations.
   - ⋆ New, unified exposition in terms of short elements in dual ideals, and formal analysis that explains prior experimental results.
   - ⋆ Insecurity is due to use of incongruous error distributions that are insufficiently "well spread" relative to the ring and its ideals.

     In particular, error coeffs have Gaussian parameter $\approx 1 \ll \sqrt{n}$.

Contributions and Findings

1. A comprehensive review of prior attacks and insecure instantiations.
   - ★ New, unified exposition in terms of short elements in dual ideals, and formal analysis that explains prior experimental results.
   - ★ Insecurity is due to use of incongruous error distributions that are insufficiently "well spread" relative to the ring and its ideals.

     In particular, error coeffs have Gaussian parameter $\approx 1 \ll \sqrt{n}$.
2. On the positive side:

Theorem

Any instantiation supported by the "worst-case hardness of search" theorem [LPR'10] (or almost so) is immune to the above class of attacks.

Contributions and Findings

1. A comprehensive review of prior attacks and insecure instantiations.
   - ⋆ New, unified exposition in terms of short elements in dual ideals, and formal analysis that explains prior experimental results.
   - ⋆ Insecurity is due to use of incongruous error distributions that are insufficiently "well spread" relative to the ring and its ideals.

     In particular, error coeffs have Gaussian parameter $\approx 1 \ll \sqrt{n}$.

2. On the positive side:

Theorem

Any instantiation supported by the "worst-case hardness of search" theorem [LPR'10] (or almost so) is immune to the above class of attacks.

- ⋆ Theorem holds for any number ring, so the rings themselves are not the source of weakness in the insecure instantiations.

Contributions and Findings

1. A comprehensive review of prior attacks and insecure instantiations.
   - ★ New, unified exposition in terms of short elements in dual ideals, and formal analysis that explains prior experimental results.
   - ★ Insecurity is due to use of incongruous error distributions that are insufficiently "well spread" relative to the ring and its ideals.

     In particular, error coeffs have Gaussian parameter $\approx 1 \ll \sqrt{n}$.
2. On the positive side:

Theorem

Any instantiation supported by the "worst-case hardness of search" theorem [LPR'10] (or almost so) is immune to the above class of attacks.

- ★ Theorem holds for any number ring, so the rings themselves are not the source of weakness in the insecure instantiations.
- ★ Hard error distributions are much wider & differently shaped than the insecure ones.

Attack Framework [EHL'14,. . .]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q}|qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).

Attack Framework [EHL'14,...]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q} | qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).
2. Given mod-qR samples (a_i, b_i), reduce modulo $\mathfrak{q}$:

$$(a_i' := a_i \bmod \mathfrak{q} \ , \ b_i' := b_i \bmod \mathfrak{q})$$

Attack Framework [EHL'14,...]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q}|qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).
2. Given mod-qR samples (a_i, b_i), reduce modulo $\mathfrak{q}$:

$$(a_i' := a_i \bmod \mathfrak{q} \ , \ b_i' := b_i \bmod \mathfrak{q})$$

3. For each $s' \in R/\mathfrak{q}$, test if $d_i := b_i' - a_i' \cdot s' \bmod \mathfrak{q}$ are non-uniform.

Attack Framework [EHL'14,...]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q} | qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).
2. Given mod-qR samples (a_i, b_i), reduce modulo $\mathfrak{q}$:

$$(a_i' := a_i \bmod \mathfrak{q} \ , \ b_i' := b_i \bmod \mathfrak{q})$$

3. For each $s' \in R/\mathfrak{q}$, test if $d_i := b_i' - a_i' \cdot s' \bmod \mathfrak{q}$ are non-uniform.

Analysis:

- For R-LWE samples and $s' = s \bmod \mathfrak{q}$, we have $d_i = e_i \bmod \mathfrak{q}$.

Attack Framework [EHL'14,...]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q} | qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).
2. Given mod-qR samples (a_i, b_i), reduce modulo $\mathfrak{q}$:

$$(a_i' := a_i \bmod \mathfrak{q} \ , \ b_i' := b_i \bmod \mathfrak{q})$$

3. For each $s' \in R/\mathfrak{q}$, test if $d_i := b_i' - a_i' \cdot s' \bmod \mathfrak{q}$ are non-uniform.

Analysis:

- For R-LWE samples and $s' = s \bmod \mathfrak{q}$, we have $d_i = e_i \bmod \mathfrak{q}$.
- For uniform samples, d_i is uniform.

Attack Framework [EHL'14,...]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q} | qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).
2. Given mod-qR samples (a_i, b_i), reduce modulo $\mathfrak{q}$:

$$(a_i' := a_i \bmod \mathfrak{q} \ , \ b_i' := b_i \bmod \mathfrak{q})$$

3. For each $s' \in R/\mathfrak{q}$, test if $d_i := b_i' - a_i' \cdot s' \bmod \mathfrak{q}$ are non-uniform.

Analysis:

- For R-LWE samples and $s' = s \bmod \mathfrak{q}$, we have $d_i = e_i \bmod \mathfrak{q}$.
- For uniform samples, d_i is uniform.
- So attack succeeds iff $\chi \bmod \mathfrak{q}$ is detectably non-uniform.

Attack Framework [EHL'14,...]

To attack 'non-dual' decision:

1. Fix an ideal $\mathfrak{q} | qR$ having small norm $\mathrm{N}(\mathfrak{q}) = |R/\mathfrak{q}|$ (possibly $\mathfrak{q} = R$).
2. Given mod-qR samples (a_i, b_i), reduce modulo $\mathfrak{q}$:

$$(a_i' := a_i \bmod \mathfrak{q} \ , \ b_i' := b_i \bmod \mathfrak{q})$$

3. For each $s' \in R/\mathfrak{q}$, test if $d_i := b_i' - a_i' \cdot s' \bmod \mathfrak{q}$ are non-uniform.

Analysis:

- For R-LWE samples and $s' = s \bmod \mathfrak{q}$, we have $d_i = e_i \bmod \mathfrak{q}$.
- For uniform samples, d_i is uniform.
- So attack succeeds iff $\chi \bmod \mathfrak{q}$ is detectably non-uniform.

Prior works [EHL'14,ELOS'15,CLS'15,'16] use theory and computer search/experiments to find insecure instantiations.
Some attacks are proven; many are only empirical.

Insecure Instantiations #1 [EHL'15,EHL'16]

- 'Non-dual' over $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$, Gaussian error param $r \approx \sqrt{pd}$.
  'Volume normalized' param $r_0 \approx d^{1/4} \to \infty$.

Insecure Instantiations #1 [EHL'15,EHL'16]

- 'Non-dual' over $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$, Gaussian error param $r \approx \sqrt{pd}$.
  'Volume normalized' param $r_0 \approx d^{1/4} \to \infty$.

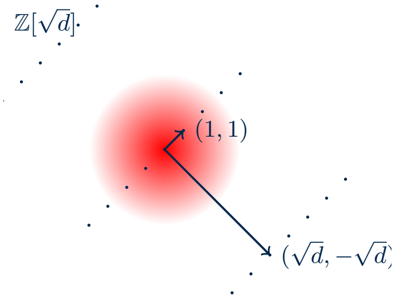

Insecure Instantiations #1 [EHL'15,EHL'16]

- 'Non-dual' over $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$, Gaussian error param $r \approx \sqrt{pd}$.
  'Volume normalized' param $r_0 \approx d^{1/4} \to \infty$.
- $R^\vee$ has $p-1$ elements of length $1/\sqrt{pd}$, so error is narrow and non-uniform mod R: many coeffs have small param ≈ 1.

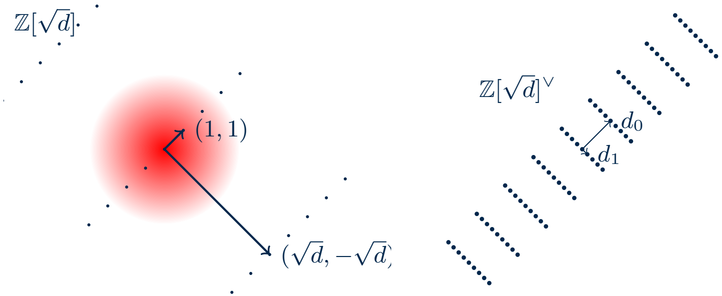

Insecure Instantiations #1 [EHL'15,EHL'16]

- 'Non-dual' over $R = \mathbb{Z}[\zeta_p, \sqrt{d}]$, Gaussian error param $r \approx \sqrt{pd}$.
  'Volume normalized' param $r_0 \approx d^{1/4} \to \infty$.
- $R^\vee$ has $p-1$ elements of length $1/\sqrt{pd}$, so error is narrow and non-uniform mod R: many coeffs have small param ≈ 1.
- Similarly for error mod $\mathfrak{q} \subset R$ (which is even sparser).

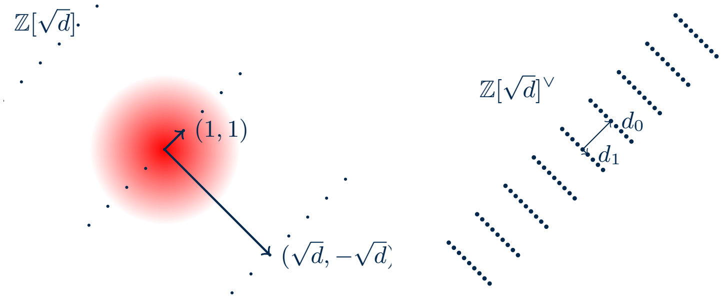

Insecure Instantiations #2

- Take $R = \mathbb{Z}[\zeta_q]$ for prime modulus q; $r \approx \sqrt{q}$. 'Normalized' $r_0 \approx 1$.

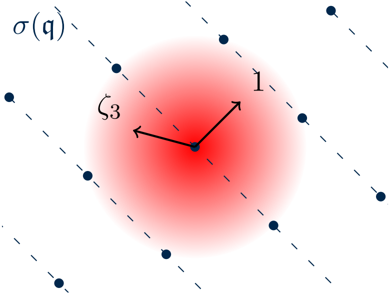

Insecure Instantiations #2

- Take $R = \mathbb{Z}[\zeta_q]$ for prime modulus q; $r \approx \sqrt{q}$. 'Normalized' $r_0 \approx 1$.
- Then $\mathfrak{q} | qR$ where $\mathfrak{q} = (1 - \zeta_q)R$, and $\mathrm{N}(\mathfrak{q}) = q$.

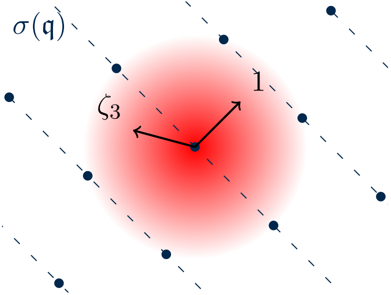

Insecure Instantiations #2

- Take $R = \mathbb{Z}[\zeta_q]$ for prime modulus q; $r \approx \sqrt{q}$. 'Normalized' $r_0 \approx 1$.
- Then $\mathfrak{q} | qR$ where $\mathfrak{q} = (1 - \zeta_q)R$, and $\mathrm{N}(\mathfrak{q}) = q$.
- $q^{-1} \in \mathfrak{q}^\vee = q^{-1}R$, has length $\approx 1/\sqrt{q}$, so error is non-uniform mod $\mathfrak{q}$.

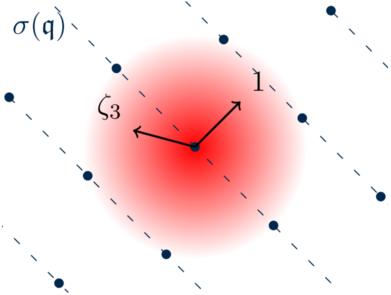

Insecure Instantiations #2

- Take $R = \mathbb{Z}[\zeta_q]$ for prime modulus q; $r \approx \sqrt{q}$. 'Normalized' $r_0 \approx 1$.
- Then $\mathfrak{q} | qR$ where $\mathfrak{q} = (1 - \zeta_q)R$, and $\mathrm{N}(\mathfrak{q}) = q$.
- $q^{-1} \in \mathfrak{q}^\vee = q^{-1}R$, has length $\approx 1/\sqrt{q}$, so error is non-uniform mod $\mathfrak{q}$.
- This formally substantiates empirical observations from [CLS'15].

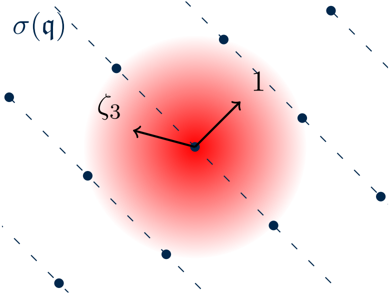

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$.

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$. 'Worst-case hardness of search' theorem applies to any $R = \mathcal{O}_K$, spherical error D_r where $r \gg 2$.

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$. 'Worst-case hardness of search' theorem applies to any $R = \mathcal{O}_K$, spherical error D_r where $r \gg 2$.

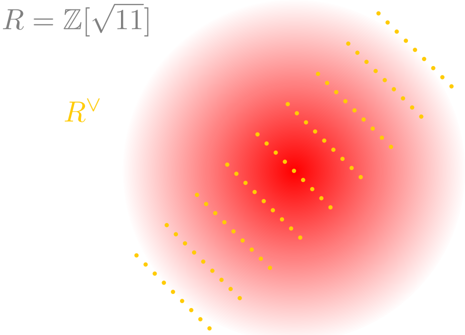

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$. 'Worst-case hardness of search' theorem applies to any $R = \mathcal{O}_K$, spherical error D_r where $r \gg 2$.
- Analogue of attack on 'non-dual' decision is:
  1. for each of the $N(\mathfrak{q})$ candidate $s' \in R^\vee / \mathfrak{q}R^\vee$,
  2. test for non-uniformity of $b_i - a_i \cdot s' \bmod \mathfrak{q}R^\vee$: should be $D_r \bmod \mathfrak{q}R^\vee$

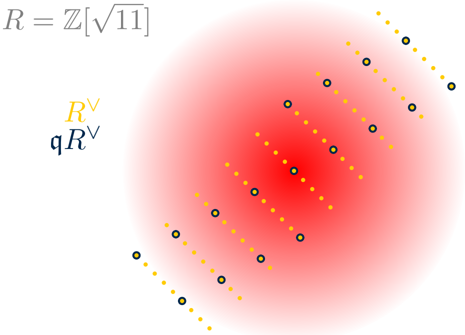

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$.
  'Worst-case hardness of search' theorem applies to any $R = \mathcal{O}_K$, spherical error D_r where $r \gg 2$.
- Analogue of attack on 'non-dual' decision is:
  1. for each of the $N(\mathfrak{q})$ candidate $s' \in R^\vee/\mathfrak{q}R^\vee$,
  2. test for non-uniformity of $b_i - a_i \cdot s' \bmod \mathfrak{q}R^\vee$: should be $D_r \bmod \mathfrak{q}R^\vee$

Theorem

For $\mathrm{N}(\mathfrak{q}) \leq 2^n$, reduced error $D_r \bmod \mathfrak{q}R^\vee$ is only 4^{-n}-far from uniform.

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$. 'Worst-case hardness of search' theorem applies to any $R = \mathcal{O}_K$, spherical error D_r where $r \gg 2$.
- Analogue of attack on 'non-dual' decision is:
  1. for each of the $N(\mathfrak{q})$ candidate $s' \in R^\vee/\mathfrak{q}R^\vee$,
  2. test for non-uniformity of $b_i - a_i \cdot s' \bmod \mathfrak{q}R^\vee$: should be $D_r \bmod \mathfrak{q}R^\vee$

Theorem

For $\mathrm{N}(\mathfrak{q}) \leq 2^n$, reduced error $D_r \bmod \mathfrak{q}R^\vee$ is only 4^{-n}-far from uniform.

Proof Idea

- Dual ideal of $\mathfrak{q}R^\vee$ is $\mathfrak{q}^{-1}$, which has $\lambda_1(\mathfrak{q}^{-1}) \geq \sqrt{n}/2$.

Invulnerability of Worst-Case-Hard Instantiations

- Recall that [LPR'10] defines 'dual' form: χ, s, b_i are modulo $qR^\vee$. 'Worst-case hardness of search' theorem applies to any $R = \mathcal{O}_K$, spherical error D_r where $r \gg 2$.
- Analogue of attack on 'non-dual' decision is:
  1. for each of the $N(\mathfrak{q})$ candidate $s' \in R^\vee/\mathfrak{q}R^\vee$,
  2. test for non-uniformity of $b_i - a_i \cdot s' \bmod \mathfrak{q}R^\vee$: should be $D_r \bmod \mathfrak{q}R^\vee$

Theorem

For $\mathrm{N}(\mathfrak{q}) \leq 2^n$, reduced error $D_r \bmod \mathfrak{q}R^\vee$ is only 4^{-n}-far from uniform.

Proof Idea

- Dual ideal of $\mathfrak{q}R^\vee$ is $\mathfrak{q}^{-1}$, which has $\lambda_1(\mathfrak{q}^{-1}) \geq \sqrt{n}/2$.
- So 'smoothing parameter' of $\mathfrak{q}R^\vee$ is ≤ 2, so $D_r \bmod \mathfrak{q}R^\vee$ is uniform.

Conclusions and Parting Thoughts

- Choice of error distribution for Ring-LWE is subtler than for LWE: must account for geometry of ring and its ideals.

Conclusions and Parting Thoughts

- Choice of error distribution for Ring-LWE is subtler than for LWE: must account for geometry of ring and its ideals.
- Some attacks need qR to have small-norm divisors, but it seems prudent not to rely on (lack of) factorization for security.
  Can 'ideal switching' make factorization irrelevant?

Conclusions and Parting Thoughts

- Choice of error distribution for Ring-LWE is subtler than for LWE: must account for geometry of ring and its ideals.
- Some attacks need qR to have small-norm divisors, but it seems prudent not to rely on (lack of) factorization for security. Can 'ideal switching' make factorization irrelevant?
- Worst-case hardness theorems yield (nearly) minimal conditions for invulnerability to a new class of attacks.

Conclusions and Parting Thoughts

- Choice of error distribution for Ring-LWE is subtler than for LWE: must account for geometry of ring and its ideals.
- Some attacks need qR to have small-norm divisors, but it seems prudent not to rely on (lack of) factorization for security.
  Can 'ideal switching' make factorization irrelevant?
- Worst-case hardness theorems yield (nearly) minimal conditions for invulnerability to a new class of attacks.

Thanks!

`http://eprint.iacr.org/2016/351`