Mod-NTRU trapdoors and applications

Alexandre Wallet

Lattices: From Theory to Practice

Simons Institute, 29/04/2020

Based on a joint work with Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé and Keita Xagawa, ePrint 2019/1456

NTT

Today's talk

A larger class of almost "optimal" trapdoors from NTRU modules

Known applications: (not detailed today)

(A) New meaningful security/efficiency trade-offs for GPV signatures
Acceptably efficient PKE/KEM à la NTRUEncrypt

(B) Extension of [DLP'14]'s IBE

(A) see our article (B) Cheon, Kim, Kim, and Son, ePrint 2019/1468

Roadmap

1. Lattice trapdoors, NTRU lattices
2. Hard NTRU lattices with half-trapdoors
3. Completing the trapdoor, application to signatures

Lattice trapdoors

Parity-check lattices

For $\mathbf{A} \in \mathbb{Z}^{m \times n}$ and $q \in \mathbb{Z}$

$$\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m \ : \ \mathbf{x}\mathbf{A} = \mathbf{0} \bmod q\}.$$

$$\mathbf{x} \ \overset{n}{\mathbf{A}}\ {}^{m} = 0 \bmod q$$

[Ajt'96] $(\Lambda_q^{\perp}(\mathbf{A}))_{\mathbf{A}}$ are "hard lattices": for $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{m \times n})$, $\mathsf{SIS}_{m,q} \geq \mathsf{SIVP}_{\mathrm{poly}(n)}$

A trapdoor is a **short** basis $\mathbf{B}$ of $\Lambda_q^{\perp}(\mathbf{A})$.

($\|\mathbf{B}\|_{\max} := \max_i \|\mathbf{b}_i\|$ is small)

$$\mathbf{B}\ \mathbf{A} = 0 \bmod q$$

What is "optimal"? $\|\widetilde{\mathbf{B}}\|_{\max} \approx \mathrm{Vol}(\Lambda_q^{\perp}(\mathbf{A}))^{1/m}$, where $\widetilde{\mathbf{B}} = \mathsf{GSO}(\mathbf{B})$.

Canonical example: GPV signatures

If $\mathbf{B}$ is basis of $\Lambda_q^{\perp}(\mathbf{A})$, then $\mathbf{BA} = \mathbf{0} \bmod q$

Simplified Sign$_{\mathbf{B}}(\mathrm{msg})$:

- $\mathbf{c}$ such that $\mathbf{cA} = \mathcal{H}(\mathrm{msg})$
- $\mathbf{v} \leftarrow D_{\mathcal{L}(\mathbf{B}),\mathbf{c},\sigma}$ with TheSampler†
- Signature: $\mathbf{s} = \mathbf{c} - \mathbf{v}$.

Simplified Verif$_{\mathbf{A}}(\mathrm{msg}, \mathbf{s})$:

- If $\|\mathbf{s}\|$ too big, refuse.
- If $\mathbf{sA} \neq \mathcal{H}(\mathrm{msg})$, refuse.
- Accept.

Requirements

σ small $\Rightarrow \widetilde{\mathbf{B}}$ short

$\mathbf{B}$ Gaussian of std.dev. $\sigma \Rightarrow \|\mathbf{s}\| \approx \sigma\sqrt{m}$
Want n and q s.t. SIS$_{m,q,\sigma\sqrt{m}}$ is hard

Hard to compute $\mathbf{B}$ from $\mathbf{A}$

Easy to generate $(\mathbf{A}, \mathbf{B})$

$\left.\right\}$ Method determines $m = m(n, q)$.

†: remember Thomas' talk

Development of lattice trapdoors

Algorithms to generate trapdoored hard lattices:

$\widetilde{\mathbf{B}} = \mathsf{GSO}(\mathbf{B})$

- [Ajt'99] $\mathbf{A}$ hard and $\|\mathbf{B}\|_{\max} = O(m^{5/2})$.

  × optimal
  × practical

- [AP'09] $\mathbf{A}$ hard, $m = \Omega(n \log q)$
  $\|\widetilde{\mathbf{B}}\|_{\max} = O(\sqrt{n \log q})$

  ✓ optimal
  × practical

Development of lattice trapdoors

Algorithms to generate trapdoored hard lattices:

$\widetilde{\mathbf{B}} = \mathsf{GSO}(\mathbf{B})$

- [Ajt'99] $\mathbf{A}$ hard and $\|\mathbf{B}\|_{\max} = O(m^{5/2})$.

  ✗ optimal
  ✗ practical

- [AP'09] $\mathbf{A}$ hard, $m = \Omega(n \log q)$
  $\|\widetilde{\mathbf{B}}\|_{\max} = O(\sqrt{n \log q})$

  ✓ optimal
  ✗ practical

- [MP'12] Meaningful improvements
  But still $\|\widetilde{\mathbf{B}}\| = O(\sqrt{n \log q})$

  getting there!

Development of lattice trapdoors

Algorithms to generate trapdoored hard lattices:

$\widetilde{\mathbf{B}} = \mathsf{GSO}(\mathbf{B})$

- [Ajt'99] $\mathbf{A}$ hard and $\|\mathbf{B}\|_{\max} = O(m^{5/2})$.

  ✗ optimal
  ✗ practical

- [AP'09] $\mathbf{A}$ hard, $m = \Omega(n \log q)$
  $\|\widetilde{\mathbf{B}}\|_{\max} = O(\sqrt{n \log q})$

  ✓ optimal
  ✗ practical

- [MP'12] Meaningful improvements
  But still $\|\widetilde{\mathbf{B}}\| = O(\sqrt{n \log q})$

  getting there!

- [DLP'14] $\mathbf{A}$ an NTRU lattice, $m = 2n$
  $\|\widetilde{\mathbf{B}}\|_{\max} \approx \sqrt{q}$

  ✓ optimal
  ✓ practical

- **Today**: $\mathbf{A}$ an NTRU lattice, $m = cn$
  $\|\widetilde{\mathbf{B}}\|_{\max} \approx q^{\frac{1}{c}}$.

NTRU modules

$R = \mathbb{Z}[X]/(\phi), \deg \phi = n$, irreducible.
q a prime

$f = \sum_i f_i X^i$
$(f_0, \ldots, f_{n-1})$ or $\mathsf{T}(f)$ multiplication matrix

$\mathbf{F} \in R^{m \times m}$ invertible mod q, $\mathbf{G} \in R^{m \times k}$

$$\mathbf{H} = \mathbf{F}^{-1} \left. \overbrace{\mathbf{G}}^{k} \right|_{m} \bmod q$$

NTRU modules

$R = \mathbb{Z}[X]/(\phi), \deg \phi = n$, irreducible.
q a prime

$f = \sum_i f_i X^i$
$(f_0, \ldots, f_{n-1})$ or $\mathsf{T}(f)$ multiplication matrix

$\mathbf{F} \in R^{m \times m}$ invertible mod q, $\mathbf{G} \in R^{m \times k}$

$$\mathbf{H} = \mathbf{F}^{-1} \underbrace{\mathbf{G}}_{k \times m} \bmod q$$

$$\mathcal{L}_{\mathsf{NTRU}}^{m,k} := \Lambda_q^{\perp}([\mathbf{H} | -\mathbf{I}_k]) = \{(\mathbf{u}, \mathbf{v}) \in R^{(m+k)} \ : \ \mathbf{uH} - \mathbf{v} = \mathbf{0} \bmod q\},$$

(full) rank $(m+k)n$ lattice with volume q^{kn}

easy (public) basis:

$$\begin{bmatrix} \mathbf{I}_{mn} & \mathsf{T}(\mathbf{H}) \\ \mathbf{0} & q\mathbf{I}_{kn} \end{bmatrix}$$

Minima, covering radius, smoothing parameter all are $\approx q^{k/(m+k)}$

Use of NTRU modules

Non exhaustive; all of these are for $m = k = 1$

PKE/KEM:

- NTRUEncrypt [HPS'98]
- NTRUEnc-HRSS [HH+'17]
- NTRUPrime [BCLV'17]

Signatures:

- NTRUSign [HHS+'03]
- Falcon (from [DLP'14] from [GPV'08])
- BLISS [DDLL'13]

Advanced:

- HE [LTV'12]
- Multilinear maps [GGH'13]
- IBE [DLP'14]

Where are we?

How to generate a useful NTRU module

Trapdoor basis $\mathbf{B} = \begin{bmatrix} \mathbf{F} & \mathbf{G} \\ * & * \end{bmatrix}$ should give us $\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max} \approx q^{k/(m+k)}$

Lemma: If $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_{m+k}]$, then:

$$\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max} = \max_i\{\|\widetilde{\mathbf{b}}_1\|, \ldots, \|\widetilde{\mathbf{b}}_{m+k}\|\} \geq q^{k/(m+k)}$$

A starter: take $s \approx q^{k/(m+k)}$

1) Sample $\mathbf{b}_i \leftarrow D_{R,s}^{m+k}$ for $1 \leq i \leq m$
2) Parse as $[\mathbf{b}_1, \ldots, \mathbf{b}_m] = [\mathbf{F}|\mathbf{G}]$; restart if $\mathbf{F}$ not invertible mod q

How to generate a useful NTRU module

Trapdoor basis $\mathbf{B} = \begin{bmatrix} \mathbf{F} & \mathbf{G} \\ * & * \end{bmatrix}$ should give us $\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max} \approx q^{k/(m+k)}$

Lemma: If $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_{m+k}]$, then:

$$\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max} = \max_i\{\|\widetilde{\mathbf{b}}_1\|, \ldots, \|\widetilde{\mathbf{b}}_{m+k}\|\} \geq q^{k/(m+k)}$$

A starter: take $s \approx q^{k/(m+k)}$

1) Sample $\mathbf{b}_i \leftarrow D_{R,s}^{m+k}$ for $1 \leq i \leq m$
2) Parse as $[\mathbf{b}_1, \ldots, \mathbf{b}_m] = [\mathbf{F}|\mathbf{G}]$; restart if $\mathbf{F}$ not invertible mod q

Caveat: orthogonal projections shrink vectors by some factor γ_i
$\Rightarrow \mathbf{b}_1$ will be maximal, completion of basis will compensate.

How to generate a useful NTRU module

Trapdoor basis $\mathbf{B} = \begin{bmatrix} \mathbf{F} & \mathbf{G} \\ * & * \end{bmatrix}$ should give us $\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max} \approx q^{k/(m+k)}$

Lemma: If $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_{m+k}]$, then:

$$\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max} = \max_i\{\|\widetilde{\mathbf{b}}_1\|, \ldots, \|\widetilde{\mathbf{b}}_{m+k}\|\} \geq q^{k/(m+k)}$$

A better start: set $s_i \approx \gamma_i \cdot q^{k/(m+k)}$

1) Sample $\mathbf{b}_i \leftarrow D_{R,s_i}^{m+k}$ for $1 \leq i \leq m$
2) Parse as $[\mathbf{b}_1, \ldots, \mathbf{b}_m] = [\mathbf{F}|\mathbf{G}]$; restart if $\mathbf{F}$ not invertible mod q

Output a half-trapdoor for $\mathbf{H} = \mathbf{F}^{-1}\mathbf{G} \bmod q$.

Remaining problems:

- Is $\Lambda_q^{\perp}(\mathbf{H})$ a hard lattice ?
- How to complete the basis?
- Will the completion be nice?

How hard are trapdoored NTRU lattices?

"NTRU assumption"

Computational

Hard to compute $\mathbf{F}, \mathbf{G}$ from $\mathbf{H}$

Well, if not, it's not a trapdoor...

Decisional

Hard to distinguish $\mathbf{H}$ from $\mathcal{U}(R_q^{m\times k})$

Needed for $\Lambda_q^{\perp}(\mathbf{H})$ to be "hard"

Call $\mathcal{E}_s$ the distribution of $\mathbf{H} = \mathbf{F}^{-1}\mathbf{G} \bmod q$

How hard are trapdoored NTRU lattices?

"NTRU assumption"

Computational

Hard to compute $\mathbf{F}, \mathbf{G}$ from $\mathbf{H}$

Well, if not, it's not a trapdoor...

Decisional

Hard to distinguish $\mathbf{H}$ from $\mathcal{U}(R_q^{m \times k})$

Needed for $\Lambda_q^{\perp}(\mathbf{H})$ to be "hard"

Call $\mathcal{E}_s$ the distribution of $\mathbf{H} = \mathbf{F}^{-1}\mathbf{G} \bmod q$

New result: $\Phi = X^n + 1$, n a power of two, $q \equiv 3 \bmod 8$, for $3k \geq m \geq 1$

When $s \geq \widetilde{O}(n \cdot q^{\frac{k}{m+k}})$, then $\mathcal{E}_s \approx \mathcal{U}(R_q^{m \times k})$

[SS'11] for $m = k = 1$, the result hold for all q.

Strongly supports hardness of the trapdoored NTRU lattices

On the uniformity of the public basis

New result: $\Phi = X^n + 1$, n a power of two, $q \equiv 3 \bmod 8$, for $3k \geq m \geq 1$,

$$\text{when } s \geq \widetilde{O}(n \cdot q^{\frac{k}{m+k}}) \text{, then } \mathcal{E}_s \approx \mathcal{U}(R_q^{m \times k})$$

Intermediate useful result:

if $q = \mathfrak{p}_1 \ldots \mathfrak{p}_r$, when $s \geq \widetilde{O}(n \cdot q^{\frac{1}{2r}})$, then $\mathbb{P}_{\mathbf{F} \leftarrow D_{R,s}^{m \times m}}[\mathbf{F} \text{ invertible} \bmod q] \geq 1 - \frac{4n}{q^{n/r}}$

On the uniformity of the public basis

New result: $\Phi = X^n + 1$, n a power of two, $q \equiv 3 \bmod 8$, for $3k \geq m \geq 1$,

when $s \geq \widetilde{O}(n \cdot q^{\frac{k}{m+k}})$, then $\mathcal{E}_s \approx \mathcal{U}(R_q^{m\times k})$

Intermediate useful result:

if $q = \mathfrak{p}_1 \ldots \mathfrak{p}_r$, when $s \geq \widetilde{O}(n \cdot q^{\frac{1}{2r}})$, then $\mathbb{P}_{\mathbf{F}\leftarrow D_{R,s}^{m\times m}}[\mathbf{F} \text{ invertible } \bmod q] \geq 1 - \frac{4n}{q^{n/r}}$

Proof ideas/tools:

- Inspired of [SS'11] and [LPR'13]
- Involve module "multi-lattices" (additive subgroups of $\mathcal{M}_m(R)$, see also [BF'11])
- {Mod q invertibles} is not a lattice; our strategy to describe it:

  inclusion/exclusion over *all* lattices containing $q\mathcal{M}_m(R)$

  (They correspond to *all* r-uples of subspaces of $(\mathbb{F}_{q^{n/r}})^m$)

Generating a somewhat short basis[1]

From now on, $k = 1$ and $m \geq 1$.

$$\mathbf{h} = \mathbf{F}^{-1}\,\mathbf{g} \quad {}_m \bmod q \qquad \text{with } [\mathbf{F}|\mathbf{g}] = [\mathbf{b}_1, \ldots, \mathbf{b}_n] \text{ and } \mathbf{b}_i \leftarrow D_{R,s_i}^{m+1}$$

Now, need $(\mathbf{f}', g') \in R^{m+1}$ such that

$$D := \det \begin{bmatrix} \mathbf{F} & \mathbf{g} \\ \mathbf{f}' & g' \end{bmatrix} = q$$

Generating a somewhat short basis[1]

From now on, $k = 1$ and $m \geq 1$.

$$\mathbf{h} = \mathbf{F}^{-1} \mathbf{g} \ \mathrm{mod}\ q \qquad \text{with } [\mathbf{F}|\mathbf{g}] = [\mathbf{b}_1, \dots, \mathbf{b}_n] \text{ and } \mathbf{b}_i \leftarrow D_{R,s_i}^{m+1}$$

Now, need $(\mathbf{f}', g') \in R^{m+1}$ such that

$$D := \det \begin{bmatrix} \mathbf{F} & \mathbf{g} \\ \mathbf{f}' & g' \end{bmatrix} = q$$

With Shur's complement and $\mathrm{adj}(\mathbf{F}) = \det(\mathbf{F}) \cdot \mathbf{F}^{-1} \in R^{m \times m}$:

$$\begin{aligned} D &= \det(\mathbf{F}) \cdot \det(g' - \mathbf{f}' \cdot \mathbf{F}^{-1} \cdot \mathbf{g}) \\ &= g' \cdot \underbrace{\det(\mathbf{F})}_{\substack{\text{known} \\ \in R}} - \mathbf{f}' \cdot \underbrace{\mathrm{adj}(\mathbf{F})\mathbf{g}}_{\substack{\text{known} \\ \in R^m}} \end{aligned}$$

Take $\mathbf{f}' = (\dots, 0, f_i', 0, \dots) \Rightarrow$ back to solving an NTRU equation

(remember Thomas' talk)

[1] For another approach, see Cheon et al. ePrint 2019/1468

Almost optimal trapdoors

Last problem: how large is $\mathbf{b}_{m+1} = (\mathbf{f}', g')$?

Fact 1: $\|\widetilde{\mathbf{b}}_{m+1}\| \geq \dfrac{q}{\prod_i \|\widetilde{\mathbf{b}_i}\|}$

Since all $\|\widetilde{\mathbf{b}_i}\|$'s are about $q^{1/(m+1)}$, $\|\widetilde{\mathbf{b}}_{m+1}\|$ should be, too.

Fact 2: $\|\widetilde{\mathbf{b}}_{m+1}\|$ computable from $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_m$ without knowing $\mathbf{b}_{m+1}$

Almost optimal trapdoors

Last problem: how large is $\mathbf{b}_{m+1} = (\mathbf{f}', g')$?

Fact 1: $\|\widetilde{\mathbf{b}}_{m+1}\| \geq \dfrac{q}{\prod_i \|\widetilde{\mathbf{b}_i}\|}$

Since all $\|\widetilde{\mathbf{b}_i}\|$'s are about $q^{1/(m+1)}$, $\|\widetilde{\mathbf{b}}_{m+1}\|$ should be, too.

Fact 2: $\|\widetilde{\mathbf{b}}_{m+1}\|$ computable from $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_m$ without knowing $\mathbf{b}_{m+1}$

Finishing the trapdoor generation:

1) for $1 \leq i \leq m$, resample any vector that is too far from $q^{1/(m+1)}$
2) Compute $\|\widetilde{\mathbf{b}}_{m+1}\|$, restart if too large
3) Compute $\mathbf{b}_{m+1}$ and output $(\mathbf{H}, \mathbf{B})$.

$\|\mathbf{b}_i\|$'s close to λ_i's, $\|\widetilde{\mathsf{T}}(\mathbf{B})\|_{\max}$ close to $\eta_\epsilon(\Lambda_q^\perp(\mathbf{H}))$

These trapdoors are almost optimal.

A practical application: Mod-Falcon[2]

	m	n	$\|\|s\|\|$	Qsec	Minimizing |sig| |vk|	Minimizing |sig| |sig|	Minimizing |sig|+|vk| |vk|	Minimizing |sig|+|vk| |sig|
Falcon-512	1	512	6598	109	897	658	28	1276
Falcon-1024	1	1024	9331	252	1793	1274	63	2508
Mod-Falcon	2	512	1512	**174**	1792	972	940	1438

security/efficiency trade-off for Falcon

[2] To appear at AsiaCCS 2020; all size expressed in bytes

A practical application: Mod-Falcon[2]

	m	n	$\|\|s\|\|$	Qsec	Minimizing |sig| |vk|	Minimizing |sig| |sig|	Minimizing |sig|+|vk| |vk|	Minimizing |sig|+|vk| |sig|
Falcon-512	1	512	6598	109	897	658	28	1276
Falcon-1024	1	1024	9331	252	1793	1274	63	2508
Mod-Falcon	2	512	1512	**174**	1792	972	940	1438

security/efficiency trade-off for Falcon

	|vk|	|sig|	Qsec
dilithium-III	1472	2701	125
qTesla-p-I	14880	2592	140
dilithium-IV	1760	3366	158
Mod-Falcon	**1792** **940**	**972** **1438**	**174**

more compact for equivalent security

[2] To appear at AsiaCCS 2020; all size expressed in bytes

Food for thoughts

Question 1: We have almost optimal trapdoors for $\mathbf{h} = \mathbf{F}^{-1}\mathbf{g}$

Can this be extended to almost optimal trapdoors for $\mathbf{H} = \mathbf{F}^{-1}\mathbf{G}$?

(main problem: how to complete the basis?)

Question 2: We can use them for signature/IBE.

Can we use these new trapdoors for something else?

Can half-trapdoors' usefulness be improved too?

Question 3: Extend uniformity results to all q's

And to more fields (Galois, all?)

Generally, find new tools/techniques to compute Gaussian mass of subsets

Also, related to repartition of algebraic numbers

Food for thoughts

Question 1: We have almost optimal trapdoors for $\mathbf{h} = \mathbf{F}^{-1}\mathbf{g}$

Can this be extended to almost optimal trapdoors for $\mathbf{H} = \mathbf{F}^{-1}\mathbf{G}$?

(main problem: how to complete the basis?)

Question 2: We can use them for signature/IBE.

Can we use these new trapdoors for something else?

Can half-trapdoors' usefulness be improved too?

Question 3: Extend uniformity results to all q's

And to more fields (Galois, all?)

Generally, find new tools/techniques to compute Gaussian mass of subsets

Also, related to repartition of algebraic numbers

Thank you!

Questions $=$ **?**