FrodoKEM

practical quantum-secure key encapsulation from generic lattices

Erdem Alkim Joppe W. Bos Léo Ducas Patrick Longa

Ilya Mironov Michael Naehrig Valeria Nikolaenko

Chris Peikert Ananth Raghunathan Douglas Stebila

FrodoKEM

FrodoKEM's security derives from plain *Learning With Errors*
on algebraically unstructured lattices,
parameterized cautiously to avoid known risk categories,
and to conform to a worst-case/average-case reduction.

FrodoKEM

FrodoKEM's security derives from plain *Learning With Errors*

on algebraically unstructured lattices,

parameterized cautiously to avoid known risk categories,

and to conform to a worst-case/average-case reduction.

FrodoKEM

FrodoKEM's security derives from plain *Learning With Errors*

on algebraically unstructured lattices,

parameterized cautiously to avoid known risk categories,

and to conform to a worst-case/average-case reduction.

FrodoKEM

FrodoKEM's security derives from plain *Learning With Errors*
on algebraically unstructured lattices,
parameterized cautiously to avoid known risk categories,
and to conform to a worst-case/average-case reduction.

FrodoPKE (IND-CPA) $\xrightarrow[\text{(generic transform)}]{\text{[FujisakiOkamoto'99,HHK'17]}}$ FrodoKEM (IND-CCA)

FrodoKEM

FrodoKEM's security derives from plain *Learning With Errors*
on algebraically unstructured lattices,
parameterized cautiously to avoid known risk categories,
and to conform to a worst-case/average-case reduction.

FrodoPKE (IND-CPA) $\xrightarrow[\text{(generic transform)}]{\text{[FujisakiOkamoto'99,HHK'17]}}$ FrodoKEM (IND-CCA)

Concrete Instantiations

1. FrodoKEM-640: targets Level 1 security ($\geq$ AES-128).
2. FrodoKEM-976: targets Level 3 security ($\geq$ AES-192).
3. Other parameterizations are easy, by changing compile-time constants.

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions:

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions: breaking random inputs $\implies$ solving famous problems on any lattice.

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions: breaking random inputs $\Longrightarrow$ solving famous problems on any lattice.

  > *"[This] assures us that attacks on the cryptographic construction are likely to be effective only for small choices of parameters and not asymptotically. In other words . . . there are no fundamental flaws in the design of our cryptographic construction."* *[MicciancioRegev'09]*

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions: breaking random inputs $\implies$ solving famous problems on any lattice.

  > *"[This] assures us that attacks on the cryptographic construction are likely to be effective only for small choices of parameters and not asymptotically. In other words ... there are no fundamental flaws in the design of our cryptographic construction."* *[MicciancioRegev'09]*

- LWE has been heavily used and cryptanalyzed by countless works.

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions: breaking random inputs $\Longrightarrow$ solving famous problems on any lattice.

  > *"[This] assures us that attacks on the cryptographic construction are likely to be effective only for small choices of parameters and not asymptotically. In other words . . . there are no fundamental flaws in the design of our cryptographic construction."* *[MicciancioRegev'09]*

- LWE has been heavily used and cryptanalyzed by countless works.

Public-Key Encryption/Key Exchange

- Many schemes with tight (CPA-)security reductions from LWE:

  [Regev'05,PVW'08,GPV'08,P'09,LP'11,. . .]

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions:
  breaking random inputs $\Longrightarrow$ solving famous problems on any lattice.

  > *"[This] assures us that attacks on the cryptographic construction are likely to be effective only for small choices of parameters and not asymptotically. In other words ... there are no fundamental flaws in the design of our cryptographic construction."* [MicciancioRegev'09]

- LWE has been heavily used and cryptanalyzed by countless works.

Public-Key Encryption/Key Exchange

- Many schemes with tight (CPA-)security reductions from LWE:
  [Regev'05,PVW'08,GPV'08,P'09,LP'11,...]

- FrodoCCS [BCDMNNRS'16] instantiated and implemented [LP'11], using pseudorandom public matrix $\mathbf{A}$ to reduce public key size.

Pedigree

Learning With Errors (LWE) [Regev'05]

- Lineage of [Ajtai'96,AjtaiDwork'97]: worst-case/average-case reductions: breaking random inputs $\implies$ solving famous problems on any lattice.

  "[This] assures us that attacks on the cryptographic construction are likely to be effective only for small choices of parameters and not asymptotically. In other words ... there are no fundamental flaws in the design of our cryptographic construction." *[MicciancioRegev'09]*

- LWE has been heavily used and cryptanalyzed by countless works.

Public-Key Encryption/Key Exchange

- Many schemes with tight (CPA-)security reductions from LWE: [Regev'05,PVW'08,GPV'08,P'09,LP'11,. . .]
- FrodoCCS [BCDMNNRS'16] instantiated and implemented [LP'11], using pseudorandom public matrix $\mathbf{A}$ to reduce public key size.
- FrodoPKE [this work]: wider error distributions, new parameters, . . .

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on ‘small’ integers.

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

  Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

$$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}] \stackrel{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

  Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

  $$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}] \overset{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

Bounded-distance decoding on a random 'q-ary' lattice defined by $\mathbf{A}$:

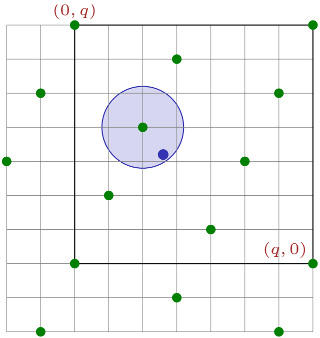

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

  Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

  $$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}] \stackrel{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

$$\mathbf{S} \leftarrow \chi^{k\times n} \qquad \xrightarrow[(\mathbf{A} = \mathsf{expand}(\mathsf{seed}_{\mathbf{A}}) \in \mathbb{Z}_q^{n\times n})]{pk = \mathsf{seed}_{\mathbf{A}}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}}$$

(Images courtesy xkcd.org)

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

  Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

$$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}] \stackrel{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

$\mathbf{S} \leftarrow \chi^{k \times n}$

$$\xrightarrow[(\mathbf{A} = \mathsf{expand}(\mathsf{seed}_{\mathbf{A}}) \in \mathbb{Z}_q^{n \times n})]{pk = \mathsf{seed}_{\mathbf{A}}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}}$$

$\mathbf{M} \in \{0,1\}^{k \times \ell}$

(Images courtesy xkcd.org)

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

$$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}] \stackrel{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

$\mathbf{S} \leftarrow \chi^{k \times n}$

$$\xrightarrow[(\mathbf{A} = \mathsf{expand}(\mathsf{seed}_{\mathbf{A}}) \in \mathbb{Z}_q^{n \times n})]{pk = \mathsf{seed}_{\mathbf{A}}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}}$$

$$\xleftarrow[\mathbf{C}' \approx \mathbf{B}\mathbf{R} + \frac{q}{2} \cdot \mathbf{M}]{\mathbf{C}\ \approx \mathbf{A}\mathbf{R}}$$

$\mathbf{M} \in \{0,1\}^{k \times \ell}$

(Images courtesy xkcd.org)

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

  Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

  $$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}] \stackrel{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

$\mathbf{S} \leftarrow \chi^{k \times n}$

$$\xrightarrow[(\mathbf{A} = \mathsf{expand}(\mathsf{seed}_{\mathbf{A}}) \in \mathbb{Z}_q^{n \times n})]{pk = \mathsf{seed}_{\mathbf{A}}\ ,\ \mathbf{B} \approx \mathbf{S}\mathbf{A}}$$

$$\xleftarrow[\mathbf{C}' \approx \mathbf{B}\mathbf{R} + \frac{q}{2} \cdot \mathbf{M}]{\mathbf{C} \approx \mathbf{A}\mathbf{R}}$$

$\mathbf{M} \in \{0,1\}^{k \times \ell}$

$\mathbf{C}' - \mathbf{S}\mathbf{C} \approx \frac{q}{2} \cdot \mathbf{M}$

(Images courtesy xkcd.org)

LWE and FrodoPKE

Learning With Errors

- Dimension n, modulus q, error distribution χ on 'small' integers.

  Assumption: for uniformly random matrix $\mathbf{A}$ over $\mathbb{Z}_q$ and $\mathbf{S}$ from χ,

$$[\mathbf{A}\ ,\ \mathbf{B} \approx \mathbf{SA}] \stackrel{c}{\equiv} \text{uniform over } \mathbb{Z}_q.$$

$$\mathbf{S} \leftarrow \chi^{k\times n}$$

$$\xrightarrow[(\mathbf{A} = \mathsf{expand}(\mathsf{seed}_{\mathbf{A}}) \in \mathbb{Z}_q^{n\times n})]{pk = \mathsf{seed}_{\mathbf{A}}\ ,\ \mathbf{B} \approx \mathbf{SA}}$$

$$\xleftarrow[\mathbf{C}' \approx \mathbf{BR} + \frac{q}{2}\cdot\mathbf{M}]{\mathbf{C} \approx \mathbf{AR}}$$

$$\mathbf{M} \in \{0,1\}^{k\times\ell}$$

$$\mathbf{C}' - \mathbf{SC} \approx \frac{q}{2}\cdot\mathbf{M}$$

$$(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{C}') \stackrel{c}{\equiv} \text{unif}$$

(Images courtesy xkcd.org)

Distinctive Features of FrodoPKE/KEM

1. Generic, algebraically unstructured lattices: plain LWE.

2. 'Semi-wide' errors conforming to a worst-case/average-case reduction from a previously studied lattice problem: BDD with DGS.

3. Simple design and constant-time implementation:
   - power-of-2 modulus q for cheap & easy modular arithmetic
   - straightforward error sampling
   - no 'reconciliation' or error-correcting codes for removing noise
   - x64 implementation: 256 lines of plain C code (+ preexisting symmetric primitives)

Distinctive Features of FrodoPKE/KEM

1. Generic, algebraically unstructured lattices: plain LWE.

2. 'Semi-wide' errors conforming to a worst-case/average-case reduction from a previously studied lattice problem: BDD with DGS.

3. Simple design and constant-time implementation:
   - power-of-2 modulus q for cheap & easy modular arithmetic
   - straightforward error sampling
   - no 'reconciliation' or error-correcting codes for removing noise
   - x64 implementation: 256 lines of plain C code (+ preexisting symmetric primitives)

Distinctive Features of FrodoPKE/KEM

1. Generic, algebraically unstructured lattices: plain LWE.
2. 'Semi-wide' errors conforming to a worst-case/average-case reduction from a previously studied lattice problem: BDD with DGS.
3. Simple design and constant-time implementation:
   - ⋆ power-of-2 modulus q for cheap & easy modular arithmetic
   - ⋆ straightforward error sampling
   - ⋆ no 'reconciliation' or error-correcting codes for removing noise
   - ⋆ x64 implementation: 256 lines of plain C code (+ preexisting symmetric primitives)

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow$ n short vectors, speeds up lattice attacks [KF'17].

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow n$ short vectors, speeds up lattice attacks [KF'17].

   (Doesn't apply to Ring/Module-LWE.)

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow$ n short vectors, speeds up lattice attacks [KF'17].
   (Doesn't apply to Ring/Module-LWE.)
2. $2^{\tilde{O}(\sqrt{n})}$-approx-SVP in qpoly-time for ideal lattices in cyclotomics [CDPR'16,CDW'17].

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow n$ short vectors, speeds up lattice attacks [KF'17].

   (Doesn't apply to Ring/Module-LWE.)

2. $2^{\tilde{O}(\sqrt{n})}$-approx-SVP in qpoly-time for ideal lattices in cyclotomics [CDPR'16,CDW'17].

   (Doesn't apply to NTRU or R/M-LWE, nor to PKE approx factors.)

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow n$ short vectors, speeds up lattice attacks [KF'17].

   (Doesn't apply to Ring/Module-LWE.)

2. $2^{\tilde{O}(\sqrt{n})}$-approx-SVP in qpoly-time for ideal lattices in cyclotomics [CDPR'16,CDW'17].

   (Doesn't apply to NTRU or R/M-LWE, nor to PKE approx factors.)

$\Longrightarrow$ May be gaps in hardness between structured and unstructured lattices.

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow n$ short vectors, speeds up lattice attacks [KF'17].

   (Doesn't apply to Ring/Module-LWE.)

2. $2^{\tilde{O}(\sqrt{n})}$-approx-SVP in qpoly-time for ideal lattices in cyclotomics [CDPR'16,CDW'17].

   (Doesn't apply to NTRU or R/M-LWE, nor to PKE approx factors.)

$\Longrightarrow$ May be gaps in hardness between structured and unstructured lattices.

Our Foundation: Plain LWE on Unstructured Lattices

- LWE is bounded-distance decoding on a lattice defined by the uniformly random, unstructured matrix $\mathbf{A}$.

Unstructured Lattices

Risk Category 1: Geometric & Algebraic Structure

1. NTRU structure $\Rightarrow$ n short vectors, speeds up lattice attacks [KF'17].

   (Doesn't apply to Ring/Module-LWE.)

2. $2^{\tilde{O}(\sqrt{n})}$-approx-SVP in qpoly-time for ideal lattices in cyclotomics [CDPR'16,CDW'17].

   (Doesn't apply to NTRU or R/M-LWE, nor to PKE approx factors.)

$\Longrightarrow$ May be gaps in hardness between structured and unstructured lattices.

Our Foundation: Plain LWE on Unstructured Lattices

- LWE is bounded-distance decoding on a lattice defined by the uniformly random, unstructured matrix $\mathbf{A}$.
- No algebraic or 'planted' geometric structure in the lattice.

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\implies$ smaller parameters q, n $\implies$ better efficiency.

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\Longrightarrow$ smaller parameters q, n $\Longrightarrow$ better efficiency.
- But how narrow can the error distribution *safely* be?

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\Longrightarrow$ smaller parameters q, n $\Longrightarrow$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14]

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\implies$ smaller parameters $q, n \implies$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14] given large-poly(n)-many samples. (PKEs don't reveal this many!)

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\implies$ smaller parameters q, n $\implies$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14] given large-poly(n)-many samples. (PKEs don't reveal this many!)
2. Worst-case-hardness theorems need Gaussian error of $\sigma > \sqrt{n}/(2\pi)$.

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\Longrightarrow$ smaller parameters q, n $\Longrightarrow$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14] given large-poly(n)-many samples. (PKEs don't reveal this many!)
2. Worst-case-hardness theorems need Gaussian error of $\sigma > \sqrt{n}/(2\pi)$. Or narrower error, but only for few LWE samples. (PKEs reveal more!)

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\Longrightarrow$ smaller parameters q, n $\Longrightarrow$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14] given large-poly(n)-many samples. (PKEs don't reveal this many!)
2. Worst-case-hardness theorems need Gaussian error of $\sigma > \sqrt{n}/(2\pi)$. Or narrower error, but only for few LWE samples. (PKEs reveal more!)

$\Longrightarrow$ Sizeable gap between known-vulnerable and worst-case-hard params.

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\implies$ smaller parameters q, n $\implies$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14] given large-poly(n)-many samples. (PKEs don't reveal this many!)
2. Worst-case-hardness theorems need Gaussian error of $\sigma > \sqrt{n}/(2\pi)$. Or narrower error, but only for few LWE samples. (PKEs reveal more!)

$\implies$ Sizeable gap between known-vulnerable and worst-case-hard params.

New Worst-Case Hardness

- A latent reduction from [R'05,PRS'17] works for our $\sigma \geq 2.3 \approx \eta(\mathbb{Z})$.

Semi-Wide Errors

Choosing an Error Distribution

- Narrower errors $\implies$ smaller parameters q, n $\implies$ better efficiency.
- But how narrow can the error distribution *safely* be?

Risk Category 2: Narrow Errors

1. LWE with $O(1)$-bounded error is poly(n)-time solvable [AG'11,ACFP'14] given large-poly(n)-many samples. (PKEs don't reveal this many!)
2. Worst-case-hardness theorems need Gaussian error of $\sigma > \sqrt{n}/(2\pi)$. Or narrower error, but only for few LWE samples. (PKEs reveal more!)

$\implies$ Sizeable gap between known-vulnerable and worst-case-hard params.

New Worst-Case Hardness

- A latent reduction from [R'05,PRS'17] works for our $\sigma \geq 2.3 \approx \eta(\mathbb{Z})$.
- Works for a bounded poly(n) number of LWE samples: covers PKEs!

New Worst-Case Hardness

Worst-Case Problem: BDD with DGS [AR'04,R'05,LLM'06,DRS'14]

- Given N samples from discrete Gaussian $D_{\mathcal{L}^*}$, decode $\mathcal{L}$ to distance d.
- State of the art is limited to distance $d < \sqrt{\ln(N)/(2\pi)}$.

New Worst-Case Hardness

Worst-Case Problem: BDD with DGS [AR'04,R'05,LLM'06,DRS'14]

- Given N samples from discrete Gaussian $D_{\mathcal{L}^*}$, decode $\mathcal{L}$ to distance d.
- State of the art is limited to distance $d < \sqrt{\ln(N)/(2\pi)}$.

Theorem (extracted from [R'05,PRS'17])

Solving LWE for Gaussian error $\sigma \geq \eta(\mathbb{Z})$ with $m = \text{poly}(n)$ samples

$$\Downarrow$$

solving BDD at distance $d = \sigma\sqrt{2\pi}$ with $N = m \cdot \text{poly}(n)$ DGS samples.

New Worst-Case Hardness

Worst-Case Problem: BDD with DGS [AR'04,R'05,LLM'06,DRS'14]

- Given N samples from discrete Gaussian $D_{\mathcal{L}^*}$, decode $\mathcal{L}$ to distance d.
- State of the art is limited to distance $d < \sqrt{\ln(N)/(2\pi)}$.

Theorem (extracted from [R'05,PRS'17])

Solving LWE for Gaussian error $\sigma \geq \eta(\mathbb{Z})$ with $m = \mathrm{poly}(n)$ samples

$$\Downarrow$$

solving BDD at distance $d = \sigma\sqrt{2\pi}$ with $N = m \cdot \mathrm{poly}(n)$ DGS samples.

Interpretation

- Theoretical support & more confidence for semi-wide Gaussian error with limited number of samples.
- Reduction is non-tight; for concrete security we use cryptanalysis.

  (Tightening the time & sample overhead is a good research direction.)

Concrete Parameters

- Use 'core-SVP' methodology [ADPS'16] to lower-bound the *first-order exponential time* (and space) of SVP in appropriate dimension.

Concrete Parameters

- Use 'core-SVP' methodology [ADPS'16] to lower-bound the *first-order exponential time* (and space) of SVP in appropriate dimension.

  This significantly underestimates the cost of known attacks, but it is prudent to expect better lower-order terms with further research.

Concrete Parameters

- Use 'core-SVP' methodology [ADPS'16] to lower-bound the *first-order exponential time* (and space) of SVP in appropriate dimension.

  This significantly underestimates the cost of known attacks, but it is prudent to expect better lower-order terms with further research.

	n	q	σ	Bits of Security	
				C $\geq$	Q $\geq$
FrodoKEM-640	640	2^{15}	2.75	143	103
FrodoKEM-976	976	2^{16}	2.3	209	150

Performance

- Speed (in kilocycles, 3.4GHz Intel Core i7-6700 Skylake, AES-NI):

	KeyGen	**Encaps**	**Decaps**
FrodoKEM-640	1,287	1,810	1,811
FrodoKEM-976	2,715	3,572	3,588

Performance

- Speed (in kilocycles, 3.4GHz Intel Core i7-6700 Skylake, AES-NI):

	KeyGen	**Encaps**	**Decaps**
FrodoKEM-640	1,287	1,810	1,811
FrodoKEM-976	2,715	3,572	3,588

- Sizes (in bytes):

	secret key	**public key**	**ciphertext**
FrodoKEM-640	10,256	9,616	9,736
FrodoKEM-976	15,640	15,632	15,768

Parting Thought

FrodoKEM's security derives from plain *Learning With Errors*
on algebraically unstructured lattices,
parameterized cautiously to avoid known risk categories,
and to conform to a worst-case/average-case reduction.

`https://FrodoKEM.org`

Thanks!