Ideal Lattices and Ring-LWE: Overview and Open Problems

Chris Peikert
Georgia Institute of Technology

ICERM
23 April 2015

Agenda

1. Ring-LWE and its hardness from ideal lattices

2. Open questions

Selected bibliography:

LPR'10 V. Lyubashevsky, C. Peikert, O. Regev.
"On Ideal Lattices and Learning with Errors Over Rings," Eurocrypt'10 and JACM'13.

LPR'13 V. Lyubashevsky, C. Peikert, O. Regev.
"A Toolkit for Ring-LWE Cryptography," Eurocrypt'13.

A Brief, Selective History of Lattice Cryptography

1996 Ajtai's worst-case/average-case reduction, one-way function & public-key encryption (very inefficient)

A Brief, Selective History of Lattice Cryptography

1996 Ajtai's worst-case/average-case reduction, one-way function & public-key encryption (very inefficient)

1996 NTRU efficient ring-based encryption (heuristic security)

A Brief, Selective History of Lattice Cryptography

1996 Ajtai's worst-case/average-case reduction, one-way function & public-key encryption (very inefficient)

1996 NTRU efficient ring-based encryption (heuristic security)

2002 Micciancio's ring-based one-way function with worst-case hardness (no encryption)

A Brief, Selective History of Lattice Cryptography

1996 Ajtai's worst-case/average-case reduction, one-way function
& public-key encryption (very inefficient)

1996 NTRU efficient ring-based encryption (heuristic security)

2002 Micciancio's ring-based one-way function
with worst-case hardness (no encryption)

2005 Regev's LWE: encryption with worst-case hardness
(inefficient)

A Brief, Selective History of Lattice Cryptography

1996 Ajtai's worst-case/average-case reduction, one-way function & public-key encryption (very inefficient)

1996 NTRU efficient ring-based encryption (heuristic security)

2002 Micciancio's ring-based one-way function with worst-case hardness (no encryption)

2005 Regev's LWE: encryption with worst-case hardness (inefficient)

2008– Countless applications of LWE (still inefficient)

A Brief, Selective History of Lattice Cryptography

1996	Ajtai's worst-case/average-case reduction, one-way function & public-key encryption	(very inefficient)
1996	NTRU efficient ring-based encryption	(heuristic security)
2002	Micciancio's ring-based one-way function with worst-case hardness	(no encryption)
2005	Regev's LWE: encryption with worst-case hardness	(inefficient)
2008–	Countless applications of LWE	(still inefficient)
2010	Ring-LWE: efficient encryption, worst-case hardness	()

Learning With Errors [Regev'05]

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

Learning With Errors [Regev'05]

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

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

Learning With Errors [Regev'05]

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

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

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

$$\vdots$$

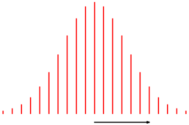

$\sqrt{n} \leq$ error $\ll q$

Learning With Errors [Regev'05]

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

$$\begin{pmatrix} \vdots \\ \mathbf{A} \\ \vdots \end{pmatrix} , \quad \begin{pmatrix} \vdots \\ \mathbf{b} \\ \vdots \end{pmatrix} = \mathbf{A}\mathbf{s} + \mathbf{e}$$

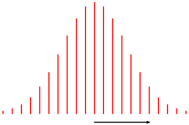

$\sqrt{n} \leq$ error $\ll q$

Learning With Errors [Regev'05]

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

$$\left(\begin{array}{c}\vdots\\ \mathbf{A}\\ \vdots\end{array}\right), \quad \left(\begin{array}{c}\vdots\\ \mathbf{b}\\ \vdots\end{array}\right) = \mathbf{A}\mathbf{s} + \mathbf{e}$$

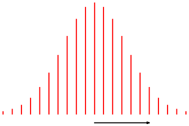

$\sqrt{n} \leq$ error $\ll q$

- **Decision:** distinguish $(\mathbf{A}\ ,\ \mathbf{b})$ from uniform $(\mathbf{A}\ ,\ \mathbf{b})$

Learning With Errors [Regev'05]

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

$$\left(\begin{array}{c} \vdots \\ \mathbf{A} \\ \vdots \end{array} \right) , \quad \left(\begin{array}{c} \vdots \\ \mathbf{b} \\ \vdots \end{array} \right) = \mathbf{A}\mathbf{s} + \mathbf{e}$$

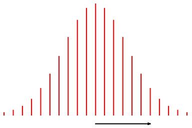

$\sqrt{n} \leq$ error $\ll q$

- **Decision:** distinguish $(\mathbf{A}\ ,\ \mathbf{b})$ from uniform $(\mathbf{A}\ ,\ \mathbf{b})$

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

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

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

Learning With Errors [Regev'05]

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

$$\begin{pmatrix} \vdots \\ \mathbf{A} \\ \vdots \end{pmatrix} , \begin{pmatrix} \vdots \\ \mathbf{b} \\ \vdots \end{pmatrix} = \mathbf{A}\mathbf{s} + \mathbf{e}$$

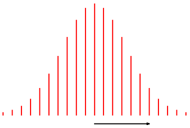

$\sqrt{n} \leq$ error $\ll q$

- **Decision:** distinguish $(\mathbf{A}\ ,\ \mathbf{b})$ from uniform $(\mathbf{A}\ ,\ \mathbf{b})$

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

worst case lattice problems $\underset{\uparrow}{\leq}$ search-LWE $\underset{\uparrow}{\leq}$ decision-LWE $\leq$ crypto

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

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

LWE is Versatile

What kinds of crypto can we do with LWE?

LWE is Versatile

What kinds of crypto can we do with LWE?

Public Key Encryption and Oblivious Transfer	[R'05,PVW'08]
Actively Secure PKE (w/o RO)	[PW'08,P'09,MP'12]

LWE is Versatile

What kinds of crypto can we do with LWE?

Public Key Encryption and Oblivious Transfer	[R'05,PVW'08]
Actively Secure PKE (w/o RO)	[PW'08,P'09,MP'12]
Identity-Based Encryption (in RO model)	[GPV'08]
Hierarchical ID-Based Encryption (w/o RO)	[CHKP'10,ABB'10]

LWE is Versatile

What kinds of crypto can we do with LWE?

Public Key Encryption and Oblivious Transfer	[R'05,PVW'08]
Actively Secure PKE (w/o RO)	[PW'08,P'09,MP'12]
Identity-Based Encryption (in RO model)	[GPV'08]
Hierarchical ID-Based Encryption (w/o RO)	[CHKP'10,ABB'10]
Leakage-Resilient Crypto	[AGV'09,DGKPV'10,GKPV'10,ADNSWW'10,...]
Fully Homomorphic Encryption	[BV'11,BGV'12,GSW'13,...]
Attribute-Based Encryption	[AFV'11,GVW'13,BGG+'14,...]
Symmetric-Key Primitives	[BPR'12,BMLR'13,BP'14,...]
Other Exotic Encryption	[ACPS'09,BHHI'10,OP'10,...]

the list goes on...

LWE is (Sort Of) Efficient

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

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

LWE is (Sort Of) Efficient

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

- 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.

LWE is (Sort Of) Efficient

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

- 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:

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

LWE is (Sort Of) Efficient

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

- 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:

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

- Can fix $\mathbf{A}$ for all users, but still $\geq n^2$ work to encrypt & decrypt an n-bit message

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,...]

LWE Over Rings, Over Simplified

- Let $\boxed{R = \mathbb{Z}[X]/(X^n + 1)}$ for n a power of two, and $\boxed{R_q = R/qR}$

LWE Over Rings, Over Simplified

- Let $\boxed{R = \mathbb{Z}[X]/(X^n + 1)}$ for n a power of two, and $\boxed{R_q = R/qR}$
  - Elements of R_q are deg $< n$ polynomials with mod-q coefficients
  - Operations in R_q are very efficient using FFT-like algorithms

LWE Over Rings, Over Simplified

- Let $\boxed{R = \mathbb{Z}[X]/(X^n + 1)}$ for n a power of two, and $\boxed{R_q = R/qR}$
  - ⋆ Elements of R_q are deg $< n$ polynomials with mod-q coefficients
  - ⋆ Operations in R_q are very efficient using FFT-like algorithms
- **Search**: find secret ring element $s(X) \in R_q$, given:

$$\begin{array}{lcl} a_1 \leftarrow R_q & , & b_1 = a_1 \cdot s + e_1 \in R_q \\ a_2 \leftarrow R_q & , & b_2 = a_2 \cdot s + e_2 \in R_q \\ a_3 \leftarrow R_q & , & b_3 = a_3 \cdot s + e_3 \in R_q \\ & \vdots & \end{array}$$

($e_i \in R$ are 'small')

LWE Over Rings, Over Simplified

- Let $\boxed{R = \mathbb{Z}[X]/(X^n + 1)}$ for n a power of two, and $\boxed{R_q = R/qR}$
  - ⋆ Elements of R_q are deg $< n$ polynomials with mod-q coefficients
  - ⋆ Operations in R_q are very efficient using FFT-like algorithms
- **Search**: find secret ring element $s(X) \in R_q$, given:

$$\begin{array}{lcl} a_1 \leftarrow R_q & , & b_1 = a_1 \cdot s + e_1 \in R_q \\ a_2 \leftarrow R_q & , & b_2 = a_2 \cdot s + e_2 \in R_q \\ a_3 \leftarrow R_q & , & b_3 = a_3 \cdot s + e_3 \in R_q \\ & \vdots & \end{array}$$

($e_i \in R$ are 'small')

Note: (a_i, b_i) are uniformly random subject to $\boxed{b_i - a_i \cdot s \approx 0}$

LWE Over Rings, Over Simplified

- Let $\boxed{R = \mathbb{Z}[X]/(X^n + 1)}$ for n a power of two, and $\boxed{R_q = R/qR}$
  - ⋆ Elements of R_q are deg $< n$ polynomials with mod-q coefficients
  - ⋆ Operations in R_q are very efficient using FFT-like algorithms
- **Search**: find secret ring element $s(X) \in R_q$, given:

$$\begin{array}{lcl} a_1 \leftarrow R_q & , & b_1 = a_1 \cdot s + e_1 \in R_q \\ a_2 \leftarrow R_q & , & b_2 = a_2 \cdot s + e_2 \in R_q \\ a_3 \leftarrow R_q & , & b_3 = a_3 \cdot s + e_3 \in R_q \\ & \vdots & \end{array}$$

($e_i \in R$ are ‘small’)

Note: (a_i, b_i) are uniformly random subject to $\boxed{b_i - a_i \cdot s \approx 0}$

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

(with noticeable advantage)

Hardness of Ring-LWE

- Two main theorems (reductions):

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

Hardness of Ring-LWE

- Two main theorems (reductions):

$$\begin{array}{c}\text{worst-case approx-SVP}\\ \text{on \textit{ideal} lattices in } R\end{array} \underset{\substack{\uparrow\\ \text{(quantum,}\\ \text{any } R=\mathcal{O}_K)}}{\leq} \text{search } R\text{-LWE} \underset{\substack{\uparrow\\ \text{(classical,}\\ \text{any cyclotomic } R)}}{\leq} \text{decision } R\text{-LWE}$$

1. If you can find s given $(a_i\ ,\ b_i)$, then you can find approximately shortest vectors in *any* ideal lattice in R (using a quantum algorithm).

Hardness of Ring-LWE

- Two main theorems (reductions):

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

1. If you can find s given $(a_i\ , b_i)$, then you can find approximately shortest vectors in *any* ideal lattice in R (using a quantum algorithm).
2. If you can distinguish $(a_i\ , b_i)$ from $(a_i\ , b_i)$, then you can find s.

Hardness of Ring-LWE

- Two main theorems (reductions):

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

1. If you can find s given $(a_i\ ,\ b_i)$, then you can find approximately shortest vectors in *any* ideal lattice in R (using a quantum algorithm).
2. If you can distinguish $(a_i\ ,\ b_i)$ from $(a_i\ ,\ b_i)$, then you can find s.

- Then:

decision R-LWE $\leq$ lots of crypto

Hardness of Ring-LWE

- Two main theorems (reductions):

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

1. If you can find s given $(a_i\ ,\ b_i)$, then you can find approximately shortest vectors in *any* ideal lattice in R (using a quantum algorithm).
2. If you can distinguish $(a_i\ ,\ b_i)$ from $(a_i\ ,\ b_i)$, then you can find s.

- Then:

decision R-LWE $\leq$ lots of crypto

★ If you can break the crypto, then you can distinguish $(a_i\ ,\ b_i)$ from $(a_i\ ,\ b_i)\ldots$

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

To get ideal lattices, embed R and its ideals into $\mathbb{R}^n$. How?

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

To get ideal lattices, embed R and its ideals into $\mathbb{R}^n$. How?

1. 'Obvious' answer: 'coefficient embedding'

$$a_0 + a_1 X + \cdots + a_{n-1} X^{n-1} \in R \quad \mapsto \quad (a_0, \ldots, a_{n-1}) \in \mathbb{Z}^n$$

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

To get ideal lattices, embed R and its ideals into $\mathbb{R}^n$. How?

1. ‘Obvious’ answer: ‘coefficient embedding’

$$a_0 + a_1 X + \cdots + a_{n-1} X^{n-1} \in R \quad \mapsto \quad (a_0, \ldots, a_{n-1}) \in \mathbb{Z}^n$$

$+$ is coordinate-wise, but analyzing $\cdot$ is cumbersome.

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

To get ideal lattices, embed R and its ideals into $\mathbb{C}^n$. How?

1. 'Obvious' answer: 'coefficient embedding'

$$a_0 + a_1 X + \cdots + a_{n-1} X^{n-1} \in R \quad \mapsto \quad (a_0, \ldots, a_{n-1}) \in \mathbb{Z}^n$$

$+$ is coordinate-wise, but analyzing $\cdot$ is cumbersome.

2. [Minkowski]: 'canonical embedding.' Let $\omega = \exp(\pi i/n) \in \mathbb{C}$, so roots of $X^n + 1$ are $\omega^1, \omega^3, \ldots, \omega^{2n-1}$. Embed:

$$a(X) \in R \quad \mapsto \quad \left(a(\omega^1)\,,\, a(\omega^3)\,, \ldots,\, a(\omega^{2n-1})\right) \in \mathbb{C}^n$$

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

To get ideal lattices, embed R and its ideals into $\mathbb{C}^n$. How?

1. 'Obvious' answer: 'coefficient embedding'

$$a_0 + a_1 X + \cdots + a_{n-1} X^{n-1} \in R \quad \mapsto \quad (a_0, \ldots, a_{n-1}) \in \mathbb{Z}^n$$

$+$ is coordinate-wise, but analyzing $\cdot$ is cumbersome.

2. [Minkowski]: 'canonical embedding.' Let $\omega = \exp(\pi i/n) \in \mathbb{C}$, so roots of $X^n + 1$ are $\omega^1, \omega^3, \ldots, \omega^{2n-1}$. Embed:

$$a(X) \in R \quad \mapsto \quad (a(\omega^1)\,,\, a(\omega^3)\,, \ldots,\, a(\omega^{2n-1})) \in \mathbb{C}^n$$

Both $+$ and $\cdot$ are coordinate-wise.

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^n + 1)$ for power-of-two n. (Or $R = \mathcal{O}_K$.)
- An ideal $\mathcal{I} \subseteq R$ is closed under $+$ and $-$, and under $\cdot$ with R.

To get ideal lattices, embed R and its ideals into $\mathbb{R}^n$. How?

1. 'Obvious' answer: 'coefficient embedding'

$$a_0 + a_1 X + \cdots + a_{n-1} X^{n-1} \in R \quad \mapsto \quad (a_0, \ldots, a_{n-1}) \in \mathbb{Z}^n$$

$+$ is coordinate-wise, but analyzing $\cdot$ is cumbersome.

2. [Minkowski]: 'canonical embedding.' Let $\omega = \exp(\pi i/n) \in \mathbb{C}$, so roots of $X^n + 1$ are $\omega^1, \omega^3, \ldots, \omega^{2n-1}$. Embed:

$$a(X) \in R \quad \mapsto \quad (a(\omega^1)\,,\, a(\omega^3)\,, \ldots,\, a(\omega^{2n-1})) \in \mathbb{C}^n$$

Both $+$ and $\cdot$ are coordinate-wise.

(NB: LWE error distribution is Gaussian in canonical embedding.)

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^2 + 1)$. Embeddings map $X \mapsto \pm i$.

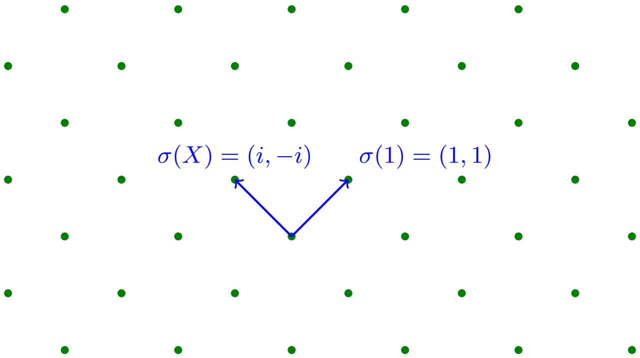

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^2+1)$. Embeddings map $X \mapsto \pm i$.
- $\mathcal{I} = \langle X-2, -3X+1 \rangle$ is an ideal in R.

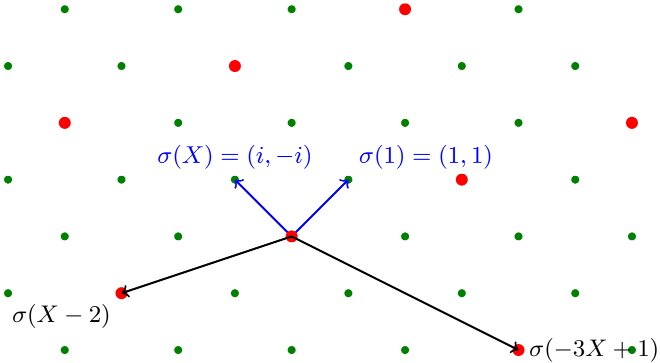

Ideal Lattices

- Say $R = \mathbb{Z}[X]/(X^2+1)$. Embeddings map $X \mapsto \pm i$.
- $\mathcal{I} = \langle X-2, -3X+1 \rangle$ is an ideal in R.

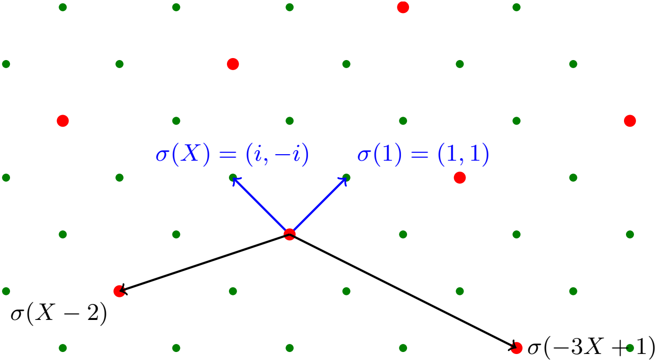

(Approximate) Shortest Vector Problem

- Given (an arbitrary basis of) an arbitrary ideal $\mathcal{I} \subseteq R$, find a nearly shortest nonzero $a \in \mathcal{I}$.

Hardness of Search Ring-LWE

Theorem 1

For any large enough q, solving search R-LWE
is as hard as quantumly solving poly(n)-approx SVP
in any (worst-case) ideal lattice in $R = \mathcal{O}_K$.

Hardness of Search Ring-LWE

Theorem 1

For any large enough q, solving search R-LWE
is as hard as quantumly solving poly(n)-approx SVP
in any (worst-case) ideal lattice in $R = \mathcal{O}_K$.

- Proof follows the template of [Regev'05] for LWE & arbitrary lattices.

  Quantum component used as 'black-box;' only classical part needs adaptation to the ring setting.

Hardness of Search Ring-LWE

Theorem 1

For any large enough q, solving search R-LWE
is as hard as quantumly solving poly(n)-approx SVP
in any (worst-case) ideal lattice in $R = \mathcal{O}_K$.

- Proof follows the template of [Regev'05] for LWE & arbitrary lattices.

  Quantum component used as 'black-box;' only classical part needs adaptation to the ring setting.

- Main technique: 'clearing ideals' while preserving R-module structure:

$$\begin{aligned} \mathcal{I}/q\mathcal{I} &\mapsto R/qR, \\ \mathcal{I}^\vee/q\mathcal{I}^\vee &\mapsto R^\vee/qR^\vee. \end{aligned}$$

  Uses Chinese remainder theorem and theory of duality for ideals.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision R-LWE in any cyclotomic $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$

(for any poly(n)-bounded prime $q = 1 \bmod m$)

is as hard as solving search R-LWE.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision R-LWE in any cyclotomic $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$

(for any poly(n)-bounded prime $q = 1 \bmod m$)

is as hard as solving search R-LWE.

Facts Used in the Proof

- $\mathbb{Z}_q^*$ has order $q - 1 = 0 \bmod m$, so has an element ω of order m.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision R-LWE in any cyclotomic $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$

(for any poly(n)-bounded prime $q = 1 \bmod m$)

is as hard as solving search R-LWE.

Facts Used in the Proof

- $\mathbb{Z}_q^*$ has order $q - 1 = 0 \bmod m$, so has an element ω of order m.
- Modulo q, $\Phi_m(X)$ has $n = \varphi(m)$ roots ω^j, for $j \in \mathbb{Z}_m^*$.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision R-LWE in any cyclotomic $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/\Phi_m(X)$

(for any poly(n)-bounded prime $q = 1 \bmod m$)

is as hard as solving search R-LWE.

Facts Used in the Proof

- $\mathbb{Z}_q^*$ has order $q - 1 = 0 \bmod m$, so has an element ω of order m.
- Modulo q, $\Phi_m(X)$ has $n = \varphi(m)$ roots ω^j, for $j \in \mathbb{Z}_m^*$.
- So there is a ring isomorphism $R_q \cong \mathbb{Z}_q^n$ given by

$$a(X) \in R_q \mapsto \quad \big(\, a(\omega^j) \,\big)_{j \in \mathbb{Z}_m^*} \in \mathbb{Z}_q^n.$$

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$
is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,, b \approx a \cdot s)$ from uniform $(a\,, b)$.

Goal: Find $s \in R_q$, given samples $(a\,, b \approx a \cdot s)$.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$
is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,, b \approx a \cdot s)$ from uniform $(a\,, b)$.

Goal: Find $s \in R_q$, given samples $(a\,, b \approx a \cdot s)$.

1. Equivalent to finding $s(\omega^j) \in \mathbb{Z}_q$ for all $j \in \mathbb{Z}_m^*$.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$
is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,, b \approx a \cdot s)$ from uniform $(a\,, b)$.

Goal: Find $s \in R_q$, given samples $(a\,, b \approx a \cdot s)$.

1. Equivalent to finding $s(\omega^j) \in \mathbb{Z}_q$ for all $j \in \mathbb{Z}_m^*$.
2. Hybrid argument: randomize one $b(\omega^j) \in \mathbb{Z}_q$; or two; or three; or . . .

   Then $\mathcal{O}$ must distinguish relative to some ω^{j^*}.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$
is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,, b \approx a \cdot s)$ from uniform $(a\,, b)$.

Goal: Find $s \in R_q$, given samples $(a\,, b \approx a \cdot s)$.

1. Equivalent to finding $s(\omega^j) \in \mathbb{Z}_q$ for all $j \in \mathbb{Z}_m^*$.
2. Hybrid argument: randomize one $b(\omega^j) \in \mathbb{Z}_q$; or two; or three; or . . .
   Then $\mathcal{O}$ must distinguish relative to some ω^{j^*}.
3. Using $\mathcal{O}$, guess-and-check to find $s(\omega^{j^*}) \in \mathbb{Z}_q$.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$ is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,, b \approx a \cdot s)$ from uniform $(a\,, b)$.

Goal: Find $s \in R_q$, given samples $(a\,, b \approx a \cdot s)$.

1. Equivalent to finding $s(\omega^j) \in \mathbb{Z}_q$ for all $j \in \mathbb{Z}_m^*$.
2. Hybrid argument: randomize one $b(\omega^j) \in \mathbb{Z}_q$; or two; or three; or ... Then $\mathcal{O}$ must distinguish relative to some ω^{j^*}.
3. Using $\mathcal{O}$, guess-and-check to find $s(\omega^{j^*}) \in \mathbb{Z}_q$.
4. How to find other $s(\omega^j)$? Couldn't $\mathcal{O}$ be useless at other roots?

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$
is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,, b \approx a \cdot s)$ from uniform $(a\,, b)$.

Goal: Find $s \in R_q$, given samples $(a\,, b \approx a \cdot s)$.

1. Equivalent to finding $s(\omega^j) \in \mathbb{Z}_q$ for all $j \in \mathbb{Z}_m^*$.
2. Hybrid argument: randomize one $b(\omega^j) \in \mathbb{Z}_q$; or two; or three; or . . .
   Then $\mathcal{O}$ must distinguish relative to some ω^{j^*}.
3. Using $\mathcal{O}$, guess-and-check to find $s(\omega^{j^*}) \in \mathbb{Z}_q$.
4. How to find other $s(\omega^j)$? Couldn't $\mathcal{O}$ be useless at other roots?
   $\omega \mapsto \omega^k$ $(k \in \mathbb{Z}_m^*)$ permutes roots of $\Phi_m(X)$, and preserves error.

Hardness of Decision Ring-LWE

Theorem 2

Solving decision Ring-LWE in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$
is as hard as solving search Ring-LWE.

Proof Sketch

Given: $\mathcal{O}$ distinguishes samples $(a\,,b \approx a \cdot s)$ from uniform $(a\,,b)$.

Goal: Find $s \in R_q$, given samples $(a\,,b \approx a \cdot s)$.

1. Equivalent to finding $s(\omega^j) \in \mathbb{Z}_q$ for all $j \in \mathbb{Z}_m^*$.
2. Hybrid argument: randomize one $b(\omega^j) \in \mathbb{Z}_q$; or two; or three; or ...
   Then $\mathcal{O}$ must distinguish relative to some ω^{j^*}.
3. Using $\mathcal{O}$, guess-and-check to find $s(\omega^{j^*}) \in \mathbb{Z}_q$.
4. How to find other $s(\omega^j)$? Couldn't $\mathcal{O}$ be useless at other roots?
   $\omega \mapsto \omega^k$ $(k \in \mathbb{Z}_m^*)$ permutes roots of $\Phi_m(X)$, and preserves error.
   So send each ω^j to ω^{j^*}, and use $\mathcal{O}$ to find $s(\omega^j)$.

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

   - ★ [P'09] reduces GapSVP (i.e., estimate $\lambda_1(\mathcal{L})$) on general lattices to plain-LWE, classically.

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

   - ⋆ [P'09] reduces GapSVP (i.e., estimate $\lambda_1(\mathcal{L})$) on general lattices to plain-LWE, classically.
   - ⋆ But estimating $\lambda_1(\mathcal{L})$ is trivially easy on ideal lattices!
     Finding short vectors is what appears hard.

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

   - ★ [P'09] reduces GapSVP (i.e., estimate $\lambda_1(\mathcal{L})$) on general lattices to plain-LWE, classically.
   - ★ But estimating $\lambda_1(\mathcal{L})$ is trivially easy on ideal lattices!
     Finding short vectors is what appears hard.

2. Search- and decision-R-LWE are equivalent in cyclotomic R.
   Does this hold in other kinds of rings?

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

   - ⋆ [P'09] reduces GapSVP (i.e., estimate $\lambda_1(\mathcal{L})$) on general lattices to plain-LWE, classically.
   - ⋆ But estimating $\lambda_1(\mathcal{L})$ is trivially easy on ideal lattices!
     Finding short vectors is what appears hard.

2. Search- and decision-R-LWE are equivalent in cyclotomic R.
   Does this hold in other kinds of rings?

   - ⋆ Yes, for any Galois number field (identical proof).

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

   - ⋆ [P'09] reduces GapSVP (i.e., estimate $\lambda_1(\mathcal{L})$) on general lattices to plain-LWE, classically.

   - ⋆ But estimating $\lambda_1(\mathcal{L})$ is trivially easy on ideal lattices!
     Finding short vectors is what appears hard.

2. Search- and decision-R-LWE are equivalent in cyclotomic R.
   Does this hold in other kinds of rings?

   - ⋆ Yes, for any Galois number field (identical proof).

   - ⋆ ~~Probably not, for carefully constructed rings S, moduli q, and errors!~~
     ~~Decision-S-LWE easily broken, but search unaffected. [EHL'14,ELOS'15]~~
     Update 8/2016: both search and decision are broken [CIV'16,P'16]

Open Problems: Reductions

1. Search-R-LWE is quantumly at least as hard as approx-R-SVP.
   Is there a classical reduction?

   - ⋆ [P'09] reduces GapSVP (i.e., estimate $\lambda_1(\mathcal{L})$) on general lattices to plain-LWE, classically.
   - ⋆ But estimating $\lambda_1(\mathcal{L})$ is trivially easy on ideal lattices!
     Finding short vectors is what appears hard.

2. Search- and decision-R-LWE are equivalent in cyclotomic R.
   Does this hold in other kinds of rings?

   - ⋆ Yes, for any Galois number field (identical proof).
   - ⋆ ~~Probably not, for carefully constructed rings S, moduli q, and errors!~~

     ~~Decision-S-LWE easily broken, but search unaffected. [EHL'14,ELOS'15]~~
     Update 8/2016: both search and decision are broken [CIV'16,P'16]

     "cyclotomic fields, used for Ring-LWE, are uniquely protected against the attacks presented in this paper"

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction? Can we solve R-LWE using an oracle for approx-R-SVP?

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction?
   Can we solve R-LWE using an oracle for approx-R-SVP?

   - ⋆ R-LWE samples $(a_i, b_i)_{i=1,\ldots,\ell}$ don't readily translate to ideals in R.

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction?
   Can we solve R-LWE using an oracle for approx-R-SVP?

   - ⋆ R-LWE samples $(a_i, b_i)_{i=1,\ldots,\ell}$ don't readily translate to ideals in R.
   - ⋆ They do yield a BDD instance on an R-module lattice:

$$\mathcal{L} = \left\{ (v_i) \ : \ v_i = a_i \cdot z \pmod{qR} \right\} \subseteq R^\ell$$

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction? Can we solve R-LWE using an oracle for approx-R-SVP?

   ⋆ R-LWE samples $(a_i, b_i)_{i=1,\ldots,\ell}$ don't readily translate to ideals in R.

   ⋆ They do yield a BDD instance on an R-module lattice:

$$\mathcal{L} = \big\{(v_i) \;:\; v_i = a_i \cdot z \pmod{qR}\big\} \subseteq R^\ell$$

2. How hard/easy is approx-R-SVP, anyway? (In cyclotomics etc.)

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction?
   Can we solve R-LWE using an oracle for approx-R-SVP?

   - ⋆ R-LWE samples $(a_i, b_i)_{i=1,\ldots,\ell}$ don't readily translate to ideals in R.
   - ⋆ They do yield a BDD instance on an R-module lattice:

   $$\mathcal{L} = \left\{(v_i) \ : \ v_i = a_i \cdot z \pmod{qR}\right\} \subseteq R^\ell$$

2. How hard/easy is approx-R-SVP, anyway? (In cyclotomics etc.)

   - ⋆ Despite abundant ring structure (e.g., subfields, Galois), no substantial improvement over attacks on general lattices.

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction? Can we solve R-LWE using an oracle for approx-R-SVP?

   - ⋆ R-LWE samples $(a_i, b_i)_{i=1,\ldots,\ell}$ don't readily translate to ideals in R.
   - ⋆ They do yield a BDD instance on an R-module lattice:

   $$\mathcal{L} = \left\{(v_i) \ : \ v_i = a_i \cdot z \pmod{qR}\right\} \subseteq R^\ell$$

2. How hard/easy is approx-R-SVP, anyway? (In cyclotomics etc.)

   - ⋆ Despite abundant ring structure (e.g., subfields, Galois), no substantial improvement over attacks on general lattices.
   - ⋆ Next up: attacks on a specialized variant: given a principal ideal $\mathcal{I}$ guaranteed to have an "unusually short" generator, find it.

Open Problems: Attacks

1. We know approx-R-SVP $\leq$ R-LWE (quantumly). Other direction? Can we solve R-LWE using an oracle for approx-R-SVP?

   - ⋆ R-LWE samples $(a_i, b_i)_{i=1,\ldots,\ell}$ don't readily translate to ideals in R.
   - ⋆ They do yield a BDD instance on an R-module lattice:

   $$\mathcal{L} = \big\{(v_i) \ : \ v_i = a_i \cdot z \pmod{qR}\big\} \subseteq R^\ell$$

2. How hard/easy is approx-R-SVP, anyway? (In cyclotomics etc.)

   - ⋆ Despite abundant ring structure (e.g., subfields, Galois), no substantial improvement over attacks on general lattices.
   - ⋆ Next up: attacks on a specialized variant: given a principal ideal $\mathcal{I}$ guaranteed to have an "unusually short" generator, find it.
   - ⋆ These conditions are extremely rare for general ideals, so (worst-case) approx-R-SVP is unaffected.