A Key-Recovery Attack against Mitaka in the t-Probing Model

Thomas Prest

PQShield

PKC 2023
https://pkc.iacr.org/2023/

Roadmap

Early lattice signatures

Roadmap

Early lattice signatures → Black-box leakage → Parallelepiped attack

Roadmap

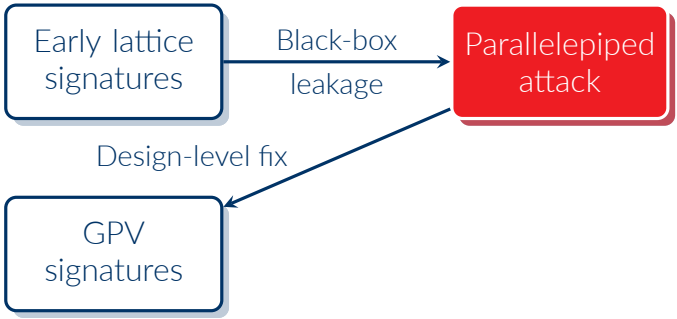
Early lattice signatures
Black-box leakage
Parallelepiped attack
Design-level fix
GPV signatures

Roadmap

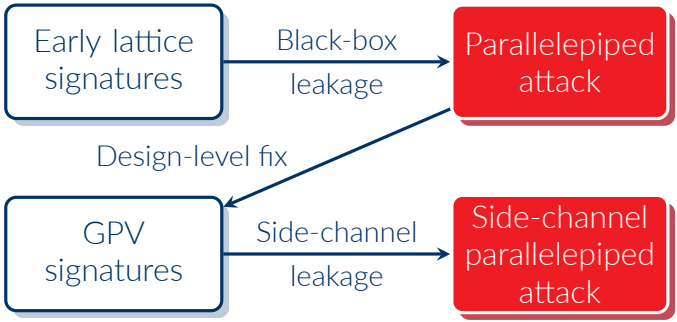
Early lattice signatures
Black-box leakage
Parallelepiped attack
Design-level fix
GPV signatures
Side-channel leakage
Side-channel parallelepiped attack

Roadmap

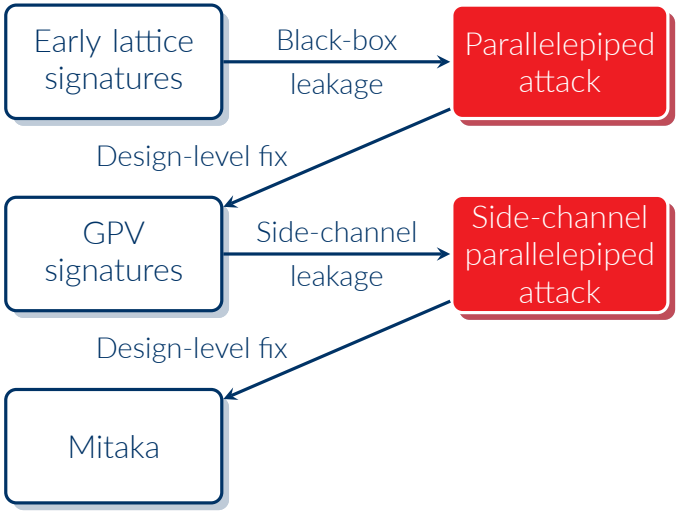
Early lattice signatures
Black-box leakage
Parallelepiped attack
Design-level fix
GPV signatures
Side-channel leakage
Side-channel parallelepiped attack
Design-level fix
Mitaka

Roadmap

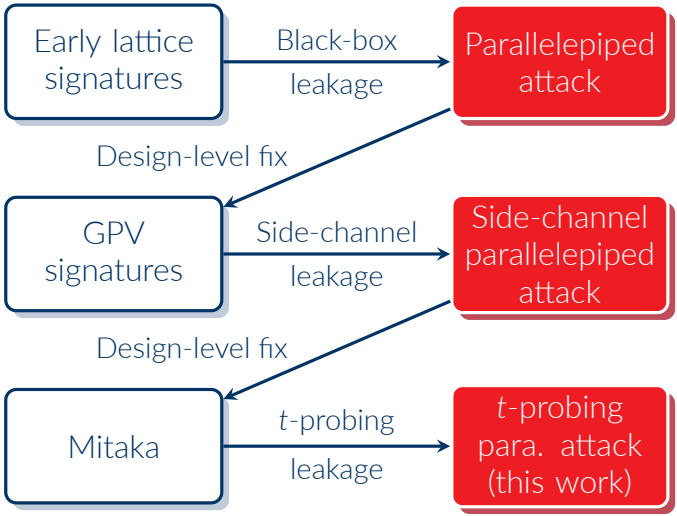
Early lattice signatures
Black-box leakage
Parallelepiped attack
Design-level fix
GPV signatures
Side-channel leakage
Side-channel parallelepiped attack
Design-level fix
Mitaka
t-probing leakage
t-probing para. attack (this work)

Hash-then-Sign

Initial attempts: NTRUSign (1997), GGHSign (2003)

Keygen(1^λ)

1. Gen. matrices $\mathbf{A}, \mathbf{B}$ such that:
   - $\mathbf{A}$ is pseudorandom
   - $\mathbf{A} \cdot \mathbf{B} = 0$
   - $\mathbf{B}$ has small coefficients
2. $\mathsf{pk} := \mathbf{A}, \mathsf{sk} := \mathbf{B}$

Sign(msg, sk = B)

1. Compute $\mathbf{c}$ such that $\mathbf{A} \cdot \mathbf{c} = H(\mathsf{msg})$
2. $\mathbf{v} := \mathbf{B} \lfloor \mathbf{B}^{-1}\mathbf{c} \rceil$
3. $\mathsf{sig} := \mathbf{s} = (\mathbf{c} - \mathbf{v})$

Verify(msg, pk = A, sig = s)

Check ($\mathbf{s}$ short) & ($\mathbf{A} \cdot \mathbf{s} = H(\mathsf{msg})$)

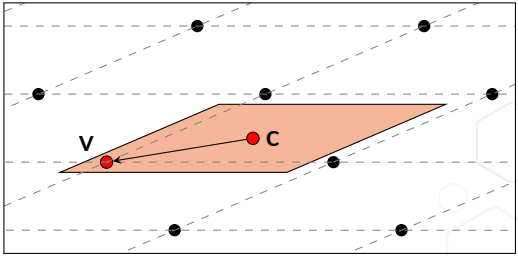

Initial attempts: NTRUSign (1997), GGHSign (2003)

Keygen(1^λ)

1. Gen. matrices $\mathbf{A}, \mathbf{B}$ such that:
   - $\mathbf{A}$ is pseudorandom
   - $\mathbf{A} \cdot \mathbf{B} = \mathbf{0}$
   - $\mathbf{B}$ has small coefficients
2. $\mathsf{pk} := \mathbf{A}, \mathsf{sk} := \mathbf{B}$

Sign(msg, sk = B)

1. Compute $\mathbf{c}$ such that $\mathbf{A} \cdot \mathbf{c} = H(\mathsf{msg})$
2. $\mathbf{v} := \mathbf{B} \lfloor \mathbf{B}^{-1}\mathbf{c} \rceil$
3. $\mathsf{sig} := \mathbf{s} = (\mathbf{c} - \mathbf{v})$

Verify(msg, pk = A, sig = s)

Check ($\mathbf{s}$ short) & ($\mathbf{A} \cdot \mathbf{s} = H(\mathsf{msg})$)

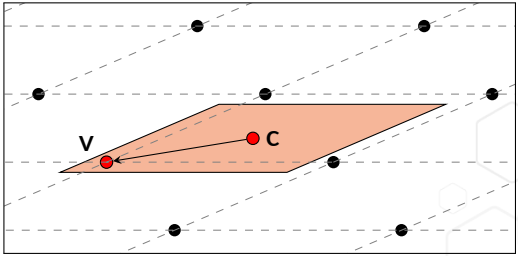

→ **Correctness:** easy

→ **Security:** Finding a <u>short</u> preimage $\mathbf{s}$ of $H(\mathsf{msg})$ should be difficult... or is it?

The parallelepiped attack

Problem: The distribution of the signature **s** is correlated to **B**

$$\mathbf{s} = \mathbf{c} - \mathbf{B}\left\lfloor \mathbf{B}^{-1}\mathbf{c} \right\rceil \in \mathbf{B} \cdot \left[-\frac{1}{2}, \frac{1}{2}\right]^n \tag{1}$$

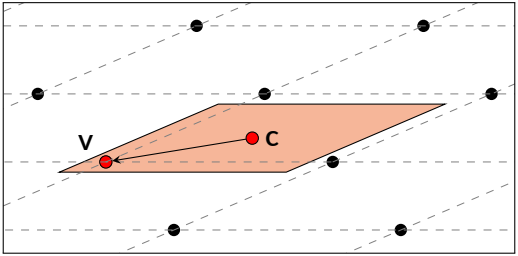

Given many signatures, **B** can be recovered using techniques from Independent Component Analysis (ICA)

→ 2006: key-recovery on NTRUSign and GGHSign

Design-level solution: trapdoor sampling à la "GPV"

$\mathbf{B}\lfloor \mathbf{B}^{-1}\mathbf{c} \rceil$

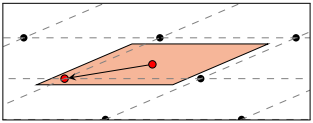

$\mathbf{B}\lfloor \mathbf{B}^{-1}\mathbf{c} \rceil_{\sigma_1}$

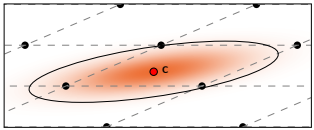

$\mathbf{B}\left\lfloor \mathbf{B}^{-1}\left(\mathbf{c} + \mathbf{M}\lfloor \mathbf{0} \rceil_{\sigma_2}\right)\right\rceil_{\sigma_1}$

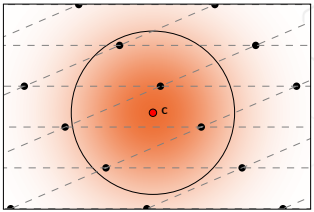

Indistinguishability: For appropriately chosen parameters[1], the rightmost procedure outputs a distribution close to a perfect Gaussian $D_{\Lambda(\mathbf{B}),\mathbf{c},\sigma}$.

Consequence: these two worlds are indistinguishable (in the ROM)

1. Sample a short vector $\mathbf{s}$, then set $H(\mathbf{msg}) = \mathbf{A} \cdot \mathbf{s}$
2. Compute $H(\mathbf{msg})$, then use $\mathbf{B}$ to sample a short preimage $\mathbf{s}$ of $H(\mathbf{msg})$

[1] It suffices $\left(\sigma_2^2 \cdot \mathbf{M}^t\mathbf{M} + \sigma_1^2 \cdot \mathbf{B}^t\mathbf{B} = \sigma^2\mathbf{I}\right)$ for σ large enough. See (Peikert, CRYPTO 2010)

Signature schemes in the GPV family

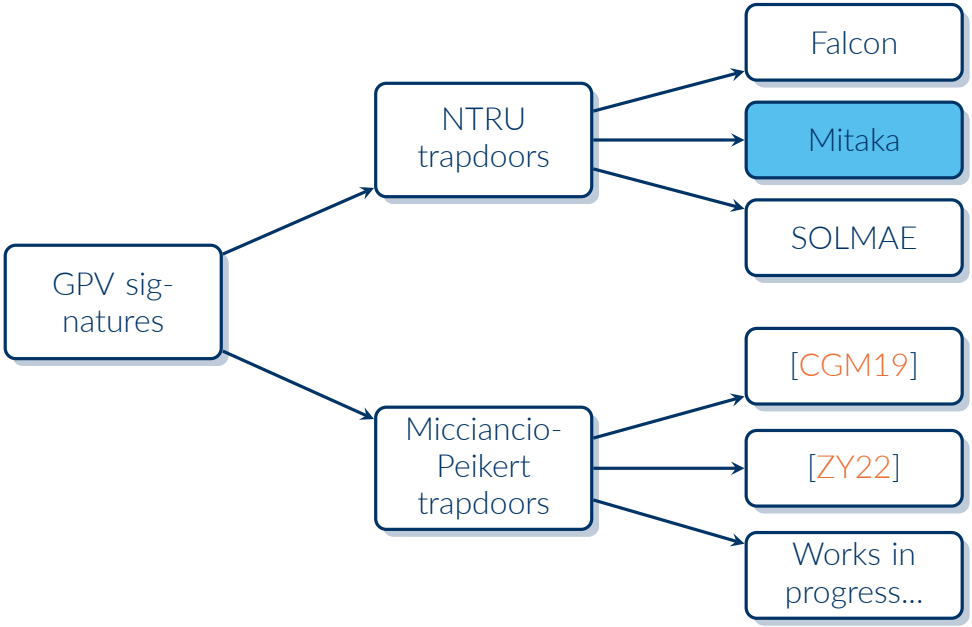

Side-Channel Attacks

Side-channel attacks in cryptography

Power consumption [KJJ99]

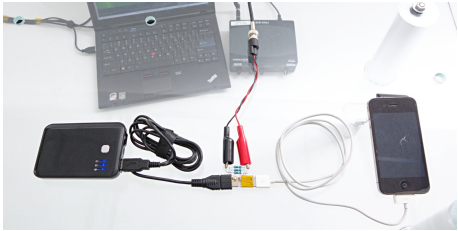

Electromagnetic emissions [Eck85]

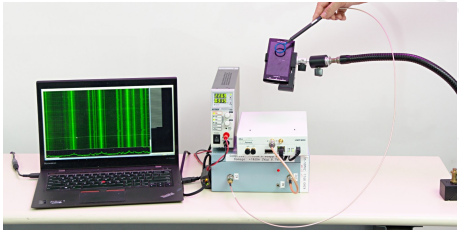

Timing measurement [Koc96]

Acoustic emissions [AA04]

The return of the parallelepiped attacks

In Falcon, a signature is $\mathbf{s} = \mathbf{c} - \mathbf{v}$, where $\mathbf{v} = \sum_i z_i \cdot \mathbf{b}_i$ for $\mathbf{B} = (\mathbf{b}_i)_i$ and $z_i \in \mathbb{Z}$
Monitoring the power consumption can provide information about the z_i.

This allows *side-channel assisted* parallelepiped attacks:

- *The Hidden Parallelepiped Is Back Again: Power Analysis Attacks on Falcon* [GMRR22] (TCHES 2022)
- *Improved Power Analysis Attacks on Falcon* [ZLYW23] (Eurocrypt 2023)

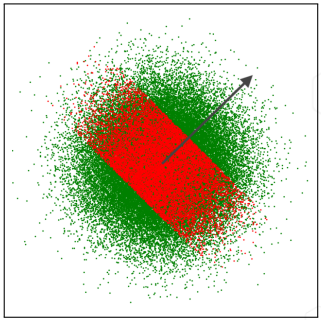

Illustration when $z_0 \in \{0, 1\}$ [GMRR22]

In general, the most robust countermeasure against side-channel attacks is masking.

The return of the parallelepiped attacks

In Falcon, a signature is $\mathbf{s} = \mathbf{c} - \mathbf{v}$, where $\mathbf{v} = \sum_i z_i \cdot \mathbf{b}_i$ for $\mathbf{B} = (\mathbf{b}_i)_i$ and $z_i \in \mathbb{Z}$
Monitoring the power consumption can provide information about the z_i.

This allows *side-channel assisted* parallelepiped attacks:

- *The Hidden Parallelepiped Is Back Again: Power Analysis Attacks on Falcon* [GMRR22] (TCHES 2022)
- *Improved Power Analysis Attacks on Falcon* [ZLYW23] (Eurocrypt 2023)

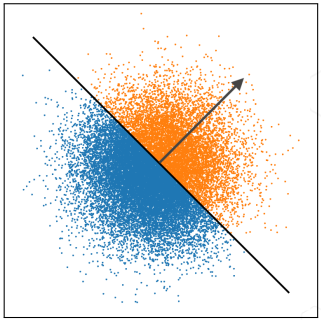

Illustration when $z_0 \in \mathbb{Z}+$ [ZLYW23]

In general, the most robust countermeasure against side-channel attacks is masking.

Masking and the t-probing model

t-probing model

- Adversary can probe t circuit values at runtime
- Unrealistic but a good starting point

Masking

- Each sensitive value x is split in d shares:

$$[\![x]\!] = (x_j)_{j \in [d]} \quad \text{such that} \quad \sum_j x_j = x \qquad (2)$$

- Perform operations using MPC techniques.
  - Linear operations $\rightarrow$ linear overhead
  - Multiplications $\rightarrow$ quadratic
  - More complex operations $\rightarrow \geq$ quadratic
- In "real life", security is exponential in d
- In t-probing model, ideally 0 leakage if $d > t$

The design of Mitaka

Mitaka (Eurocrypt 2022) is designed to be easy maskable (proof in the t-probing model):

→ (simplified) The signing procedure is

$$\mathbf{v} \leftarrow \mathbf{B} \left\lfloor \mathbf{B}^{-1} \left(\mathbf{c} + \mathbf{M} \lfloor \mathbf{0} \rceil_{\sigma_2} \right) \right\rceil_{\sigma_1}$$

The design of Mitaka

Mitaka (Eurocrypt 2022) is designed to be easy maskable (proof in the t-probing model):

- (simplified) The signing procedure is

$$\mathbf{v} \leftarrow \mathbf{B} \left\lfloor \mathbf{B}^{-1}\left(\mathbf{c} + \mathbf{M}\overbrace{\lfloor \mathbf{0} \rceil_{\sigma_2}}^{\text{offline}}\right) \right\rceil_{\sigma_1}$$

- $\lfloor \mathbf{0} \rceil_{\sigma_2}$ is sampled offline

The design of Mitaka

Mitaka (Eurocrypt 2022) is designed to be easy maskable (proof in the t-probing model):

- (simplified) The signing procedure is

$$\mathbf{v} \leftarrow \mathbf{B} \left\lfloor \underbrace{\mathbf{B}^{-1}\left(\mathbf{c} + \mathbf{M}\overbrace{\lfloor \mathbf{0} \rceil_{\sigma_2}}^{\text{offline}}\right)}_{\text{Easy-ish to mask}} \right\rceil_{\sigma_1}$$

- $\lfloor \mathbf{0} \rceil_{\sigma_2}$ is sampled offline
- Multiplications by $\mathbf{M}, \mathbf{B}, \mathbf{B}^{-1}$ are easy-ish to mask

The design of Mitaka

Mitaka (Eurocrypt 2022) is designed to be easy maskable (proof in the t-probing model):

- (simplified) The signing procedure is

$$\mathbf{v} \leftarrow \underbrace{\mathbf{B} \left\lfloor \underbrace{\mathbf{B}^{-1}\left(\mathbf{c} + \mathbf{M}\overbrace{\lfloor \mathbf{0} \rceil_{\sigma_2}}^{\text{offline}}\right)}_{\text{Easy-ish to mask}} \right\rceil_{\sigma_1}}_{?}$$

- $\lfloor \mathbf{0} \rceil_{\sigma_2}$ is sampled offline
- Multiplications by $\mathbf{M}, \mathbf{B}, \mathbf{B}^{-1}$ are easy-ish to mask
- What about $\lfloor \cdots \rceil_{\sigma_1}$? (next slide)

Masked Gaussian sampling

Algorithm 3 GaussShareByShare($[\![c]\!], r) \to [\![z]\!]$

Require: A standard deviation r, an arithmetic masking $[\![c]\!]$ for $c \in \frac{1}{C} \cdot \mathbb{Z}$, $B = \left\lceil \sqrt{2d} \right\rceil$.

Ensure: An arithmetic masking $[\![z]\!]$, where $z \overset{s}{\sim} D_{\mathbb{Z},c,r}$

```
repeat
    for j ∈ [d] do
        z_j ← D_{(1/B)·Z, c_j, r/√d}
    end for
    acc := Decode((z_j mod 1)_{j∈[d]})
until acc = 0
return [[z]] := (z_j)_{j∈[d]}
```

Figure 1: Masked discrete Gaussian sampling (in $\mathbb{Z}$) in Mitaka

- **Main idea:** sum of d Gaussians is a $\sqrt{d}$-times larger Gaussian[2]
- **Complexity:** $O(d)$ per sample

[2] For simplicity, we ignore **acc** (we suppose **acc** is always 0)

The Attack

Proof outline and its flaw

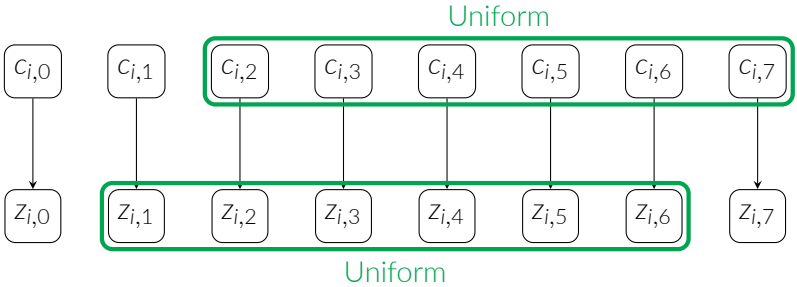

Proof outline:

- ✔ The input is a uniform encoding $[\![c_i]\!] = (c_{i,j})_{j \in [d]}$
- ✔ Same comment for the output $[\![z_i]\!] = (z_{i,j})_{j \in [d]}$
- ? So probing should leak nothing... right?

Proof outline and its flaw

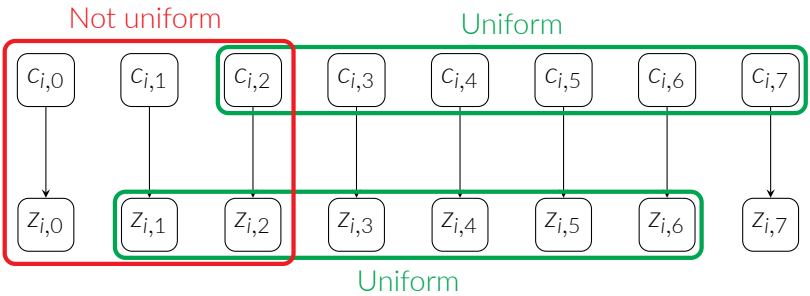

Flaw/contradiction:

- For each j, $(c_{i,j} - z_{i,j})$ is Gaussian
- $c_i - z_i = \sum_j (c_{i,j} - z_{i,j})$ is correlated to $c_{i,j} - z_{i,j}$

The attack

Let us note $\mathbf{v} = \mathbf{B} \cdot \mathbf{z} = \sum_i z_i \cdot \mathbf{b}_i$ the output of the trapdoor sampler.
$\mathbf{v}$ is distributed as a discrete Gaussian centered over $\mathbf{c}$.

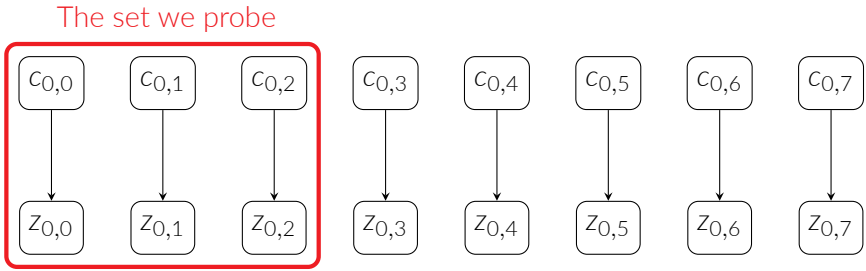

For each signing call we do this:

1. Probe $(z_{0,j}, c_{0,j})$ for $j \in [t_1]$, where $t_1 = \lfloor \frac{d-1}{2} \rfloor$
2. Compute $w = \sum_{j \in [t_1]} (c_{0,j} - z_{0,j})$
3. Since $w \cdot \mathbf{b}_0$ is an additive component of $\mathbf{s} = \mathbf{c} - \mathbf{v}$, this value tends to be > 0:

$$\langle w \cdot \mathbf{b}_0, \mathbf{s} \rangle = \langle w \cdot \mathbf{s}, \mathbf{b}_0 \rangle \tag{3}$$

This means the value $w \cdot \mathbf{s}$ is biased in the same direction as $\mathbf{b}_0$.

The (return of the)^2 parallelepiped attack

Estimator for $\mathbf{b}_0$:

$$\hat{\mathbf{b}}_0 = \frac{1}{\left(\sum_{\ell \in [N]} w_\ell^2\right)} \cdot \left(\sum_{\ell \in [N]} w_\ell \cdot \mathbf{s}_\ell\right). \tag{4}$$

One can show that $\hat{\mathbf{b}}_0 \sim \mathbf{b}_0 + X$, where X is a Gaussian of parameter σ_X:

$$\sigma_X \ll \sigma \cdot \sqrt{\frac{d}{t_1 \cdot N}} \tag{5}$$

where N is the number of signatures and traces.

Comments:

1. When $N \gtrsim 2^{21}$, we recover $\mathbf{b}_0$ using rounding (+ guessing)
2. When $N \lesssim 2^{21}$, we recover $\mathbf{b}_0$ using lattice reduction
3. When $N = \Omega(d)$,[3] we can recover $\mathbf{b}_0$ even with $t = O(1)$

[3] Large constant

Summary

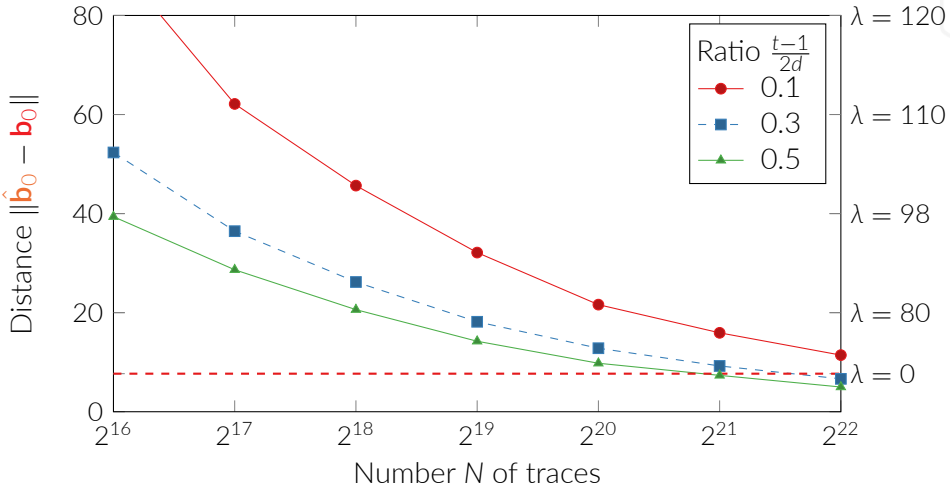

Figure 2: Distance $\|\hat{\mathbf{b}}_0 - \mathbf{b}_0\|$ as a function of # of traces (x-axis) and the ratio $\frac{t-1}{2d}$.
Right-side marks $\{\lambda = x\}$: core-SVP hardness of lattice reduction.
Under {- - -} line: immediate key-recovery via rounding (+ guessing).

Bibliography

First attempts

- NTRUSign [HHP+03]
- GGHSign [GGH97]

The parallelepiped attack

- Initial attack [NR06]
- Breaking countermeasures [DN12]

The GPV framework

- Trapdoor sampling [GPV08]
- A bunch of trapdoor samplers [Pei10, Pre15, DP16]

GPV-based signatures

- Falcon [PFH+20]
- Mitaka [EFG+22]

Side-channel parallelepiped attacks

- Power analysis attacks [GMRR22, ZLYW23]
- This paper

Thank you!

https://ia.cr/2023/157
https://tprest.github.io/

Dmitri Asonov and Rakesh Agrawal.
Keyboard acoustic emanations.
In *2004 IEEE Symposium on Security and Privacy*, pages 3–11. IEEE Computer Society Press, May 2004.

Yilei Chen, Nicholas Genise, and Pratyay Mukherjee.
Approximate trapdoors for lattices and smaller hash-and-sign signatures.
In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 3–32. Springer, Heidelberg, December 2019.

Léo Ducas and Phong Q. Nguyen.
Learning a zonotope and more: Cryptanalysis of NTRUSign countermeasures.
In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 433–450. Springer, Heidelberg, December 2012.

Léo Ducas and Thomas Prest.
Fast fourier orthogonalization.
In Sergei A. Abramov, Eugene V. Zima, and Xiao-Shan Gao, editors, *Proceedings of the ACM on International Symposium on Symbolic and Algebraic Computation, ISSAC 2016, Waterloo, ON, Canada, July 19-22, 2016*, pages 191–198. ACM, 2016.

Wim Van Eck.
Electromagnetic radiation from video display units: An eavesdropping risk?
Computers & Security, 4:269–286, 1985.

Thomas Espitau, Pierre-Alain Fouque, François Gérard, Mélissa Rossi, Akira Takahashi, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu.
Mitaka: A simpler, parallelizable, maskable variant of falcon.
In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 222–253. Springer, Heidelberg, May / June 2022.

Oded Goldreich, Shafi Goldwasser, and Shai Halevi.
Eliminating decryption errors in the Ajtai-Dwork cryptosystem.
In Burton S. Kaliski Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 105–111. Springer, Heidelberg, August 1997.

Morgane Guerreau, Ange Martinelli, Thomas Ricosset, and Mélissa Rossi.
The hidden parallelepiped is back again: Power analysis attacks on falcon.
IACR TCHES, 2022(3):141–164, 2022.

Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan.
Trapdoors for hard lattices and new cryptographic constructions.
In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.

Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte.
NTRUSIGN: Digital signatures using the NTRU lattice.
In Marc Joye, editor, *CT-RSA 2003*, volume 2612 of *LNCS*, pages 122–140. Springer, Heidelberg, April 2003.

Paul C. Kocher, Joshua Jaffe, and Benjamin Jun.
Differential power analysis.
In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 388–397. Springer, Heidelberg, August 1999.

Paul C. Kocher.
Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems.
In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 104–113. Springer, Heidelberg, August 1996.

Phong Q. Nguyen and Oded Regev.
Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures.
In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 271–288. Springer, Heidelberg, May / June 2006.

Chris Peikert.
An efficient and parallel Gaussian sampler for lattices.
In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 80–97. Springer, Heidelberg, August 2010.

Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang.
FALCON.
Technical report, National Institute of Standards and Technology, 2020.

available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions.

Thomas Prest.
Gaussian Sampling in Lattice-Based Cryptography.
PhD thesis, École Normale Supérieure, Paris, France, 2015.

Shiduo Zhang, Xiuhan Lin, Yang Yu, and Weijia Wang.
Improved power analysis attacks on falcon.
Cryptology ePrint Archive, Paper 2023/224, 2023.
https://eprint.iacr.org/2023/224.

Shiduo Zhang and Yang Yu.
Towards a simpler lattice gadget toolkit.
In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part I*, volume 13177 of *LNCS*, pages 498–520. Springer, Heidelberg, March 2022.