Module Learning with Errors with non-uniform matrices

Joint Online Crypto Seminar - 12 January 2026

Katharina Boudgoust, joint work with Hannah Keller

@ Aarhus Crypto

Module Learning With Errors

Langlois, Stehlé DCC'15

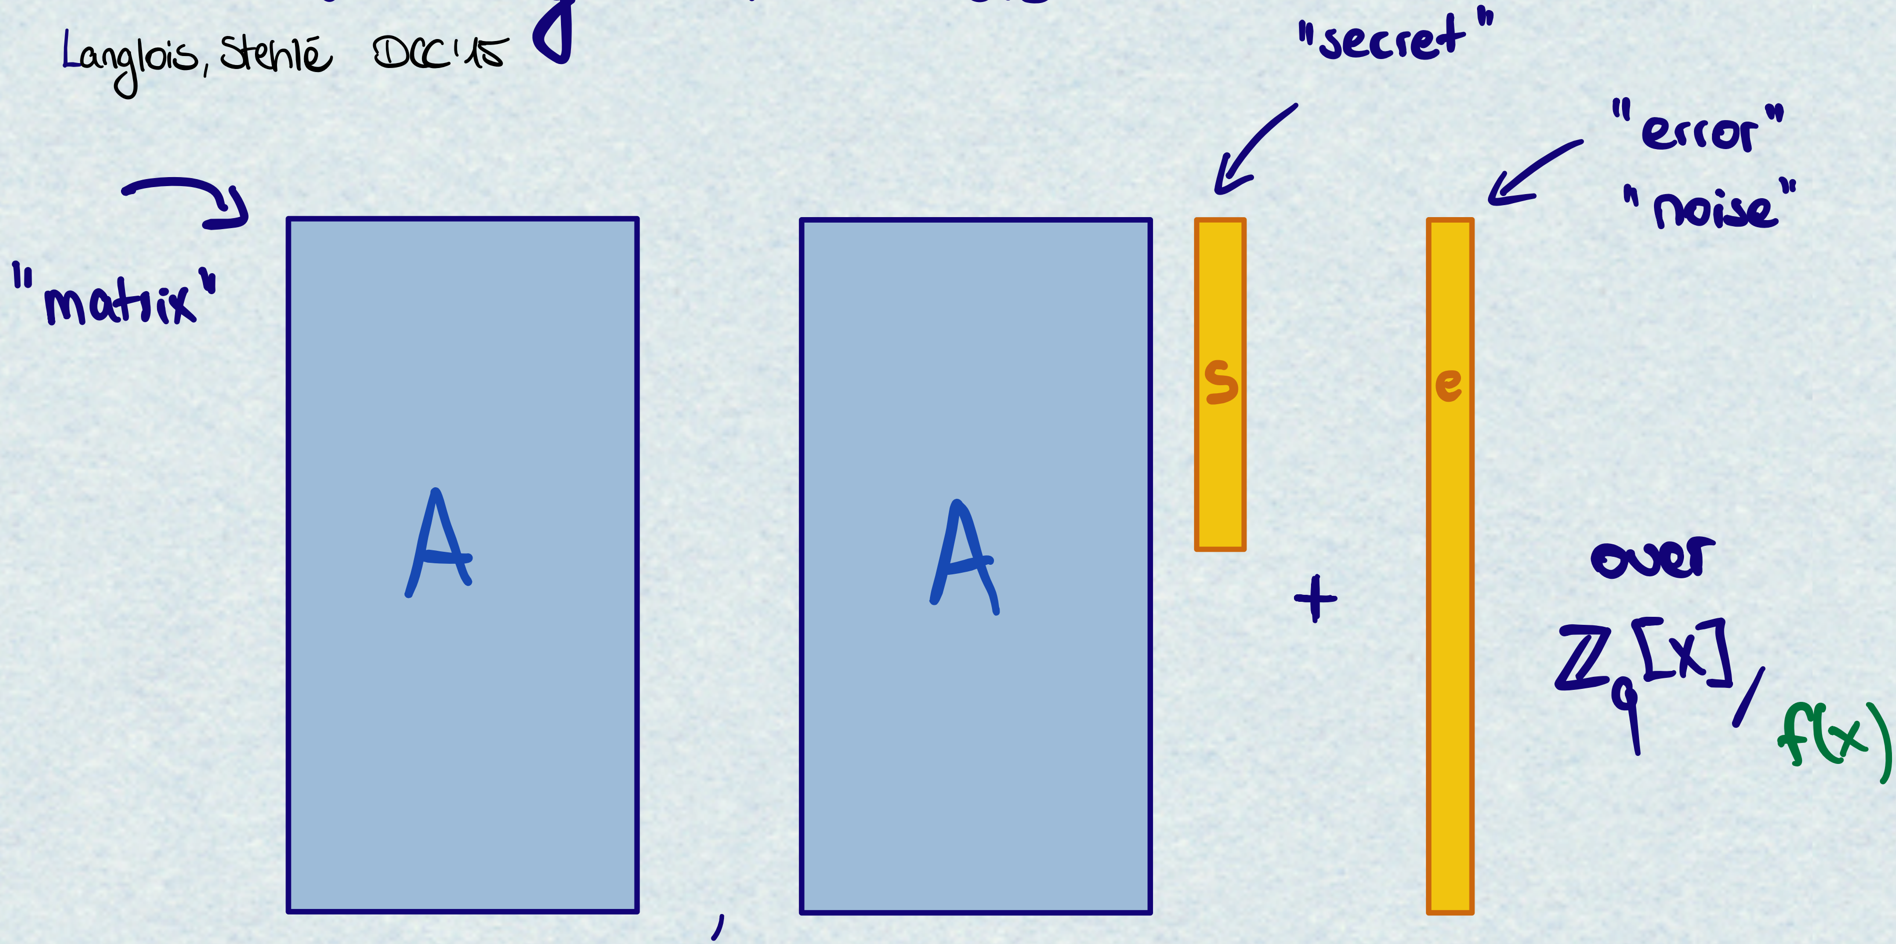

Search: find s (or e)

decision: distinguish from (A, unif)

Module Learning With Errors

Langlois, Stehlé DCC'15

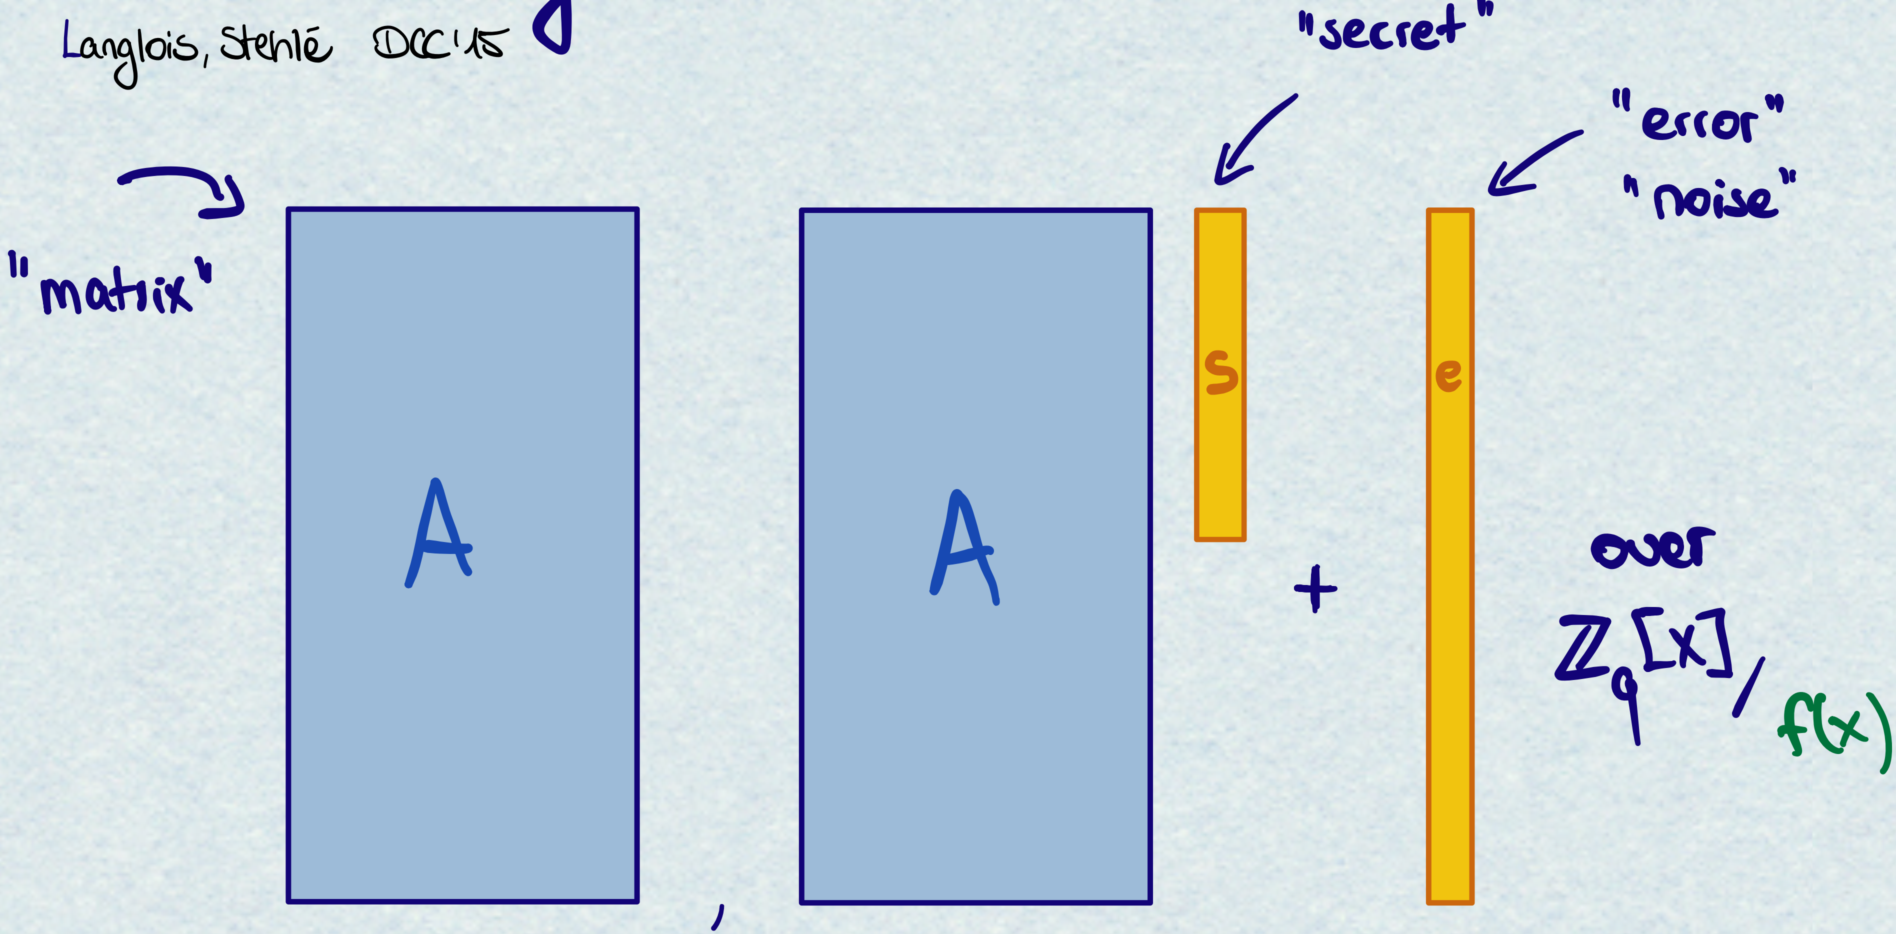

Choices:

* polynomial $f(x)$
* distribution for s
* distribution for e
* distribution for A

flexible usage but non-trivial security analysis

Standard Module Learning With Errors

Langlois, Stehlé DCC'15

Shortest Independent Vector Problem
for ANY Module Lattices

↓ quantum: Langlois, Stehlé DCC'15
classical: Boudgoust, Jeudy, Roux-Langlois, Wen Asiacrypt'20

* $f(x)$ cyclotomic polynomial
* s uniform over $\mathbb{Z}_q[x]/f(x)$
* e discrete Gaussian
* A uniform over $\mathbb{Z}_q[x]/f(x)$

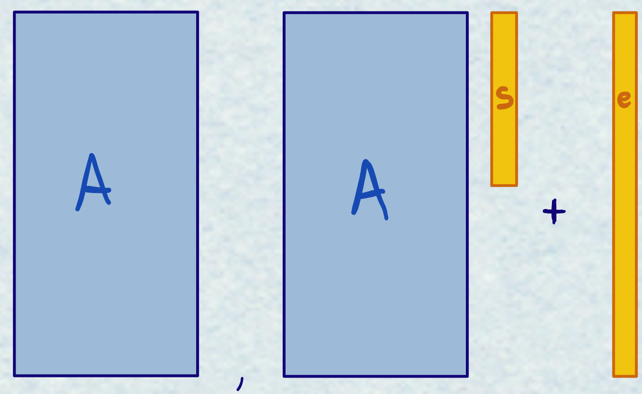

↓

Variants

Module Learning With Errors in Hermite Normal Form

Applebaum, Cash, Peikert, Sahai Crypto'09

s and e follow the same distribution

(both are short)

Module Learning With Errors in Hermite Normal Form

Applebaum, Cash, Peikert, Sahai Crypto'09

Same than noise distribution

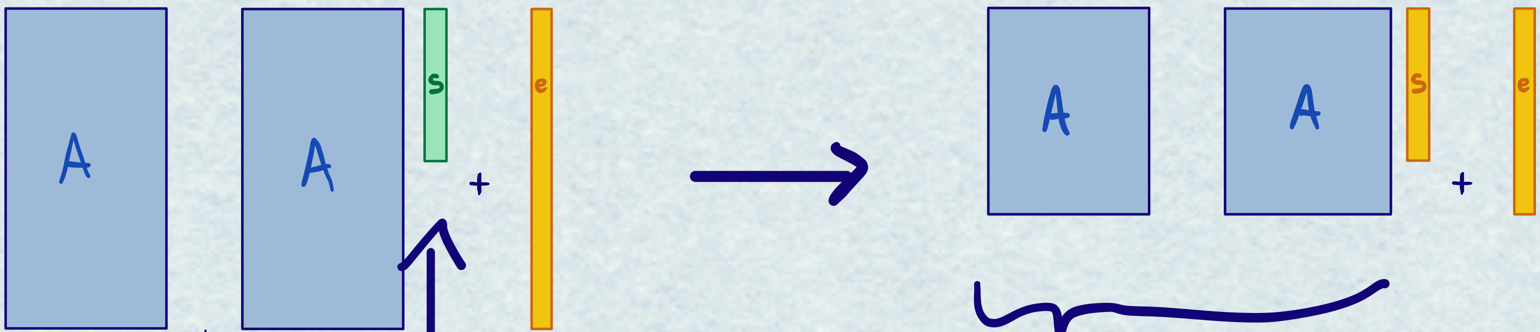

arbitrary secret distribution

with fewer* samples

* exact number depends on the underlying ring $\mathbb{Z}_q[x]/f(x)$

* Present in most practical schemes
* Used in the Lattice Estimator

Module Learning With Errors : Variants

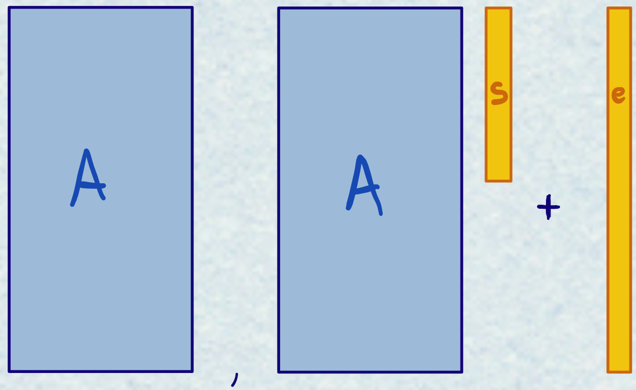

* polynomial $f(x)$: for now, all cyclotomics seem equally secure, but reductions are missing !

* distribution s : any, as long as enough min-entropy

Brakerski, Döttling TCC'20
Boudgoust, Jeudy, Roux-Langlois, Wen Indocrypt'22
Lin, Wang, Zhuang, Wang TCS'24

* distribution e : — " — and bounded norm

Boudgoust, Jeudy, Tairi, Wen E-print 2025/1472

Often tighter reductions for specific distributions

What about non-uniform matrices A ?

Some results ...

~~Module~~ Learning With Errors with binary matrix

Boneh, Lewi, Montgomery, Raghunathan Crypto'13

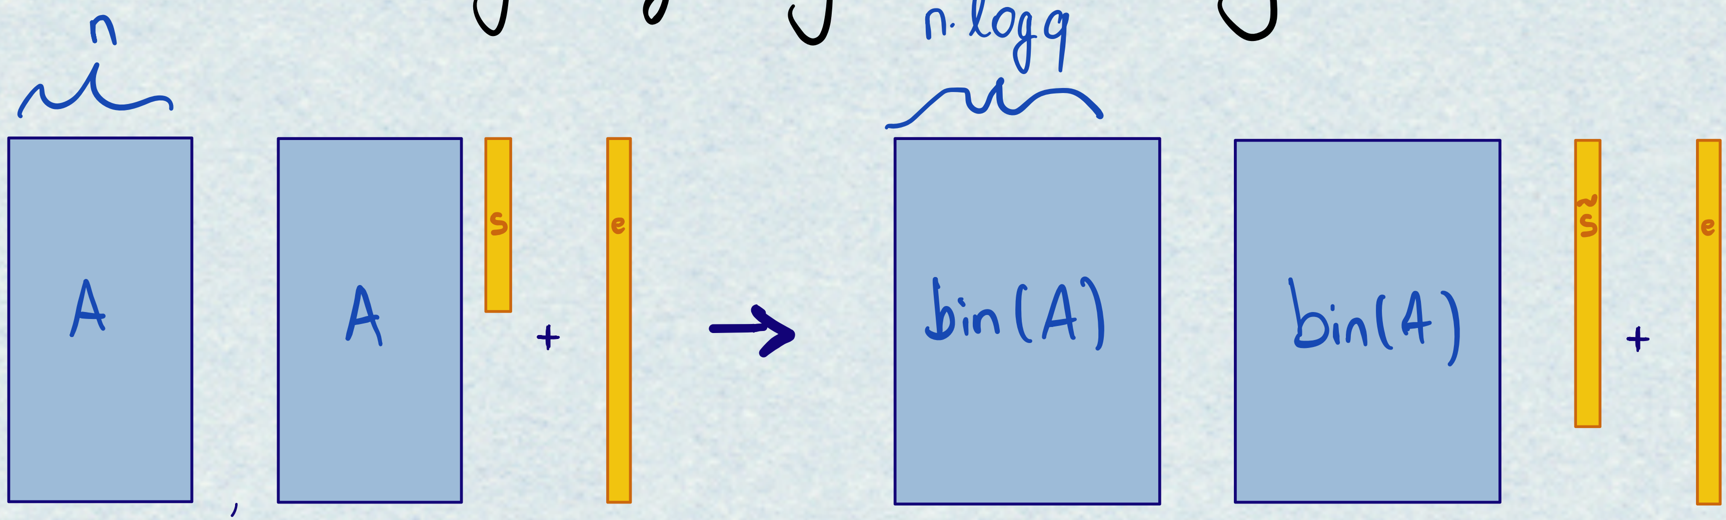

over $\mathbb{Z}_q$

$$\tilde{s} = \begin{pmatrix} \begin{matrix} 1 \\ 2 \\ \vdots \\ 2^{\log q} \end{matrix} & & & \\ & \begin{matrix} 1 \\ 2 \\ \vdots \\ 2^{\log q} \end{matrix} & & \\ & & \ddots & \\ & & & \begin{matrix} 1 \\ 2 \\ \vdots \\ 2^{\log q} \end{matrix} \end{pmatrix} \cdot s$$

"gadget" matrix

~~Module~~ Learning With Errors with binary matrix

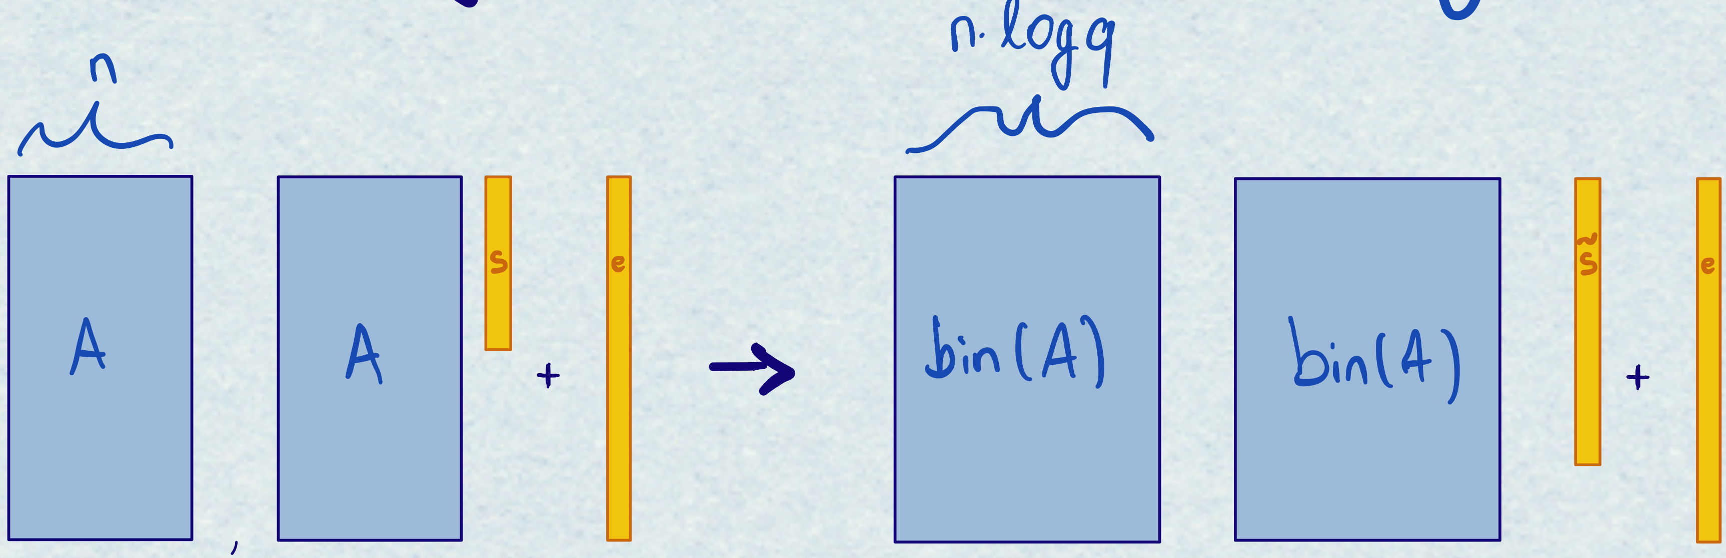

$\tilde{s}$ is large now

Actually, A, s and e can not all three be small at the same time

("Integer LWE" easy to solve)

Module Learning With Errors with

* A computationally close to uniform

Goldwasser, Kalai, Peikert, Vaikuntanathan ICS'10

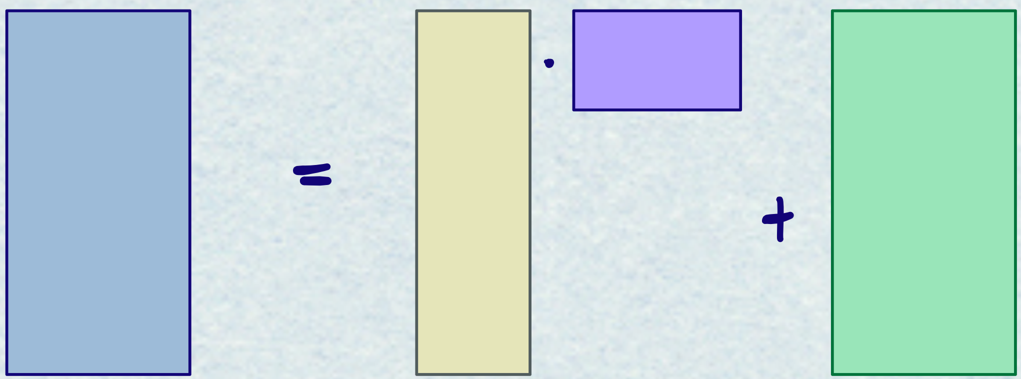

std multi-secret M-LWE instance

* A statistically / Rényi close to uniform using the Leftover Hash Lemma

Regev STOC'05

Bai, Lepoint, Roux-Langlois, Sakzad, Stehlé, Steinfeld JoC'18

* A sparse matrix

Jain, Lin, Saha Crypto'24

In our work:

for powers-of-2 cyclotomics

truncate the c low-order bits of A

$\mathcal{U}$ uniform over $\mathbb{Z}_q[x]/_{f(x)}$

$$A = \mathrm{Trunc}(\mathcal{U}, c) = \mathcal{U} - (\mathcal{U} \bmod 2^c)$$

c low-order bits of $\mathcal{U}$

$$\Rightarrow A = 0 \bmod 2^c$$

trivial setup: 2^c divides q and $\|e\| < 2^c$

$$\Rightarrow As + e \bmod 2^c = e$$

Why looking at this variant? pk-compression

In PKE:

$$pk = \mathrm{Trunc}(A, As+e) = A'$$

$$ct = A' \cdot r + f + \text{encoded}(msg)$$

A' is truncated

security argument based on truncated M-LWE

(was proposed for Kyber, then discarded)

* an alternative solution:

add random low-order bits (in the ROM)

Research Goal:

Reduce hardness of

truncated M-LWE

from

standard M-LWE

$$A = \mathrm{Trunc}(\mathcal{U}, c) = \mathcal{U} - N_{\mathcal{U}}$$

Approach 1

$$(A, As+e) \rightsquigarrow (\mathcal{U}, \mathcal{U}\cdot s + \underbrace{e - N_{\mathcal{U}}\cdot s}_{\approx e'})$$

use noise flooding to argue

$e - N_{\mathcal{U}}\cdot s \approx e'$ fresh noise

Statistical: super-poly $e \Rightarrow$ super-poly q

Rényi divergence: poly e and q, but only search in the paper

$A = \mathrm{Trunc}(\mathcal{U}, c) = \mathcal{U} - N_{\mathcal{U}}$

Approach 1

$$(A, As + e) \rightsquigarrow (\mathcal{U}, \mathcal{U} \cdot s + \underbrace{e - N_{\mathcal{U}} \cdot s}_{\approx e'})$$

use noise flooding to argue

$e - N_{\mathcal{U}} \cdot s \approx e'$ fresh noise

+ Rényi noise flooding applies to many distributions

+ $\|e\| \gg \underbrace{\|N_u\|}_{\sim 2^c} \cdot \|s\|$ $\rightsquigarrow$ disallows trivial setup

− not for decision (yet): large distance between $\mathcal{U}$ and A

What about hardness

of decision

truncated M-LWE ?

$$A = \mathrm{Trunc}(\mathcal{U}, c) = \mathcal{U} - N_{\mathcal{U}}$$

Approach 2 via M-LWE with hints on s

$(A, As+e, H, Hs+f)$ f noise

$\approx$ $\|H\|$ bounded

$(A, \mathrm{Unif}, H, Hs+f)$

adversarially chosen depending on A

HNF M-LWE
with Gaussians $\Rightarrow$ M-LWE with hints on s

$\Downarrow$

Bermudo Mera, Karmakar, Marc, Soleimanian
PKC'22

M-LWE with truncated A

$$\underbrace{A = \mathrm{Trunc}(\mathcal{U}, c) = \mathcal{U} - N_{\mathcal{U}}}$$

Approach 2

M-LWE with hints on s

$\Downarrow$

M-LWE with truncated A

$$(\mathcal{U}, \mathcal{U}s + e_1, \underbrace{-N_{\mathcal{U}}, -N_{\mathcal{U}}s + e_2}_{\text{hint on secret } s}) \rightsquigarrow (A, As + e)$$

$e = e_1 + e_2$

Gaussian composition

Zoom out

Work	Assumption	Variant	Distr. s	Distr. e
Jia, Zhang, Wang IET'23	Module NTRU	Search	any, enough min-entropy	Gaussian
Approach 1)	Module LWE	Search	Bounded	Rényi-close
Approach 2)	Module LWE	Decision	Gaussian	Gaussian

Open Questions:

* decision hardness for non-Gaussians
* additive and multiplicative transformations

Zoom out

Thanks!

Work	Assumption	Variant	Distr. s	Distr. e
Jia, Zhang, Wang IET '23	Module NTRU	Search	any, enough min-entropy	Gaussian
Approach 1)	Module LWE	Search	Bounded	Rényi-close
Approach 2)	Module LWE	Decision	Gaussian	Gaussian

ia.cr/2025/120

Open Questions:

* decision hardness for non-Gaussians
* additive and multiplicative transformations