Lattice Assumptions in Crypto: Status Update

Chris Peikert
University of Michigan

(covers work with Oded Regev and Noah Stephens-Davidowitz
to appear, STOC'17)

10 March 2017

Lattice-Based Cryptography

$y = g^x \bmod p$

$m^e \bmod N$

$N = p \cdot q$

$e(g^a, g^b)$

$\Longrightarrow$

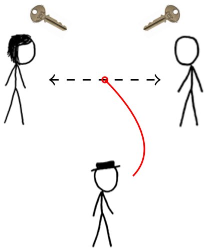

(Images courtesy xkcd.org)

Lattice-Based Cryptography

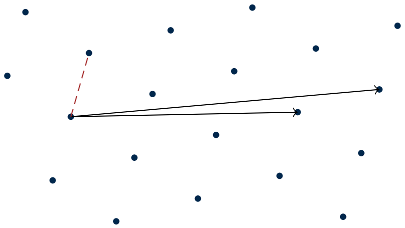

$\Longrightarrow$

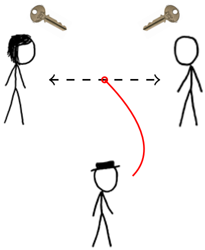

(Images courtesy xkcd.org)

Lattice-Based Cryptography

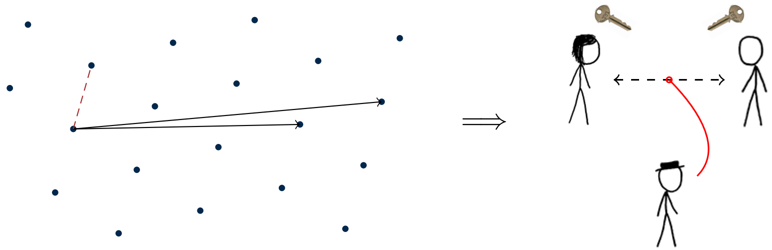

Main Attractions

- Efficient: linear, embarrassingly parallel operations

(Images courtesy xkcd.org)

Lattice-Based Cryptography

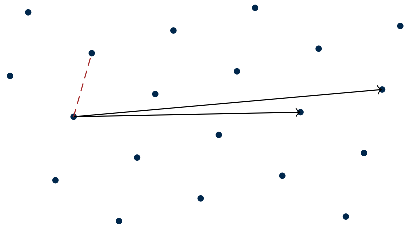

$\Longrightarrow$

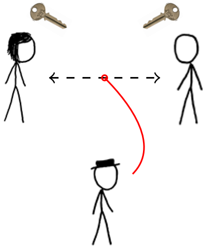

Main Attractions

- Efficient: linear, embarrassingly parallel operations
- Resists quantum attacks (so far)

(Images courtesy xkcd.org)

Lattice-Based Cryptography

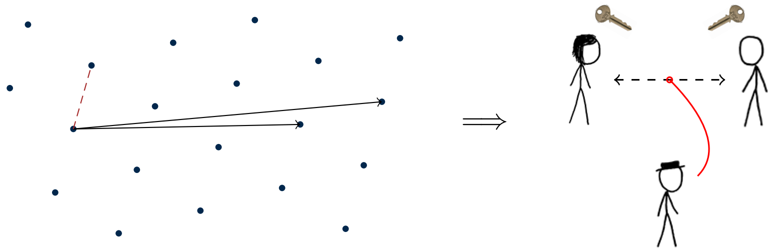

Main Attractions

- Efficient: linear, embarrassingly parallel operations
- Resists quantum attacks (so far)
- Security from worst-case assumptions

(Images courtesy xkcd.org)

Lattice-Based Cryptography

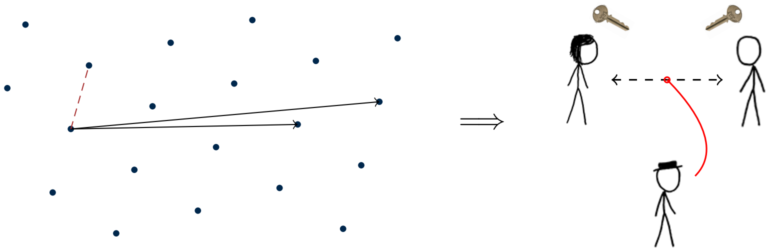

Main Attractions

- Efficient: linear, embarrassingly parallel operations
- Resists quantum attacks (so far)
- Security from worst-case assumptions
- Solutions to ‘holy grail’ problems in crypto: FHE and related

(Images courtesy xkcd.org)

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus q, error 'rate' α

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus q, error 'rate' α
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

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

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus q, error 'rate' α
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

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

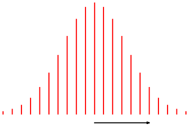

width αq

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus q, error 'rate' α
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

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

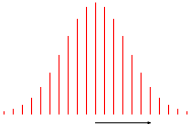

width αq

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

Learning With Errors [Regev'05]

- Parameters: dimension n, integer modulus q, error 'rate' α
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

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

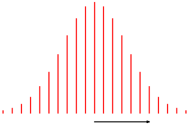

- **Decision:** distinguish (a_i , b_i) from uniform (a_i , b_i)

LWE is Hard and Versatile

worst case

(n/α)-SIVP on n-dim lattices $\leq$ 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, error 'rate' α
- **Search:** find secret $\mathbf{s} \in \mathbb{Z}_q^n$ given many 'noisy inner products'

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

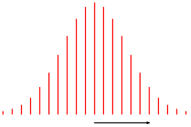

width αq

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

LWE is Hard and Versatile

worst case

(n/α)-SIVP on n-dim lattices $\leq$ search-LWE $\leq$ decision-LWE $\leq$ much crypto

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

- *Classically*, GapSVP $\leq$ search-LWE (worse params) [P'09,BLPRS'13]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - ⋆ Any prime $q = \text{poly}(n)$ [R'05]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - ★ Any prime $q = \text{poly}(n)$ [R'05]
  - ★ Any "somewhat smooth" $q = p_1 \cdots p_t$ (large enough primes p_i) [P'09]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - Any prime $q = \text{poly}(n)$ [R'05]
  - Any "somewhat smooth" $q = p_1 \cdots p_t$ (large enough primes p_i) [P'09]
  - Any $q = p^e$ for large enough prime p [ACPS'09]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - ★ Any prime $q = \text{poly}(n)$ [R'05]
  - ★ Any "somewhat smooth" $q = p_1 \cdots p_t$ (large enough primes p_i) [P'09]
  - ★ Any $q = p^e$ for large enough prime p [ACPS'09]
  - ★ Any $q = p^e$ with uniform error mod p^i [MM'11]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - ★ Any prime $q = \mathrm{poly}(n)$ [R'05]
  - ★ Any "somewhat smooth" $q = p_1 \cdots p_t$ (large enough primes p_i) [P'09]
  - ★ Any $q = p^e$ for large enough prime p [ACPS'09]
  - ★ Any $q = p^e$ with uniform error mod p^i [MM'11]
  - ★ Any $q = p^e$ — but increases α [MP'12]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - ★ Any prime $q = \text{poly}(n)$ [R'05]
  - ★ Any "somewhat smooth" $q = p_1 \cdots p_t$ (large enough primes p_i) [P'09]
  - ★ Any $q = p^e$ for large enough prime p [ACPS'09]
  - ★ Any $q = p^e$ with uniform error mod p^i [MM'11]
  - ★ Any $q = p^e$ — but increases α [MP'12]
  - ★ Any q via "mod-switching" — but increases α [P'09,BV'11,BLPRS'13]

LWE Hardness and Parameters

- Parameters: dimension n, integer modulus q, error 'rate' α

Worst case SIVP $\leq$ Search-LWE

- One reduction for best known parameters: any $q \geq \sqrt{n}/\alpha$ [R'05]

Search-LWE $\leq$ Decision-LWE

- Messy. Many incomparable reductions for different forms of q:
  - ★ Any prime $q = \text{poly}(n)$ [R'05]
  - ★ Any "somewhat smooth" $q = p_1 \cdots p_t$ (large enough primes p_i) [P'09]
  - ★ Any $q = p^e$ for large enough prime p [ACPS'09]
  - ★ Any $q = p^e$ with uniform error mod p^i [MM'11]
  - ★ Any $q = p^e$ — but increases α [MP'12]
  - ★ Any q via "mod-switching" — but increases α [P'09,BV'11,BLPRS'13]
- Increasing q, α yields a weaker ultimate hardness guarantee.

LWE is Efficient (Sort Of)

$$\begin{pmatrix} \cdots & \mathbf{a}_i & \cdots \end{pmatrix} \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + e \; = \; b \in \mathbb{Z}_q$$

- Getting one pseudorandom scalar requires an n-dim inner product mod q

LWE is Efficient (Sort Of)

$$\left(\cdots \mathbf{a}_i \cdots\right) \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + e \; = \; b \in \mathbb{Z}_q$$

- Getting one pseudorandom scalar requires an n-dim inner product mod q
- Can amortize each $\mathbf{a}_i$ over many secrets $\mathbf{s}_j$, but still $\tilde{O}(n)$ work per scalar output.

LWE is Efficient (Sort Of)

$$\begin{pmatrix} \cdots & \mathbf{a}_i & \cdots \end{pmatrix} \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + e \; = \; b \in \mathbb{Z}_q$$

- Getting one pseudorandom scalar requires an n-dim inner product mod q
- Can amortize each $\mathbf{a}_i$ over many secrets $\mathbf{s}_j$, but still $\tilde{O}(n)$ work per scalar output.

- Cryptosystems have rather large keys: $\Omega(n^2 \log^2 q)$ bits:

$$pk = \left.\underbrace{\begin{pmatrix} & \vdots & \\ & \mathbf{A} & \\ & \vdots & \end{pmatrix}}_{n} \quad , \quad \begin{pmatrix} \vdots \\ \mathbf{b} \\ \vdots \end{pmatrix}\right\} \Omega(n)$$

Wishful Thinking...

$$\begin{pmatrix} \vdots \\ \mathbf{a}_i \\ \vdots \end{pmatrix} \star \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + \begin{pmatrix} \vdots \\ \mathbf{e}_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{b}_i \\ \vdots \end{pmatrix} \in \mathbb{Z}_q^n$$

- Get n pseudorandom scalars from just one cheap product operation?

Wishful Thinking...

$$\begin{pmatrix} \vdots \\ \mathbf{a}_i \\ \vdots \end{pmatrix} \star \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + \begin{pmatrix} \vdots \\ \mathbf{e}_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{b}_i \\ \vdots \end{pmatrix} \in \mathbb{Z}_q^n$$

- Get n pseudorandom scalars from just one cheap product operation?

Question

- How to define the product '$\star$' so that $(\mathbf{a}_i, \mathbf{b}_i)$ is pseudorandom?

Wishful Thinking...

$$\begin{pmatrix} \vdots \\ \mathbf{a}_i \\ \vdots \end{pmatrix} \star \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + \begin{pmatrix} \vdots \\ \mathbf{e}_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{b}_i \\ \vdots \end{pmatrix} \in \mathbb{Z}_q^n$$

- Get n pseudorandom scalars from just one cheap product operation?

Question

- How to define the product '$\star$' so that $(\mathbf{a}_i, \mathbf{b}_i)$ is pseudorandom?
- Careful! With small error, coordinate-wise multiplication is insecure!

Wishful Thinking...

$$\begin{pmatrix} \vdots \\ \mathbf{a}_i \\ \vdots \end{pmatrix} \star \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + \begin{pmatrix} \vdots \\ \mathbf{e}_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{b}_i \\ \vdots \end{pmatrix} \in \mathbb{Z}_q^n$$

- Get n pseudorandom scalars from just one cheap product operation?

Question

- How to define the product '$\star$' so that $(\mathbf{a}_i, \mathbf{b}_i)$ is pseudorandom?
- Careful! With small error, coordinate-wise multiplication is insecure!

Answer

- '$\star$' $=$ multiplication in a polynomial ring: e.g., $\mathbb{Z}_q[X]/(X^n + 1)$.

  Fast and practical with FFT: $n \log n$ operations mod q.

Wishful Thinking...

$$\begin{pmatrix} \vdots \\ \mathbf{a}_i \\ \vdots \end{pmatrix} \star \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + \begin{pmatrix} \vdots \\ \mathbf{e}_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{b}_i \\ \vdots \end{pmatrix} \in \mathbb{Z}_q^n$$

- Get n pseudorandom scalars from just one cheap product operation?

Question

- How to define the product '$\star$' so that $(\mathbf{a}_i, \mathbf{b}_i)$ is pseudorandom?
- Careful! With small error, coordinate-wise multiplication is insecure!

Answer

- '$\star$' $=$ multiplication in a polynomial ring: e.g., $\mathbb{Z}_q[X]/(X^n + 1)$.

  Fast and practical with FFT: $n \log n$ operations mod q.
- Same ring structures used in NTRU cryptosystem [HPS'98],

  & in compact one-way / CR hash functions [Mic'02,PR'06,LM'06,...]

Wishful Thinking...

$$\begin{pmatrix} \vdots \\ \mathbf{a}_i \\ \vdots \end{pmatrix} \star \begin{pmatrix} \vdots \\ \mathbf{s} \\ \vdots \end{pmatrix} + \begin{pmatrix} \vdots \\ \mathbf{e}_i \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{b}_i \\ \vdots \end{pmatrix} \in \mathbb{Z}_q^n$$

- Get n pseudorandom scalars from just one cheap product operation?

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 $R = \mathcal{O}_K$)
  Has a 'dual ideal' $R^\vee$ (w.r.t. '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 $R = \mathcal{O}_K$)
  Has a 'dual ideal' $R^\vee$ (w.r.t. 'canonical' geometry)
- Integer modulus q defining $R_q := R/qR$ and $R_q^\vee := R^\vee/qR^\vee$

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 $R = \mathcal{O}_K$)
  Has a 'dual ideal' $R^\vee$ (w.r.t. 'canonical' geometry)
- Integer modulus q defining $R_q := R/qR$ and $R_q^\vee := R^\vee/qR^\vee$
- Gaussian error of width $\approx \alpha q$ over $R^\vee$

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 $R = \mathcal{O}_K$)
  Has a 'dual ideal' $R^\vee$ (w.r.t. 'canonical' geometry)
- Integer modulus q defining $R_q := R/qR$ and $R_q^\vee := R^\vee/qR^\vee$
- Gaussian error of width $\approx \alpha q$ over $R^\vee$

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

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

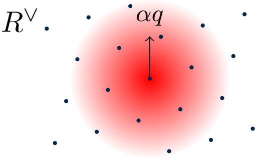

Learning With Errors over Rings (Ring-LWE) [LPR'10]

- Ring R, often $R = \mathbb{Z}[X]/(f(X))$ for irred. f of degree n (or $R = \mathcal{O}_K$)
  Has a 'dual ideal' $R^\vee$ (w.r.t. 'canonical' geometry)
- Integer modulus q defining $R_q := R/qR$ and $R_q^\vee := R^\vee/qR^\vee$
- Gaussian error of width $\approx \alpha q$ over $R^\vee$

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

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

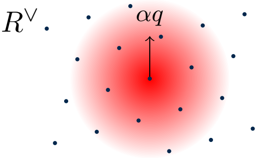

Decision: distinguish $(a_i\ ,\ b_i)$ from uniform $(a_i\ ,\ b_i) \in R_q \times R_q^\vee$

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

(classical, any *Galois* R)

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

(classical, any *Galois* R)

(Ideal $\mathcal{I} \subseteq R$: additive subgroup, $x \cdot r \in \mathcal{I}$ for all $x \in \mathcal{I}, r \in R$.)

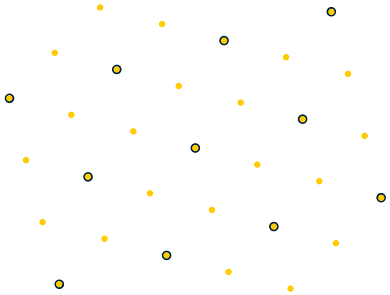

$R = \mathbb{Z}[X]/(1 + X + X^2)$

ideal $\mathcal{I} = 3R + (1 - X)R \subset R$

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

(classical, any *Galois* R)

Large disparity in known hardness of search versus decision:

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

(classical, any *Galois* R)

Large disparity in known hardness of search versus decision:

Search: any number ring, any $q \geq n^c/\alpha$.

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

(classical, any *Galois* R)

Large disparity in known hardness of search versus decision:

Search: any number ring, any $q \geq n^c/\alpha$.

Decision: any Galois number ring (e.g., cyclotomic),
any highly splitting prime $q = \mathrm{poly}(n)$.

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

Large disparity in known hardness of search versus decision:

Search: any number ring, any $q \geq n^c/\alpha$.

Decision: any Galois number ring (e.g., cyclotomic), any highly splitting prime $q = \mathrm{poly}(n)$.

Can then get any q by mod-switching, but increases α [LS'15]

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

(classical, any *Galois* R)

Large disparity in known hardness of search versus decision:

Search: any number ring, any $q \geq n^c/\alpha$.

Decision: any Galois number ring (e.g., cyclotomic), any highly splitting prime $q = \text{poly}(n)$.

Can then get any q by mod-switching, but increases α [LS'15]

- Decision has no known worst-case hardness in non-Galois rings.

Hardness of Ring-LWE [LPR'10]

worst-case (n^c/α)-SIVP on *ideal* lattices in R $\leq$ search R-LWE$_{q,\alpha}$ $\leq$ decision R-LWE$_{q,\alpha}$

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

Large disparity in known hardness of search versus decision:

Search: any number ring, any $q \geq n^c/\alpha$.

Decision: any Galois number ring (e.g., cyclotomic), any highly splitting prime $q = \text{poly}(n)$.

Can then get any q by mod-switching, but increases α [LS'15]

- Decision has no known worst-case hardness in non-Galois rings.
- But no examples of easy(er) decision when search is worst-case hard!

New Results [PRS'17]

Main Theorem: Ring-LWE is Pseudorandom in Any Ring

worst-case (n^c/α)-SIVP on ideal lattices in R $\leq$ decision R-LWE$_{q,\alpha}$

quantum,
any $R = \mathcal{O}_K$, any $q \geq n^{c-1/2}/\alpha$

New Results [PRS'17]

Main Theorem: Ring-LWE is Pseudorandom in Any Ring

worst-case (n^c/α)-SIVP on ideal lattices in R $\leq$ decision R-LWE$_{q,\alpha}$

quantum,
any $R = \mathcal{O}_K$, *any* $q \geq n^{c-1/2}/\alpha$

Bonus Theorem: LWE is Pseudorandom for Any Modulus

worst case (n/α)-SIVP on n-dim lattices $\leq$ decision-LWE$_{q,\alpha}$

quantum, *any* $q \geq \sqrt{n}/\alpha$

New Results [PRS'17]

Main Theorem: Ring-LWE is Pseudorandom in Any Ring

worst-case (n^c/α)-SIVP on ideal lattices in R $\leq$ decision R-LWE$_{q,\alpha}$

quantum,
any $R = \mathcal{O}_K$, any $q \geq n^{c-1/2}/\alpha$

Bonus Theorem: LWE is Pseudorandom for Any Modulus

worst case (n/α)-SIVP on n-dim lattices $\leq$ decision-LWE$_{q,\alpha}$

quantum, any $q \geq \sqrt{n}/\alpha$

- Both theorems match or improve the previous best params:

New Results [PRS'17]

Main Theorem: Ring-LWE is Pseudorandom in Any Ring

worst-case (n^c/α)-SIVP on ideal lattices in R $\leq$ decision R-$\mathsf{LWE}_{q,\alpha}$

quantum, any $R = \mathcal{O}_K$, any $q \geq n^{c-1/2}/\alpha$

Bonus Theorem: LWE is Pseudorandom for Any Modulus

worst case (n/α)-SIVP on n-dim lattices $\leq$ decision-$\mathsf{LWE}_{q,\alpha}$

quantum, any $q \geq \sqrt{n}/\alpha$

- Both theorems match or improve the previous best params:

One reduction to rule them all.

New Results [PRS'17]

Main Theorem: Ring-LWE is Pseudorandom in Any Ring

worst-case (n^c/α)-SIVP on ideal lattices in R $\leq$ decision R-LWE$_{q,\alpha}$

quantum,
any $R = \mathcal{O}_K$, *any* $q \geq n^{c-1/2}/\alpha$

Bonus Theorem: LWE is Pseudorandom for Any Modulus

worst case (n/α)-SIVP on n-dim lattices $\leq$ decision-LWE$_{q,\alpha}$

quantum, *any* $q \geq \sqrt{n}/\alpha$

- Both theorems match or improve the previous best params:

  One reduction to rule them all.

- Seems to adapt to 'module' lattices/LWE w/techniques from [LS'15]

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Progress on Ideal-SIVP

- Quantum poly-time $\exp(\tilde{O}(\sqrt{n}))$-Ideal-SIVP in prime-power cyclotomics (modulo heuristics) [CGS'14,BS'16,CDPR'16,CDW'17]

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Progress on Ideal-SIVP

- Quantum poly-time $\exp(\tilde{O}(\sqrt{n}))$-Ideal-SIVP in prime-power cyclotomics (modulo heuristics) [CGS'14,BS'16,CDPR'16,CDW'17]
- Quite far from the (quasi-)poly(n) factors typically used for crypto

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Progress on Ideal-SIVP

- Quantum poly-time $\exp(\tilde{O}(\sqrt{n}))$-Ideal-SIVP in prime-power cyclotomics (modulo heuristics) [CGS'14,BS'16,CDPR'16,CDW'17]
- Quite far from the (quasi-)poly(n) factors typically used for crypto
- Doesn't apply to R-LWE or NTRU (unknown if R-LWE $\leq$ Ideal-SIVP)

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Progress on Ideal-SIVP

- Quantum poly-time $\exp(\tilde{O}(\sqrt{n}))$-Ideal-SIVP in prime-power cyclotomics (modulo heuristics) [CGS'14,BS'16,CDPR'16,CDW'17]
- Quite far from the (quasi-)poly(n) factors typically used for crypto
- Doesn't apply to R-LWE or NTRU (unknown if R-LWE $\leq$ Ideal-SIVP)

Options

- Keep using R-LWE over cyclotomics

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Progress on Ideal-SIVP

- Quantum poly-time $\exp(\tilde{O}(\sqrt{n}))$-Ideal-SIVP in prime-power cyclotomics (modulo heuristics) [CGS'14,BS'16,CDPR'16,CDW'17]
- Quite far from the (quasi-)poly(n) factors typically used for crypto
- Doesn't apply to R-LWE or NTRU (unknown if R-LWE $\leq$ Ideal-SIVP)

Options

- Keep using R-LWE over cyclotomics
- Use R-LWE over (slower) rings like $\mathbb{Z}[X]/(X^p - X - 1)$ [BCLvV'16]

Which Rings To Use?

- Our results don't give any guidance: they work within a single ring R, lower-bounding the hardness of R-LWE by R-Ideal-SIVP
- We have no nontrivial relations between lattice problems over different rings. (Great open question!)

Progress on Ideal-SIVP

- Quantum poly-time $\exp(\tilde{O}(\sqrt{n}))$-Ideal-SIVP in prime-power cyclotomics (modulo heuristics) [CGS'14,BS'16,CDPR'16,CDW'17]
- Quite far from the (quasi-)poly(n) factors typically used for crypto
- Doesn't apply to R-LWE or NTRU (unknown if R-LWE $\leq$ Ideal-SIVP)

Options

- Keep using R-LWE over cyclotomics
- Use R-LWE over (slower) rings like $\mathbb{Z}[X]/(X^p - X - 1)$ [BCLvV'16]
- Use 'higher rank' problem Module-LWE over cyclotomics/others

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- **Reduction strategy**: 'play with' α, detect when it decreases.

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- **Reduction strategy**: 'play with' α, detect when it decreases.

Suppose $\mathcal{O}$ solves decision-LWE$_{q,\alpha}$ with non-negl advantage. Define

$$p(\beta) = \Pr[\mathcal{O} \text{ accepts on LWE}_{q,\exp(\beta)} \text{ samples}].$$

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- **Reduction strategy**: 'play with' α, detect when it decreases.

Suppose $\mathcal{O}$ solves decision-LWE$_{q,\alpha}$ with non-negl advantage. Define

$$p(\beta) = \Pr[\mathcal{O} \text{ accepts on LWE}_{q,\exp(\beta)} \text{ samples}].$$

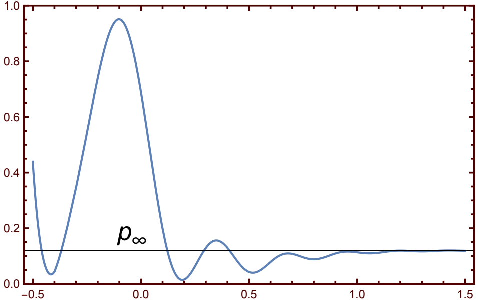

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- **Reduction strategy**: 'play with' α, detect when it decreases.

Suppose $\mathcal{O}$ solves decision-LWE$_{q,\alpha}$ with non-negl advantage. Define

$$p(\beta) = \Pr[\mathcal{O} \text{ accepts on LWE}_{q,\exp(\beta)} \text{ samples}].$$

Key Properties

1. $p(\beta)$ is 'smooth' (Lipschitz) because D_σ, D_τ are $(\frac{\tau}{\sigma} - 1)$-close.

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- **Reduction strategy**: 'play with' α, detect when it decreases.

Suppose $\mathcal{O}$ solves decision-LWE$_{q,\alpha}$ with non-negl advantage. Define

$$p(\beta) = \Pr[\mathcal{O} \text{ accepts on LWE}_{q,\exp(\beta)} \text{ samples}].$$

Key Properties

1. $p(\beta)$ is 'smooth' (Lipschitz) because D_σ, D_τ are $(\frac{\tau}{\sigma} - 1)$-close.
2. For all $\beta \geq \log n$, $p(\beta) \approx p(\infty) = \Pr[\mathcal{O}$ accepts on uniform samples$]$, because huge Gaussian error is near-uniform mod $q\mathbb{Z}$.

Overview of LWE Reduction

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- **Reduction strategy**: 'play with' α, detect when it decreases.

Suppose $\mathcal{O}$ solves decision-LWE$_{q,\alpha}$ with non-negl advantage. Define

$$p(\beta) = \Pr[\mathcal{O} \text{ accepts on LWE}_{q,\exp(\beta)} \text{ samples}].$$

Key Properties

1. $p(\beta)$ is 'smooth' (Lipschitz) because D_σ, D_τ are $(\frac{\tau}{\sigma} - 1)$-close.
2. For all $\beta \geq \log n$, $p(\beta) \approx p(\infty) = \Pr[\mathcal{O} \text{ accepts on uniform samples}]$, because huge Gaussian error is near-uniform mod $q\mathbb{Z}$.
3. $p(\log \alpha) - p(\infty)$ is noticeable, so there is a noticeable change in p somewhere between $\log \alpha$ and $\log n$.

Exploiting the Oracle

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$

Exploiting the Oracle

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- Classical part of [Regev'05] reduction:

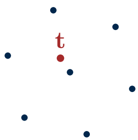

BDD$_{\mathcal{L}^*}$, dist d

$+$

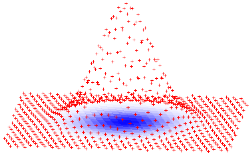

$D_{\mathcal{L},r}$ samples

$\Longrightarrow$ LWE$_{q,\alpha}$ samples

$\alpha = dr/q$

Exploiting the Oracle

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- Classical part of [Regev'05] reduction:

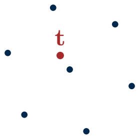

BDD$_{\mathcal{L}^*}$, dist d

$+$

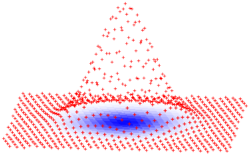

$D_{\mathcal{L},r}$ samples

$\Longrightarrow$ LWE$_{q,\alpha}$ samples

$\alpha = dr/q$

($D_{\mathcal{L},r}$ samples come from previous iteration, quantumly. They're eventually narrow enough to solve SIVP on $\mathcal{L}$.)

Exploiting the Oracle

- **Theorem:** quantumly, (n/α)-SIVP $\leq$ decision-LWE$_{q,\alpha}$ $\quad \forall\, q \geq \sqrt{n}/\alpha$
- Classical part of [Regev'05] reduction:

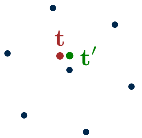

$+$

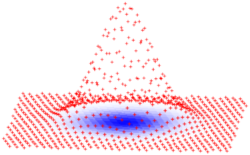

$\Longrightarrow$ LWE$_{q,\alpha}$ samples

$\alpha = dr/q$

BDD$_{\mathcal{L}^*}$, dist d $\qquad$ $D_{\mathcal{L},r}$ samples

- Idea: perturb $\mathbf{t}$, use $\mathcal{O}$ to check whether we're closer to $\mathcal{L}^*$ by how $\alpha = dr/q$ changes.

  We get a 'suffix' of $p(\cdot)$.

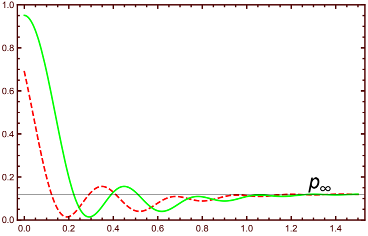

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α
- R-LWE proof has n-parameter BDD offset $\mathbf{e} \Leftrightarrow$ params $\boldsymbol{\alpha} = (\alpha_i)$. Gaussian error rate of α_i in the ith dimension.

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α
- R-LWE proof has n-parameter BDD offset $\mathbf{e} \Leftrightarrow$ params $\boldsymbol{\alpha} = (\alpha_i)$. Gaussian error rate of α_i in the ith dimension.
- Classical part of [LPR'10] reduction:

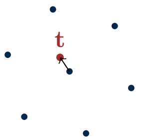

BDD$_{\mathcal{I}^*}$, offset $\mathbf{e}$

$+$

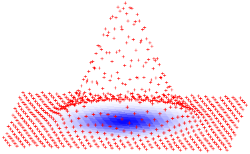

$D_{\mathcal{I},\mathbf{r}}$ samples

$\Longrightarrow$

R-LWE$_{q,\boldsymbol{\alpha}}$ samples

$\alpha_i = |e_i| r_i / q$

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α
- R-LWE proof has n-parameter BDD offset $\mathbf{e} \Leftrightarrow$ params $\boldsymbol{\alpha} = (\alpha_i)$. Gaussian error rate of α_i in the ith dimension.
- Classical part of [LPR'10] reduction:

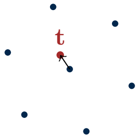

BDD$_{\mathcal{I}^*}$, offset $\mathbf{e}$

$+$

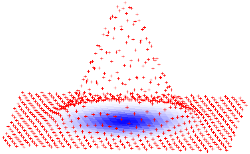

$D_{\mathcal{I},\mathbf{r}}$ samples

$\Longrightarrow$

R-LWE$_{q,\boldsymbol{\alpha}}$ samples

$\alpha_i = |e_i| r_i / q$

- Now oracle's acceptance prob. is $p(\boldsymbol{\beta})$, mapping $(\mathbb{R}^+)^n \to [0,1]$.
  - ⋆ $\lim_{\beta_i \to \infty} p(\boldsymbol{\beta}) = p(\infty)$: huge error in one dim is 'smooth' mod $R^\vee$.

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α
- R-LWE proof has n-parameter BDD offset $\mathbf{e} \Leftrightarrow$ params $\boldsymbol{\alpha} = (\alpha_i)$. Gaussian error rate of α_i in the ith dimension.
- Classical part of [LPR'10] reduction:

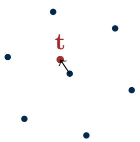

$+$

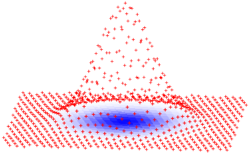

$\Longrightarrow$ R-LWE$_{q,\boldsymbol{\alpha}}$ samples

$\alpha_i = |e_i| r_i / q$

BDD$_{\mathcal{I}^*}$, offset $\mathbf{e}$ $\qquad$ $D_{\mathcal{I},\mathbf{r}}$ samples

- Now oracle's acceptance prob. is $p(\boldsymbol{\beta})$, mapping $(\mathbb{R}^+)^n \to [0,1]$.
  - ⋆ $\lim_{\beta_i \to \infty} p(\boldsymbol{\beta}) = p(\infty)$: huge error in one dim is 'smooth' mod $R^\vee$.
  - ⋆ Problem: Reduction never* produces spherical error (all α_i equal), so it's hard to get anything useful from $\mathcal{O}$.

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α
- R-LWE proof has n-parameter BDD offset $\mathbf{e} \Leftrightarrow$ params $\boldsymbol{\alpha} = (\alpha_i)$. Gaussian error rate of α_i in the ith dimension.
- Classical part of [LPR'10] reduction:

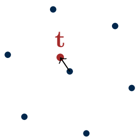

$+$

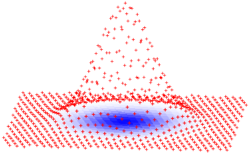

$\Longrightarrow$ R-LWE$_{q,\boldsymbol{\alpha}}$ samples

$\alpha_i = |e_i| r_i / q$

BDD$_{\mathcal{I}^*}$, offset $\mathbf{e}$ $\quad$ $D_{\mathcal{I},\mathbf{r}}$ samples

- Now oracle's acceptance prob. is $p(\boldsymbol{\beta})$, mapping $(\mathbb{R}^+)^n \to [0,1]$.
  - ★ $\lim_{\beta_i \to \infty} p(\boldsymbol{\beta}) = p(\infty)$: huge error in one dim is 'smooth' mod $R^\vee$.
  - ★ Problem: Reduction never* produces spherical error (all α_i equal), so it's hard to get anything useful from $\mathcal{O}$.
  - ★ Solution from [LPR'10]: randomize the α_i: increase by $n^{1/4}$ factor.

Extending to the Ring Setting

- The LWE proof relies on 1-parameter BDD distance $d \Leftrightarrow$ error rate α
- R-LWE proof has n-parameter BDD offset $\mathbf{e} \Leftrightarrow$ params $\boldsymbol{\alpha} = (\alpha_i)$.
  Gaussian error rate of α_i in the ith dimension.
- Classical part of [LPR'10] reduction:

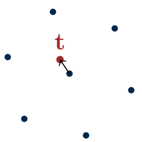

$+$

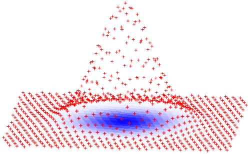

$\Longrightarrow$ R-LWE$_{q,\boldsymbol{\alpha}}$ samples

$\alpha_i = |e_i| r_i / q$

BDD$_{\mathcal{I}^*}$, offset $\mathbf{e}$ $\qquad$ $D_{\mathcal{I},\mathbf{r}}$ samples

- Now oracle's acceptance prob. is $p(\boldsymbol{\beta})$, mapping $(\mathbb{R}^+)^n \to [0,1]$.
  - ⋆ $\lim_{\beta_i \to \infty} p(\boldsymbol{\beta}) = p(\infty)$: huge error in one dim is 'smooth' mod $R^\vee$.
  - ⋆ Problem: Reduction never* produces spherical error (all α_i equal), so it's hard to get anything useful from $\mathcal{O}$.
  - ⋆ Solution from [LPR'10]: randomize the α_i: increase by $n^{1/4}$ factor.
  - ⋆ Improvement: randomization increases α_i by only $\omega(1)$ factor.

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Open Questions

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Open Questions

1. Hardness for spherical error:
   - Avoid $n^{1/4}$ degradation in the α_i rates?
   - Support unbounded samples?

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Open Questions

1. Hardness for spherical error:
   - Avoid $n^{1/4}$ degradation in the α_i rates?
   - Support unbounded samples?
2. Hardness for smaller error with fewer samples? (Extend [MP'13]?)

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Open Questions

1. Hardness for spherical error:
   - Avoid $n^{1/4}$ degradation in the α_i rates?
   - Support unbounded samples?
2. Hardness for smaller error with fewer samples? (Extend [MP'13]?)
3. Nontrivially relate Ideal-SIVP or Ring-LWE for different rings?

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Open Questions

1. Hardness for spherical error:
   - Avoid $n^{1/4}$ degradation in the α_i rates?
   - Support unbounded samples?
2. Hardness for smaller error with fewer samples? (Extend [MP'13]?)
3. Nontrivially relate Ideal-SIVP or Ring-LWE for different rings?
4. Evidence for/against Ring-LWE $\leq$ Ideal-SIVP?

Final Thoughts and Open Problems

- decision-R-LWE$_{q,\alpha}$ is worst-case hard for *any* $R = \mathcal{O}_K$, modulus q
- decision-LWE$_{q,\alpha}$ is hard for any q; approx factor independent of q

Open Questions

1. Hardness for spherical error:
   - ⋆ Avoid $n^{1/4}$ degradation in the α_i rates?
   - ⋆ Support unbounded samples?
2. Hardness for smaller error with fewer samples? (Extend [MP'13]?)
3. Nontrivially relate Ideal-SIVP or Ring-LWE for different rings?
4. Evidence for/against Ring-LWE $\leq$ Ideal-SIVP?
5. Classical reduction matching params of quantum reductions?