Concrete Security Estimates and Parameter Selection for LWE-based schemes

Fernando Virdia

Information Security Group,
Royal Holloway, University of London,
United Kingdom

November 25, 2020

In this presentation I'll give an overview of the status of concrete cryptanalysis of LWE-based constructions.

In this presentation I'll give an overview of the status of concrete cryptanalysis of LWE-based constructions.

Solving Learning With Errors

- Effective strategies for LWE
- Costing lattice reduction

In this presentation I'll give an overview of the status of concrete cryptanalysis of LWE-based constructions.

Solving Learning With Errors

- Effective strategies for LWE
- Costing lattice reduction

Other considerations

- Key-reuse/CCA attacks
- Scripts for estimating attacks

In this presentation I'll give an overview of the status of concrete cryptanalysis of LWE-based constructions.

Solving Learning With Errors

- Effective strategies for LWE
- Costing lattice reduction

Other considerations

- Key-reuse/CCA attacks
- Scripts for estimating attacks

Throughout, references in [blue] are PROMETHEUS papers!

Learning With Errors

Given $(A, \vec{b} \equiv A\vec{s} + \vec{e} \bmod q)$ where $A \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, $\vec{s} \xleftarrow{\chi_s} \mathbb{Z}_q^n$, $\vec{e} \xleftarrow{\chi_e} \mathbb{Z}_q^m$, q power of prime, distinguish them from uniform (Decision LWE) or recover $\vec{s}$ (Search LWE).

Learning With Errors

Given $(A, \vec{b} \equiv A\vec{s} + \vec{e} \bmod q)$ where $A \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, $\vec{s} \xleftarrow{\chi_s} \mathbb{Z}_q^n$, $\vec{e} \xleftarrow{\chi_e} \mathbb{Z}_q^m$, q power of prime, distinguish them from uniform (Decision LWE) or recover $\vec{s}$ (Search LWE).

Lots of variants:

- replace $\mathbb{Z}_q^n$ with a polynomial ring $R = \mathbb{Z}_q/(f)$, or with an R-module.
  $\rightarrow$ No significant speedups against popular choices of f.

Learning With Errors

Given $(A, \vec{b} \equiv A\vec{s} + \vec{e} \bmod q)$ where $A \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$, $\vec{s} \xleftarrow{\chi_s} \mathbb{Z}_q^n$, $\vec{e} \xleftarrow{\chi_e} \mathbb{Z}_q^m$, q power of prime, distinguish them from uniform (Decision LWE) or recover $\vec{s}$ (Search LWE).

Lots of variants:

- replace $\mathbb{Z}_q^n$ with a polynomial ring $R = \mathbb{Z}_q/(f)$, or with an R-module.
  $\rightarrow$ No significant speedups against popular choices of f.
- χ_s uniform vs narrowly distributed around 0, with or without low Hamming wt.
  $\rightarrow$ χ_s uniform not intrinsically safer.
  $\rightarrow$ low Hamming wt. χ_s susceptible to "hybrid" combinatorial-lattice strategies.

Learning With Errors

Given $(A, \vec{b} \equiv A\vec{s} + \vec{e} \bmod q)$ where $A \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\vec{s} \xleftarrow{\chi_s} \mathbb{Z}_q^n$, $\vec{e} \xleftarrow{\chi_e} \mathbb{Z}_q^m$, q power of prime, distinguish them from uniform (Decision LWE) or recover $\vec{s}$ (Search LWE).

Lots of variants:

- replace $\mathbb{Z}_q^n$ with a polynomial ring $R = \mathbb{Z}_q/(f)$, or with an R-module.
  $\rightarrow$ No significant speedups against popular choices of f.
- χ_s uniform vs narrowly distributed around 0, with or without low Hamming wt.
  $\rightarrow$ χ_s uniform not intrinsically safer.
  $\rightarrow$ low Hamming wt. χ_s susceptible to "hybrid" combinatorial-lattice strategies.

Three attacks are currently considered "practical" for an attacker:
primal, dual, hybrid.

The primal attack

- Solves Search LWE by finding the unique closest vector to $\vec{b}$ in
$$\Lambda_q = \{A\vec{x} \bmod q \colon \vec{x} \in \mathbb{Z}^n\}.$$
- The state of the art strategy is to reduce a basis of Λ_q using something like BKZ-k, picking the smallest sufficient block size k according to [ADPS16].

The primal attack

- Solves Search LWE by finding the unique closest vector to $\vec{b}$ in
$$\Lambda_q = \{A\vec{x} \bmod q \colon \vec{x} \in \mathbb{Z}^n\}.$$

- The state of the art strategy is to reduce a basis of Λ_q using something like BKZ-k, picking the smallest sufficient block size k according to [ADPS16].

Recent results (incremental)

- [DSDGR20] studies how to exploit side-channel information.

- [P**V**20] investigates the success probability of smaller than expected block sizes.

No significant differences in estimated costing strategy since [ADPS16].

The dual attack

- Solves Decision LWE by finding a short vector in
$$\Lambda_q^\perp = \{(\vec{y}, \vec{x}) \in \mathbb{Z}^m \times \mathbb{Z}^n \colon \vec{y}A \equiv \vec{x} \bmod q\}.$$

- The state of the art strategy [Alb17] is to reduce a basis of $\Lambda_q^\perp$ using something like BKZ-k, picking the smallest sufficient block size k according to [Che13]. It accounts for narrow secret distributions.

- No significant improvements published recently (caveat; see next slide).

The hybrid attack

- Solves Search LWE by solving unique SVP with target $(\vec{s}, \vec{e}, 1)$.
- The state of the art has been stable since [HG07]:
  $\rightarrow$ Exploit a highly structured target vector by guessing some of its components.
  $\rightarrow$ Speed up the guessing step by using MITM techniques or quantum search.
- Currently does not affect small KEM parameters but maybe significant for FHE.

The hybrid attack

- Solves Search LWE by solving unique SVP with target $(\vec{s}, \vec{e}, 1)$.
- The state of the art has been stable since [HG07]:
  $\rightarrow$ Exploit a highly structured target vector by guessing some of its components.
  $\rightarrow$ Speed up the guessing step by using MITM techniques or quantum search.
- Currently does not affect small KEM parameters but maybe significant for FHE.

Recent results

[CHHS19, EJK20] introduce similar hybrid MITM techniques in the dual attack setting.

Hybrid attacks are where the impact of very sparse distributions may be most likely seen.

Costing lattice reduction

- Lattice reduction is a fundamental tool for solving LWE.
- Popular algorithms are block based: BKZ [SE91], Slide reduction [GN08], Progressive BKZ [AWHT16], Self-Dual BKZ [MW16].

Costing lattice reduction

- Lattice reduction is a fundamental tool for solving LWE.
- Popular algorithms are block based: BKZ [SE91], Slide reduction [GN08], Progressive BKZ [AWHT16], Self-Dual BKZ [MW16].

Block reduction cost

Given a block size k, one needs to solve SVP on k-dimensional sub-lattices.

- How many calls to the SVP oracle are needed?
- How much does each call to the oracle cost?

Number of SVP calls per lattice reduction

In [ACDDPP**V**W18], we surveyed the first round submissions to NIST.

- Some consider 16 tours of BKZ $\approx 8 \cdot rk(\Lambda)$ calls the oracle.
- Some follow the "core-SVP" [ADPS16] notion: lower bound by assuming 1 call.
- The Homomorphic Encryption Standard [ACC$^+$19] uses both!

Number of SVP calls per lattice reduction

In [ACDDPP**V**W18], we surveyed the first round submissions to NIST.

- Some consider 16 tours of BKZ $\approx 8 \cdot rk(\Lambda)$ calls the oracle.
- Some follow the "core-SVP" [ADPS16] notion: lower bound by assuming 1 call.
- The Homomorphic Encryption Standard [ACC^{+}19] uses both!

Personally, not clear that "16 tours" appropriately accounts for all the variants of block reduction.
$\rightarrow$ Core-SVP may be a safe compromise.
$\rightarrow$ The last Kyber spec discusses ways to push cryptanalysis "beyond core-SVP".

Practical SVP oracles: pruned enumeration

- $2^{\Theta(k \log k)}$ time, poly memory.
- Easily parallelisable.
- Quantum variant [ANS18]: asymptotic quadratic speedup (no concrete analysis).

Practical SVP oracles: pruned enumeration

- $2^{\Theta(k \log k)}$ time, poly memory.
- Easily parallelisable.
- Quantum variant [ANS18]: asymptotic quadratic speedup (no concrete analysis).

Recent results

- $2^{\frac{1}{8} k \log k + o(k)}$ heuristic lower bound achieved [ABF+20] (previously $2^{\frac{1}{2e} k \log k + o(k)}$).
- [ABLR20] further lowers the exponent $2^{\frac{1}{8} k \log k - 0.547k + 10.4}$ to $2^{\frac{1}{8} k \log k - 0.654k + 25.84}$ using approximate HSVP oracles.

While these results improve enumeration-based lattice reduction, they don't beat sieving runtimes.

Practical SVP oracles: lattice sieving

- $2^{\Theta(k)}$ time, $2^{\Theta(k)}$ memory.
- Record holder[1] for SVP, $k = 176$.
- "Concrete" analysis suggests quantum speedups minimal in practice [AGPS19].

[1] https://www.latticechallenge.org/svp-challenge/

Practical SVP oracles: lattice sieving

- $2^{\Theta(k)}$ time, $2^{\Theta(k)}$ memory.
- Record holder[1] for SVP, $k = 176$.
- "Concrete" analysis suggests quantum speedups minimal in practice [AGPS19].

Recent results

- [Duc18] proposes the "dimensions for free" technique.
- [ADH$^+$19] introduce and implement a generalisation of block-reduction in the case of sieving, using some of the multiple short vectors returned by the SVP oracle.

While these results improve experimental sieving runtimes, they don't beat asymptotically cheapest sieves.

[1] https://www.latticechallenge.org/svp-challenge/

Decryption failure attacks

On decryption, most lattice schemes return message + error $\approx$ message.

Decryption failure attacks

On decryption, most lattice schemes return message + error $\approx$ message.

- [Flu16] first key-reuse attack on LWE (on NTRU they date back to [JJ00]).
- It uses a Ciphertext Validity Oracle (CVO) to tell if a ciphertext correctly decrypts.

Decryption failure attacks

On decryption, most lattice schemes return message + error ≈ message.

- [Flu16] first key-reuse attack on LWE (on NTRU they date back to [JJ00]).
- It uses a Ciphertext Validity Oracle (CVO) to tell if a ciphertext correctly decrypts.

Naive analysis ("one-shot")

- Pick one key pair, compute the probability δ that a random valid ciphertext fails.
- Finding F failing ciphertexts takes F/δ queries to the CVO.

Failure boosting

[DVV18, DGJ$^+$19] introduce "failure boosting":

- Precompute many ciphertexts, query to CVO only those predicted likely to fail.
  - $\rightarrow$ lower query complexity.
  - $\rightarrow$ parallelisable/Groverisable precomputation.

Failure boosting

[DVV18, DGJ+19] introduce “failure boosting”:

- Precompute many ciphertexts, query to CVO only those predicted likely to fail.
  - $\rightarrow$ lower query complexity.
  - $\rightarrow$ parallelisable/Groverisable precomputation.

Significant recent advances

- 2 Dec 2019: non-failing ciphertexts can inform search for failing ones [BS20].
- 3 Dec 2019: once first failing ciphertext found, next are cheaper to find [DR**V**20].

Failure boosting

[DVV18, DGJ+19] introduce "failure boosting":

- Precompute many ciphertexts, query to CVO only those predicted likely to fail.
  - → lower query complexity.
  - → parallelisable/Groverisable precomputation.

Significant recent advances

- 2 Dec 2019: non-failing ciphertexts can inform search for failing ones [BS20].
- 3 Dec 2019: once first failing ciphertext found, next are cheaper to find [DR**V**20].

Take-home messages:

→ failure probability should be negligible for long term keys.

→ one-shot analysis not enough.

Choosing parameters in practice

So, you want to pick parameters. What options are out there for costing attacks?

$\rightarrow$ LWE estimator.

$\rightarrow$ NIST PQC submission packages.

$\rightarrow$ Increasingly, new attack papers are providing scripts.

Choosing parameters in practice

So, you want to pick parameters. What options are out there for costing attacks?
→ LWE estimator.
→ NIST PQC submission packages.
→ Increasingly, new attack papers are providing scripts.

Gotchas

- They are community efforts, often buggy.
- Not every attack/model is implemented.
  → And comparisons across scripts may be misleading!
- They still require understanding of the attacks to avoid pitfalls.
- Hit and miss documentation.

The LWE estimator [APS15]

- Covers primal and dual attack (with exhaustive guessing).
- Can specify non-Gaussian secret distributions (small uniform, fixed weight).
- Returns "rop", or ring ($\mathbb{Z}_q$) operations.
- LGPLv3+ license.

The LWE estimator [APS15]

- Covers primal and dual attack (with exhaustive guessing).
- Can specify non-Gaussian secret distributions (small uniform, fixed weight).
- Returns "rop", or ring ($\mathbb{Z}_q$) operations.
- LGPLv3+ license.

Watch out!

- Sometimes lags behind some improvements (no "dimensions for free" yet :(, still uses GSA).
- It does not cost the hybrid primal and dual attacks.
- It does not cost CCA attacks.

General advice

If you use a third party estimator

Code review it! (And contribute fixes/features back!)

General advice

If you use a third party estimator

Code review it! (And contribute fixes/features back!)

If you write a paper with estimates

Simplify peer review:

$\rightarrow$ "What commit of the script did you use?"

$\rightarrow$ Provide example code! How did you call the estimator script?

General advice

If you use a third party estimator

Code review it! (And contribute fixes/features back!)

If you write a paper with estimates

Simplify peer review:
$\rightarrow$ "What commit of the script did you use?"
$\rightarrow$ Provide example code! How did you call the estimator script?

If you write a new estimator

$\rightarrow$ Make it public!
$\rightarrow$ Make it good! Put comments, unit tests, some documentation.

OMG so many references

There's lots of moving parts and it could be tempting to cut corners.

Still, stay safe out there!

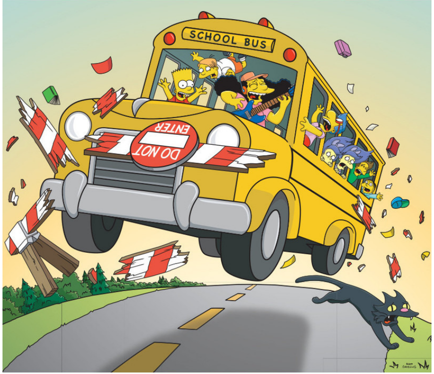

You can find a copy of the slides with bibliography at https://fundamental.domains

Martin R Albrecht, Shi Bai, Pierre-Alain Fouque, Paul Kirchner, Damien Stehlé, and Weiqiang Wen.
Faster enumeration-based lattice reduction: Root hermite factor $k^{1/(2k)}$ time $k^{k/8+o(k)}$.
In *Annual International Cryptology Conference*, pages 186–212. Springer, 2020.

Martin R. Albrecht, Shi Bai, Jianwei Li, and Joe Rowell.
Lattice reduction with approximate enumeration oracles: Practical algorithms and concrete performance.
Cryptology ePrint Archive, Report 2020/1260, 2020.
https://eprint.iacr.org/2020/1260.

Martin R Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin E Lauter, et al.
Homomorphic encryption standard.
IACR Cryptol. ePrint Arch., 2019:939, 2019.

Martin R Albrecht, Benjamin R Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W Postlethwaite, Fernando Virdia, and Thomas Wunderer.
Estimate all the {LWE, NTRU} schemes!
In *SCN*, 2018.

Martin R. Albrecht, Léo Ducas, Gottfried Herold, Elena Kirshanova, Eamonn W. Postlethwaite, and Marc Stevens.
The general sieve kernel and new records in lattice reduction.
In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 717–746. Springer, Heidelberg, May 2019.

Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe.

Post-quantum key exchange - A new hope.
In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 327–343. USENIX Association, August 2016.

Martin R. Albrecht, Vlad Gheorghiu, Eamonn W. Postlethwaite, and John M. Schanck.
Estimating quantum speedups for lattice sieves.
Cryptology ePrint Archive, Report 2019/1161, 2019.
https://eprint.iacr.org/2019/1161.

Martin R. Albrecht.
On dual lattice attacks against small-secret LWE and parameter choices in HElib and SEAL.
In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 103–129. Springer, Heidelberg, April / May 2017.

Yoshinori Aono, Phong Q. Nguyen, and Yixin Shen.
Quantum lattice enumeration and tweaking discrete pruning.
In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 405–434. Springer, Heidelberg, December 2018.

Martin R Albrecht, Rachel Player, and Sam Scott.
On the concrete hardness of learning with errors.
JMC, 2015.

Yoshinori Aono, Yuntao Wang, Takuya Hayashi, and Tsuyoshi Takagi.
Improved progressive bkz algorithms and their precise cost estimation by sharp simulator.
In *EUROCRYPT*, 2016.

Nina Bindel and John M Schanck.
Decryption failure is more likely after success.
In *International Conference on Post-Quantum Cryptography*, pages 206–225. Springer, 2020.

Yuanmi Chen.
Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe.
PhD thesis, Université Paris Diderot, 2013.

Jung Hee Cheon, Minki Hhan, Seungwan Hong, and Yongha Son.
A hybrid of dual and meet-in-the-middle attack on sparse and ternary secret lwe.
IEEE Access, 7:89497–89506, 2019.

Yuanmi Chen and Phong Q. Nguyen.
BKZ 2.0: Better lattice security estimates.
In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 1–20. Springer, Heidelberg, December 2011.

Jan-Pieter D'Anvers, Qian Guo, Thomas Johansson, Alexander Nilsson, Frederik Vercauteren, and Ingrid Verbauwhede.
Decryption failure attacks on IND-CCA secure lattice-based schemes.
In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 565–598. Springer, Heidelberg, April 2019.

Jan-Pieter D'Anvers, Mélissa Rossi, and Fernando Virdia.
(one) failure is not an option: Bootstrapping the search for failures in lattice-based encryption schemes.
In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 3–33. Springer, 2020.

Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi.
LWE with side information: Attacks and concrete security estimation.
In *CRYPTO*, 2020.

Léo Ducas.
Shortest vector from lattice sieving: A few dimensions for free.
In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 125–145. Springer, Heidelberg, April / May 2018.

Jan-Pieter D'Anvers, Frederik Vercauteren, and Ingrid Verbauwhede.
On the impact of decryption failures on the security of LWE/LWR based schemes.
Cryptology ePrint Archive, Report 2018/1089, 2018.
https://eprint.iacr.org/2018/1089.

Thomas Espitau, Antoine Joux, and Natalia Kharchenko.
On a hybrid approach to solve small secret lwe.
Cryptology ePrint Archive, Report 2020/515, 2020.
https://eprint.iacr.org/2020/515.

Scott Fluhrer.
Cryptanalysis of ring-LWE based key exchange with key share reuse.
Cryptology ePrint Archive, Report 2016/085, 2016.
http://eprint.iacr.org/2016/085.

Nicolas Gama and Phong Q Nguyen.
Finding short lattice vectors within mordell's inequality.

In *STOC*, 2008.

Rosario Gennaro and Matthew J. B. Robshaw, editors.
CRYPTO 2015, Part I, volume 9215 of *LNCS*. Springer, Heidelberg, August 2015.

Nick Howgrave-Graham.
A hybrid lattice-reduction and meet-in-the-middle attack against ntru.
In *CRYPTO*, 2007.

Éliane Jaulmes and Antoine Joux.
A chosen-ciphertext attack against NTRU.
In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 20–35. Springer, Heidelberg, August 2000.

Daniele Micciancio and Michael Walter.
Practical, predictable lattice basis reduction.
In *EUROCRYPT*, 2016.

Eamonn W. Postlethwaite and Fernando Virdia.
On the success probability of solving unique SVP via BKZ.
Cryptology ePrint Archive, Report 2020/1308, 2020.
https://eprint.iacr.org/2020/1308.

Claus-Peter Schnorr and M Euchner.
Lattice basis reduction: Improved practical algorithms and solving subset sum problems.
In *FCT*, 1991.