The Hidden Parallelepiped is Back Again: Power Analysis Attacks on Falcon

Morgane Guerreau¹

Ange Martinelli² Thom Mélissa Rossi²

Thomas Ricosset¹

¹Thales, ²ANSSI

September 20, 2022







Fast-Fourier Lattice-based Compact Signatures over NTRU



Two Power Analysis attacks on Falcon:

- Efficient DPA attack on the preimage computation
- STA on the trapdoor sampler leading to HPP attack



CVP is easy to solve with a good basis, but hard with a bad basis.



CVP is easy to solve with a good basis, but hard with a bad basis.

Basic signature scheme:

• Convert the message to sign to a vector \mathbf{c} in \mathcal{R}^n



CVP is easy to solve with a good basis, but hard with a bad basis.

Basic signature scheme:

- Convert the message to sign to a vector ${f c}$ in ${\cal R}^n$
- Use the good basis (secret key) to solve CVP



CVP is easy to solve with a good basis, but hard with a bad basis.

Basic signature scheme:

- Convert the message to sign to a vector c in Rⁿ
- Use the good basis (secret key) to solve CVP
- Anyone can verify the signature **v** with a bad basis (public key)



CVP is easy to solve with a good basis, but hard with a bad basis.

Basic signature scheme:

- Convert the message to sign to a vector c in Rⁿ
- Use the good basis (secret key) to solve CVP
- Anyone can verify the signature **v** with a bad basis (public key)

Note: It is hard to derivate the good basis from the bad basis.

• Message *m* is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$

• • = • • = •

- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{cB}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{c} \mathbf{B}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

 \implies **s** - **c** belongs to the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$.

- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{c} \mathbf{B}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

 \implies **s** - **c** belongs to the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$.



- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{c} \mathbf{B}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

 \implies **s** - **c** belongs to the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$.



Hidden Parallelepiped Problem: Recover **B** from independent samples drawn uniformly in $\mathcal{P}(\mathbf{B})$ [NR06].

- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{c} \mathbf{B}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

 \implies **s** - **c** belongs to the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$.



Hidden Parallelepiped Problem: Recover **B** from independent samples drawn uniformly in $\mathcal{P}(\mathbf{B})$ [NR06].

Deformed Parallelepiped Problem: Same as HPP when a *partial* perturbation is applied during round-off algorithm [DN12].

- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{c} \mathbf{B}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

 \implies **s** – **c** belongs to the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$.



Hidden Parallelepiped Problem: Recover **B** from independent samples drawn uniformly in $\mathcal{P}(\mathbf{B})$ [NR06].

Deformed Parallelepiped Problem: Same as HPP when a *partial* perturbation is applied during round-off algorithm [DN12].

GPV Framework: Replace the round-off algorithm by a trapdoor sampler: **s** is not always the closest vector [GPV08].

- Message m is hashed to a point $\mathbf{c} := H(m) \in \mathbb{Z}^n$
- Use round-off algorithm to solve CVP:

Signature: $\mathbf{s} = \lfloor \mathbf{c} \mathbf{B}^{-1} \rceil \mathbf{B}$ with \mathbf{B} the private basis

 $\Longrightarrow s-c$ belongs to the fundamental parallelepiped $\mathcal{P}(B).$



Hidden Parallelepiped Problem: Recover **B** from independent samples drawn uniformly in $\mathcal{P}(\mathbf{B})$ [NR06].

Deformed Parallelepiped Problem: Same as HPP when a *partial* perturbation is applied during round-off algorithm [DN12].

GPV Framework: Replace the round-off algorithm by a trapdoor sampler: **s** is not always the closest vector [GPV08].



Instantiation of GPV framework with NTRU lattices

> Instantiation of GPV framework with NTRU lattices
> Let R := Z[x]/(xⁿ + 1).
Private key: f,g,F,G ∈ R with fG - gF = q mod xⁿ + 1

> Instantiation of GPV framework with NTRU lattices > Let $\mathcal{R} := \mathbb{Z}[x]/(x^n + 1)$. Private key: $f, g, F, G \in \mathcal{R}$ with $fG - gF = q \mod x^n + 1$ Private basis: $\mathbf{B} := \begin{bmatrix} g & -f \\ G & -F \end{bmatrix} \Longrightarrow \mathbf{b_0} = (g_0, \dots, g_{n-1}, -f_0, \dots, -f_{n-1})$

> Instantiation of GPV framework with NTRU lattices > Let $\mathcal{R} := \mathbb{Z}[x]/(x^n + 1)$. Private key: $f, g, F, G \in \mathcal{R}$ with $fG - gF = q \mod x^n + 1$ Private basis: $\mathbf{B} := \begin{bmatrix} g & -f \\ G & -F \end{bmatrix} \Longrightarrow \mathbf{b_0} = (g_0, \dots, g_{n-1}, -f_0, \dots, -f_{n-1})$

Sign(m, B):

- 1. $r \leftarrow random salt$
- 2. $\mathbf{c} \leftarrow \text{HashToPoint}(r||m)$
- 3. $\mathbf{t} \leftarrow \mathbf{c} \cdot \mathbf{B}^{-1}$
- 4. $v \leftarrow \texttt{ffSampling}(t, B)$

5.
$$\mathbf{s} \leftarrow (\mathbf{t} - \mathbf{v}) \cdot \mathbf{B}$$

6. return (*r*, **s**)

- preimage computation
- trapdoor sampler

> Instantiation of GPV framework with NTRU lattices > Let $\mathcal{R} := \mathbb{Z}[x]/(x^n + 1)$. Private key: $f, g, F, G \in \mathcal{R}$ with $fG - gF = q \mod x^n + 1$ Private basis: $\mathbf{B} := \begin{bmatrix} g & -f \\ G & -F \end{bmatrix} \Longrightarrow \mathbf{b_0} = (g_0, \dots, g_{n-1}, -f_0, \dots, -f_{n-1})$



► < Ξ ► < Ξ ► Ξ Ξ </p>

Power Analysis on the preimage computation

Sign(m, B):

- 2. $\mathbf{c} \leftarrow \text{HashToPoint}(r||m)$
- 3. $\mathbf{t} \leftarrow \mathbf{c} \cdot \mathbf{B}^{-1}$
- 4. $v \leftarrow \texttt{ffSampling}(t, B)$
- 5. $\mathbf{s} \leftarrow (\mathbf{t} \mathbf{v}) \cdot \mathbf{B}$

- preimage computation
- trapdoor sampler

Original attack: DPA on a polynomial multiplication in FFT between a public digest c and a private polynomial f [KA21].

• • = • • = •

Original attack: DPA on a polynomial multiplication in FFT between a public digest c and a private polynomial f [KA21].

Three improvements:

Lowering the complexity of exhaustive search: double precision is unnecessary to recover the key.

Original attack: DPA on a polynomial multiplication in FFT between a public digest c and a private polynomial f [KA21].

Three improvements:

Lowering the complexity of exhaustive search:

double precision is unnecessary to recover the key.

Halving the number of required traces by combinating patterns: complex multiplications involve a lot of operations.

Original attack: DPA on a polynomial multiplication in FFT between a public digest c and a private polynomial f [KA21].

Three improvements:

Lowering the complexity of exhaustive search:

double precision is unnecessary to recover the key.

Halving the number of required traces by combinating patterns: complex multiplications involve a lot of operations.

> Mitigating the noise by grouping similar challenges:

we average power traces if challenges are the same (less precision).

Original attack: DPA on a polynomial multiplication in FFT between a public digest c and a private polynomial f [KA21].

Three improvements:

Lowering the complexity of exhaustive search:

double precision is unnecessary to recover the key.

> Halving the number of required traces by combinating patterns: complex multiplications involve a lot of operations.

> Mitigating the noise by grouping similar challenges:

we average power traces if challenges are the same (less precision).



Hidden Parallelepiped attack on the trapdoor sampler

Sign(m, B): 2. $\mathbf{c} \leftarrow \text{HashToPoint}(r||m)$ 3. $\mathbf{t} \leftarrow \mathbf{c} \cdot \mathbf{B}^{-1}$ 4. $\mathbf{v} \leftarrow \text{ffSampling}(\mathbf{t}, \mathbf{B})$ 5. $\mathbf{s} \leftarrow (\mathbf{t} - \mathbf{v}) \cdot \mathbf{B}$

- preimage computation
- trapdoor sampler

- 1. Side-channel analysis on the BaseSampler to recover samples
- 2. Utilisation of the samples to disclose a *deformed* parallelepiped
- 3. Application of HPP solver on filtered signatures
- 4. Private key recovering (possibly with lattice magic)





BaseSampler():

- 1. $u \leftarrow \text{UniformBits}(72)$
- 2. $z^+ \leftarrow 0$
- 3. for $i = 0 \dots 16$ do
- 4. $z^+ \leftarrow z^+ + \llbracket u < \mathsf{RCDT}[i] \rrbracket$
- 5. end
- 6. return z^+



BaseSampler():

1.
$$u \leftarrow \texttt{UniformBits}(72)$$

2.
$$z^+ \leftarrow 0$$

3. for
$$i = 0 \dots 16$$
 do

4.
$$z^+ \leftarrow z^+ + \llbracket u < \mathsf{RCDT}[i] \rrbracket$$

```
6. return z^+
```

Comparison on line 4 is in fact three successive substractions of 24 bits values, exploiting register underflow.



BaseSampler():

1.
$$u \leftarrow \texttt{UniformBits}(72)$$

2.
$$z^+ \leftarrow 0$$

3. for
$$i = 0 \dots 16$$
 do

4.
$$z^+ \leftarrow z^+ + \llbracket u < \mathsf{RCDT}[i] \rrbracket$$

Comparison on line 4 is in fact three successive substractions of 24 bits values, exploiting register underflow.

 \implies High difference in Hamming weight [KH18]



BaseSampler():

1.
$$u \leftarrow \texttt{UniformBits}(72)$$

2.
$$z^+ \leftarrow 0$$

3. for
$$i = 0 \dots 16$$
 do

4.
$$z^+ \leftarrow z^+ + \llbracket u < \mathsf{RCDT}[i] \rrbracket$$

Comparison on line 4 is in fact three successive substractions of 24 bits values, exploiting register underflow. \implies High difference in Hamming weight [KH18]

We are able to retrieve the value of z^+ through STA.

"Shifted" Babai's nearest-plane algorithm: A Gaussian translation vector \vec{z} is applied to t before solving CVP.



"Shifted" Babai's nearest-plane algorithm: A Gaussian translation vector \vec{z} is applied to t before solving CVP.



What happens when $\vec{z} := (z_0^+, \dots, z_{n-1}^+) = (0, \dots, 0)$?

"Shifted" Babai's nearest-plane algorithm: A Gaussian translation vector \vec{z} is applied to t before solving CVP.



What happens when $\vec{z} := (z_0^+, \dots, z_{n-1}^+) = (0, \dots, 0)$?



Filtering with all $z_i^+ = 0$

"Shifted" Babai's nearest-plane algorithm: A Gaussian translation vector \vec{z} is applied to t before solving CVP.



What happens when $\vec{z} := (z_0^+, \dots, z_{n-1}^+) = (0, \dots, 0)$?



 $\mathbb{P}[z_i^+=0] \approx \operatorname{erf}(\frac{\sqrt{2}}{2\sigma_i}) \in [0.4111, 0.5613]$ for all $i \in [0, n-1]$ (because of rejection sampling, not all z_i^+ are kept)

 $\mathbb{P}[z_i^+=0] \approx \operatorname{erf}(\frac{\sqrt{2}}{2\sigma_i}) \in [0.4111, 0.5613] \text{ for all } i \in [0, n-1]$ (because of rejection sampling, not all z_i^+ are kept)

 \Longrightarrow we apply HPP solver from [DN12] on 40 to 55% of the signatures

 $\mathbb{P}[z_i^+=0] \approx \operatorname{erf}(\frac{\sqrt{2}}{2\sigma_i}) \in [0.4111, 0.5613] \text{ for all } i \in [0, n-1]$ (because of rejection sampling, not all z_i^+ are kept)

 \Longrightarrow we apply HPP solver from [DN12] on 40 to 55% of the signatures

Baibai's nearest-plane algorithm: uses $\tilde{\mathbf{B}} := \text{GSO}(\mathbf{B})$ instead of \mathbf{B} to solve CVP.

 $\mathbb{P}[z_i^+=0] \approx \operatorname{erf}(\frac{\sqrt{2}}{2\sigma_i}) \in [0.4111, 0.5613] \text{ for all } i \in [0, n-1]$ (because of rejection sampling, not all z_i^+ are kept)

 \Longrightarrow we apply HPP solver from [DN12] on 40 to 55% of the signatures

Baibai's nearest-plane algorithm: uses $\tilde{\mathbf{B}} := \text{GSO}(\mathbf{B})$ instead of \mathbf{B} to solve CVP.

 \implies signatures belong to $\mathcal{P}(\mathbf{\tilde{B}})$, not $\mathcal{P}(\mathbf{B})$.

 $\mathbb{P}[z_i^+=0] \approx \operatorname{erf}(\frac{\sqrt{2}}{2\sigma_i}) \in [0.4111, 0.5613] \text{ for all } i \in [0, n-1]$ (because of rejection sampling, not all z_i^+ are kept)

 \Longrightarrow we apply HPP solver from [DN12] on 40 to 55% of the signatures

Baibai's nearest-plane algorithm: uses $\tilde{\mathbf{B}} := \text{GSO}(\mathbf{B})$ instead of \mathbf{B} to solve CVP.

- \implies signatures belong to $\mathcal{P}(\mathbf{\tilde{B}})$, not $\mathcal{P}(\mathbf{B})$.
- \implies we can only retrieve rows of $\tilde{\mathbf{B}}$ with HPP solver.

 $\mathbb{P}[z_i^+=0] \approx \operatorname{erf}(\frac{\sqrt{2}}{2\sigma_i}) \in [0.4111, 0.5613] \text{ for all } i \in [0, n-1]$ (because of rejection sampling, not all z_i^+ are kept)

 \Longrightarrow we apply HPP solver from [DN12] on 40 to 55% of the signatures

Baibai's nearest-plane algorithm: uses $\tilde{\mathbf{B}} := \text{GSO}(\mathbf{B})$ instead of \mathbf{B} to solve CVP.

- \implies signatures belong to $\mathcal{P}(\mathbf{\tilde{B}})$, not $\mathcal{P}(\mathbf{B})$.
- \implies we can only retrieve rows of $\tilde{\mathbf{B}}$ with HPP solver.

Useful observation: Because of the algorithm used in Falcon to compute the GSO (ffLDL algorithm), we have the following: $\tilde{\mathbf{b}}_0, \ldots, \tilde{\mathbf{b}}_3 \approx \mathbf{b}_0, \ldots, \mathbf{b}_3$ and $\tilde{\mathbf{b}}_n, \ldots, \tilde{\mathbf{b}}_{n+3} \approx \mathbf{b}_n, \ldots, \mathbf{b}_{n+3}$

레이 이 리아 이 리아 드레이

We combine several rows b_i to attenuate the noise on f, g. Note: ulterior iterations of HPP solver are much less costly.

We combine several rows b_i to attenuate the noise on f, g. Note: ulterior iterations of HPP solver are much less costly.

Then, two possible ways to recover the exact private key (f, g): Mere rounding when $\sigma_{(f',g')-(f,g)}$ is small enough

We combine several rows b_i to attenuate the noise on f, g. Note: ulterior iterations of HPP solver are much less costly.

Then, two possible ways to recover the exact private key (f, g): > Mere rounding when $\sigma_{(f',g')-(f,g)}$ is small enough > Solve DBDD instance with Leaky LWE/NTRU tool [Dac+20]

We combine several rows b_i to attenuate the noise on f, g. Note: ulterior iterations of HPP solver are much less costly.

Then, two possible ways to recover the exact private key (f, g):
Mere rounding when σ_{(f',g')−(f,g)} is small enough
Solve DBDD instance with Leaky LWE/NTRU tool [Dac+20]



Sign(m, B): 2. $\mathbf{c} \leftarrow \text{HashToPoint}(r||m)$ 3. $\mathbf{t} \leftarrow \mathbf{c} \cdot \mathbf{B}^{-1}$ 4. $\mathbf{v} \leftarrow \text{ffSampling}(\mathbf{t}, \mathbf{B})$ 5. $\mathbf{s} \leftarrow (\mathbf{t} - \mathbf{v}) \cdot \mathbf{B}$

- preimage computation
- trapdoor sampler

Preimage computation: Improvement of State-of-the-Art attack. Trapdoor sampler: Novel attack combining SCA and HPP.

Sign(m, B): 2. $\mathbf{c} \leftarrow \text{HashToPoint}(r||m)$ 3. $\mathbf{t} \leftarrow \mathbf{c} \cdot \mathbf{B}^{-1}$ 4. $\mathbf{v} \leftarrow \text{ffSampling}(\mathbf{t}, \mathbf{B})$ 5. $\mathbf{s} \leftarrow (\mathbf{t} - \mathbf{v}) \cdot \mathbf{B}$

- preimage computation
- trapdoor sampler

Preimage computation: Improvement of State-of-the-Art attack. Trapdoor sampler: Novel attack combining SCA and HPP.

Future works:

- Template attack on the SamplerZ
- Combination with [Fou+20] (replacing timing attack by STA)

Sign(m, B): 2. $\mathbf{c} \leftarrow \text{HashToPoint}(r||m)$ 3. $\mathbf{t} \leftarrow \mathbf{c} \cdot \mathbf{B}^{-1}$ 4. $\mathbf{v} \leftarrow \text{ffSampling}(\mathbf{t}, \mathbf{B})$ 5. $\mathbf{s} \leftarrow (\mathbf{t} - \mathbf{v}) \cdot \mathbf{B}$

- preimage computation
- trapdoor sampler

Preimage computation: Improvement of State-of-the-Art attack. Trapdoor sampler: Novel attack combining SCA and HPP.

Future works:

- Template attack on the SamplerZ
- Combination with [Fou+20] (replacing timing attack by STA)

Questions ?

References I

- [Alb+19] Martin R. Albrecht et al. "The General Sieve Kernel and New Records in Lattice Reduction". In: EUROCRYPT 2019, Part II. Ed. by Yuval Ishai and Vincent Rijmen. Vol. 11477. LNCS. Springer, Heidelberg, May 2019, pp. 717–746.
- [Dac+20] Dana Dachman-Soled et al. "LWE with Side Information: Attacks and Concrete Security Estimation". In: CRYPTO 2020, Part II. Ed. by Daniele Micciancio and Thomas Ristenpart. Vol. 12171. LNCS. Springer, Heidelberg, Aug. 2020, pp. 329–358.

References II

[DN12] Léo Ducas and Phong Q. Nguyen. "Learning a Zonotope and More: Cryptanalysis of NTRUSign Countermeasures". In: ASIACRYPT 2012. Ed. by Xiaoyun Wang and Kazue Sako. Vol. 7658. LNCS. Springer, Heidelberg, Dec. 2012, pp. 433–450.

[Fou+20] Pierre-Alain Fouque et al. "Key Recovery from Gram-Schmidt Norm Leakage in Hash-and-Sign Signatures over NTRU Lattices". In: EUROCRYPT 2020, Part III. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12107. LNCS. Springer, Heidelberg, May 2020, pp. 34–63.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. "Trapdoors for hard lattices and new cryptographic constructions". In: ACM STOC 40 (2008), pp. 197–206.

References III

- [KA21] Emre Karabulut and Aydin Aysu. "FALCON Down: Breaking FALCON Post-Quantum Signature Scheme through Side-Channel Attacks". In: 2021 58th ACM/IEEE Design Automation Conference (DAC). 2021, pp. 691–696.
- [KH18] Suhri Kim and Seokhie Hong. "Single Trace Analysis on Constant Time CDT Sampler and Its Countermeasure". In: Applied Sciences 8 (Oct. 2018), p. 1809.
- [NR06] Phong Q. Nguyen and Oded Regev. "Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures". In: *EUROCRYPT 2006* 4004 (2006), pp. 271–288.

Partial countermeasure for BaseSampler

Main idea: invert the sign of the operands to replace the (hardware) underflow by a (logical) overflow.

Replace the last substraction by the following:

- 1. $b \leftarrow \texttt{Oxfffff}$
- 2. $b := b \overline{u} + \overline{\mathsf{RCDT}[i]} + c$
- 3. return $b \gg 24$

State of the register before the last operation:



State of the register after the last operation (original implementation):



State of the register after the last operation (with countermeasure):

