Estimate all the {LWE, NTRU} schemes!

Martin R. Albrecht[1], Benjamin R. Curtis[1], Amit Deo[1], Alex Davidson[1], **Rachel Player**[1,2], Eamonn Postlethwaite[1], Fernando Virdia[1], Thomas Wunderer[3]

April 12, 2018

[1] Information Security Group, Royal Holloway, University of London, UK
[2] Sorbonne Université, CNRS, INRIA
Laboratoire d'Informatique de Paris 6, LIP6, Équipe PolSys, France
[3] Technische Universität Darmstadt, Germany

Security of LWE- and NTRU-based NIST proposals

- Several approaches for solving LWE and NTRU problems
- Most require lattice reduction
- Disagreement in the literature about estimating lattice reduction
- More precisely, disagreement in **cost model** of BKZ

By cost model we mean the combination of the cost of solving SVP in dimension β and the number of SVP oracle calls required.

Cost models used in NIST proposals

Model	Cost
Core-Sieve	$2^{0.292\beta}$
Q-Core-Sieve	$2^{0.265\beta}$
Core-Sieve+$O(1)$	$2^{0.292\beta+16.4}$
Q-Core-Sieve+$O(1)$	$2^{0.265\beta+16.4}$
Core-Sieve (min. space)	$2^{0.368\beta}$
Q-Core-Sieve (min. space)	$2^{0.2975\beta}$
β-Sieve	$\beta\, 2^{0.292\beta}$
Q-β-Sieve	$\beta\, 2^{0.265\beta}$
$8d$-Sieve+$O(1)$	$8d\, 2^{0.292\beta+16.4}$
Q-$8d$-Sieve+$O(1)$	$8d\, 2^{0.265\beta+16.4}$
Core-Enum+$O(1)$	$2^{0.187\beta \log \beta - 1.019\beta + 16.1}$
Q-Core-Enum+$O(1)$	$2^{(0.187\beta \log \beta - 1.019\beta + 16.1)/2}$
$8d$-Enum (quadratic fit)+$O(1)$	$8d\, 2^{0.000784\beta^2 + 0.366\beta - 0.9}$
LOTUS-Enum	$2^{0.125\beta \log \beta - 0.755\beta + 2.25}$

This work

- We consider all LWE- and NTRU- based proposals
- We identify each of the cost models used
- We estimate the security of each proposal according to each of the cost models

Our goal is not to declare a favourite scheme, a favourite cost model, a favourite methodology, etc. Instead we are showing the discrepancies in the concrete security estimation space.

Our scripts wraps the LWE estimator [APS15]

In this project, we added support for

- arbitrary balanced bounded uniform (including sparse) distributions
- rotations of the secret vector during hybrid attacks, needed for tighter NTRU estimates

Pressing open problem: LWE estimator would benefit from code review!

M. R. Albrecht, R. P. and S. Scott. On the concrete hardness of Learning with Errors. In *Journal of Mathematical Cryptology, 9(3):169–203*, 2015.

https://estimate-all-the-lwe-ntru-schemes.github.io/

Estimate all the {LWE, NTRU} schemes!

Complexity estimates for running the primal-uSVP and dual attacks against all LWE-based, and the primal-uSVP attack against all NTRU-based, Round 1 schemes proposed as part of the PQC process run by NIST. We make use of the [APS15] estimator. The code for generating this table is available on Github, as well as the paper. Clicking on a particular estimate cell in the table will provide with stand-alone Sagemath code for reproducing the estimate.

Below, we provide LWE-equivalent parameters, where n = LWE secret dimension, k = MLWE rank (if any), q = modulo, σ = standard deviation of the error, $\mathbb{Z}_q/(\phi)$ is the ring (if any). For NTRU schemes we provide ‖f‖, ‖g‖ = lengths of the short polynomials. If you spot a mistake in a parameter set or cost model, please feel free to open a ticket or to make a pull-request.

LWE n samples | LWE 2n samples | NTRU | 14 selected | Search:

Scheme	Assumption	Primitive	Parameters	Claimed security	NIST Category	Attack	Proposed BKZ cost models: Q-Core-Sieve	Q-Core-Sieve + O(1)	Q-Core-Sieve (min space)	Q-β-Sieve	Q-8d-Sieve + O(1)	Core-Sieve
BabyBear	ILWE	KEM	n = 624, k =	152	2	primal	153	169	172	163	183	169
BabyBear	ILWE	KEM	n = 624, k =	152	2	dual	193	206	211	202	218	207
BabyBear	ILWE	KEM	n = 624, k =	141	2	primal	143	159	160	152	172	157
BabyBear	ILWE	KEM	n = 624, k =	141	2	dual	180	191	197	186	205	193
CRYSTALS-Dilithium	MLWE	SIG	n = 768, k =	91	1	primal	92	108	104	101	122	102
CRYSTALS-Dilithium	MLWE	SIG	n = 768, k =	91	1	dual	110	123	120	117	135	119
CRYSTALS-Dilithium	MLWE	SIG	n = 1024, k =	125	2	primal	130	146	146	139	160	143
CRYSTALS-Dilithium	MLWE	SIG	n = 1024, k =	125	2	dual	149	163	165	158	176	163
CRYSTALS-Dilithium	MLWE	SIG	n = 1280, k =	158	3	primal	159	175	179	168	190	175
CRYSTALS-Dilithium	MLWE	SIG	n = 1280, k =	158	3	dual	179	193	199	187	206	195
CRYSTALS-Kyber	MLWE	KEM, PKE	n = 512, k =	102	1	primal	103	119	115	111	132	113

Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn Postlethwaite, Fernando Virdia, Thomas Wunderer.

Estimate all the {LWE, NTRU} schemes! 

Complexity estimates for running the primal-uSVP and dual attacks against all LWE-based, and the primal-uSVP attack against all NTRU-based, Round 1 schemes proposed as part of the PQC process run by NIST. We make use of the [APS15] estimator. The code for generating this table is available on Github, as well as the paper. Clicking on a particular estimate cell in the table will provide with stand-alone Sagemath code for reproducing the estimate.

Below, we provide LWE-equivalent parameters, where n = LWE secret dimension, k = MLWE rank (if any), q = modulo, σ = standard deviation of the error, $\mathbb{Z}_q/(\phi)$ is the ring (if any). For NTRU schemes we provide ‖f‖, ‖g‖ = lengths of the short polynomials. If you spot a mistake in a parameter set or cost model, please feel free to open a ticket or to make a pull-request.

LWE n samples | LWE 2n samples | NTRU | 14 selected | Search:

Scheme	Assumption	Primitive	Parameters	Claimed security	NIST Category	Attack	Proposed BKZ cost models					
												Core-Sieve
BabyBear	ILWE	KEM	n = 624, k =	152	2							169
BabyBear	ILWE	KEM	n = 624, k =	152	2							207
BabyBear	ILWE	KEM	n = 624, k =	141	2							157
BabyBear	ILWE	KEM	n = 624, k =	141	2							193
CRYSTALS-Dilithium	MLWE	SIG	n = 768, k =	91	1							102
CRYSTALS-Dilithium	MLWE	SIG	n = 768, k =	91	1							119
CRYSTALS-Dilithium	MLWE	SIG	n = 1024, k =	125	2							143
CRYSTALS-Dilithium	MLWE	SIG	n = 1024, k =	125	2	dual						163
CRYSTALS-Dilithium	MLWE	SIG	n = 1280, k =	158	3	primal	159	175	179	168	190	175
CRYSTALS-Dilithium	MLWE	SIG	n = 1280, k =	158	3	dual	179	193	199	187	206	195
CRYSTALS-Kyber	MLWE	KEM, PKE	n = 512, k =	102	1	primal	103	119	115	111	132	113

BabyBear – Q-Core-Sieve

```
# To reproduce the estimate run this snippet on http://aleph.sagemath.org/
# Ring ops: 143
# Block size: 536
# Dimension: 1115
load('https://bitbucket.org/malb/lwe-estimator/raw/HEAD/estimator.py')
n = 624
sd = 0.7905694150420949
q = 1024
alpha = sqrt(2*pi)*sd/RR(q)
m = n
secret_distribution = "normal"
success_probability = 0.99
reduction_cost_model =  lambda beta, d, B: ZZ(2)**RR(0.265*beta)
primal_usvp(n, alpha, q, secret_distribution=secret_distribution,
m=m, success_probability=success_probability,
```

Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn Postlethwaite, Fernando Virdia, Thomas Wunderer.

Comparing quantum cost estimates

- NIST proposed maximum quantum circuit depth which not all schemes take into account
- Instead some proposals use asymptotic Q- cost model
- Different ways to interpret cost model e.g. for goal of "AES128 key recovery" hardness:
  - Aim for Q-cost $\geq 2^{128} \approx 128$ "quantum-bits" security
  - Aim for Q-cost $\geq 2^{64} \approx$ cost of Grover for AES128 key search

Pressing open problem: agree on how to estimate quantum security

Cost model swaps: what?

- There are many examples where under one cost model, scheme A appears harder to break than scheme B, while under another cost model, scheme B appears harder to break

Scheme	Parameter set	Core-Sieve	Core-Enum+O(1)
EMBLEM	$n = 611$	76	142
uRound2.KEM	$n = 500$	84	126

Table 1: Example highlighted by Bernstein on PQC forum.

Different cost models give different tradeoffs

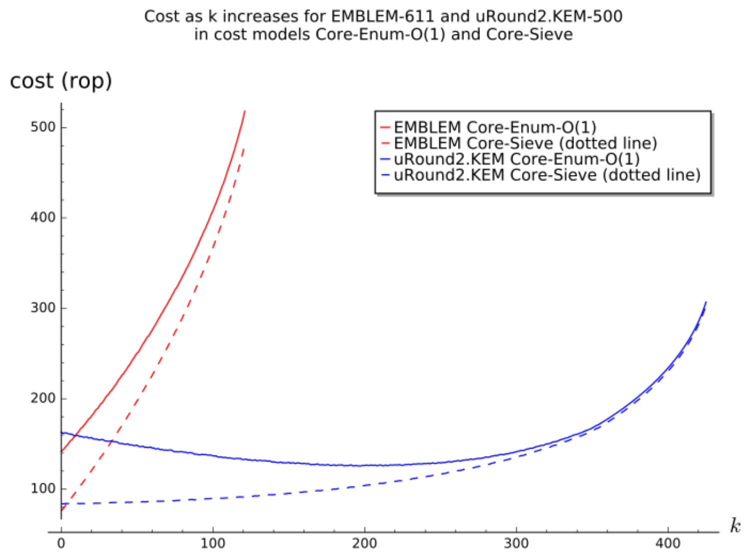

Cost model swaps: why?

- Tradeoff: larger dimensional lattice reduction, or smaller dimensional lattice reduction and repeat
- Optimising for cost depends on the choice of cost model
- E.g. hybrid attack: steeper slope means the tradeoff will be more favourable as the number of guesses increases and dimension of lattice reduction decreases

Pressing open problem: from which β does sieving beat enumeration?

Our data visualised in graphs by Mike Hamburg

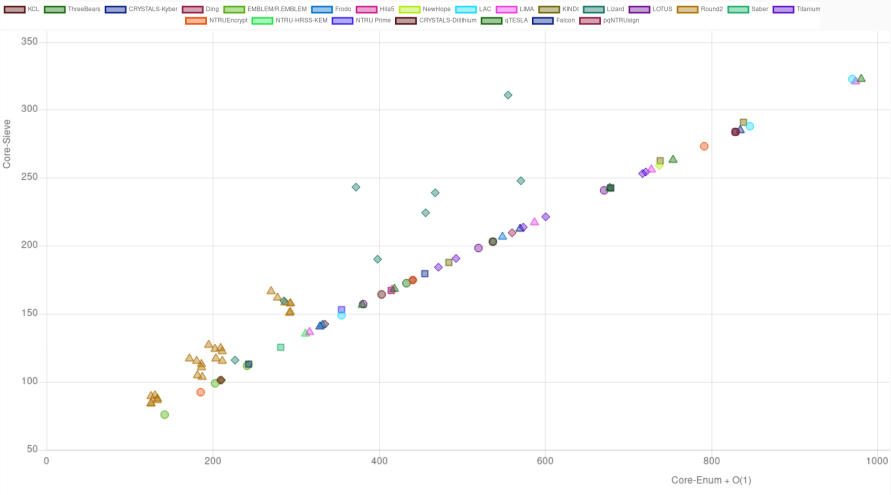

Graph generated at `https://bitwiseshiftleft.github.io/estimate-all-the-lwe-ntru-schemes.github.io/graphs`. Hamburg's page also uses performance data from the PQC lounge team, see `https://www.safecrypto.eu/pqclounge/`

Conclusion + thank you

Summary of open problems:

- Code review [APS15] estimator
- Better cost models for low β
- Agree on quantum security estimation

Website: `https://estimate-all-the-lwe-ntru-schemes.github.io`

Email: `rachel.player@lip6.fr`

Twitter: `@yayworthy`