

Plover

Muhammed F. Esgin
Monash University

Thomas Espitau
PQShield

Guilhem Niot
PQShield
University Rennes 1






Thomas Prest¹
PQShield

Amid Sakzad
Monash University

Ron Steinfeld
Monash University

¹Thanks to Thomas for letting me reuse some of his slides on Raccoon.

Signature schemes strike a balance between:

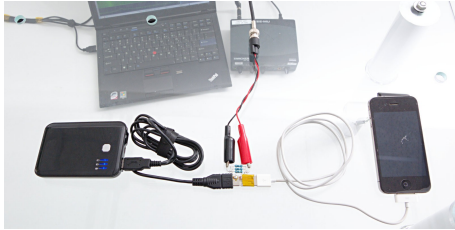
-  Sizes (verification key and signatures)
-  Speed (signing, verification)
-  Portability
-  Conservative assumptions
-  **Resistance against side-channel attacks**

And so on...

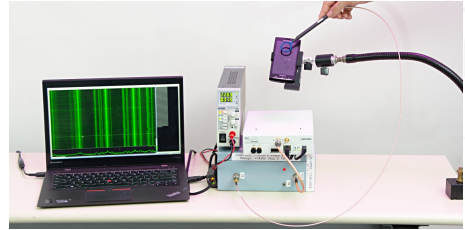
Criteria					
Dilithium	★★★	★★★★	★★★★	★★	🛡️
Falcon	★★★★	★★★★	★★	★★	🛡️
SPHINCS+	★★	★★	★★	★★★★	🛡️
Raccoon	★★	★★★★	★★★★	★★	★★★★
Plover	★★	★★★★	★★★★	★★	★★★★

Side-Channel Attacks

Power consumption [KJJ99]



Electromagnetic emissions [Eck85]



Timing measurement [Koc96]



Acoustic emissions [AA04]



In Falcon, a signature **sig** is distributed as a Gaussian.

The signing key **sk** should remain private.

The power consumption leaks information about the dot product $\langle \mathbf{sig}, \mathbf{sk} \rangle$, or **sk** itself.

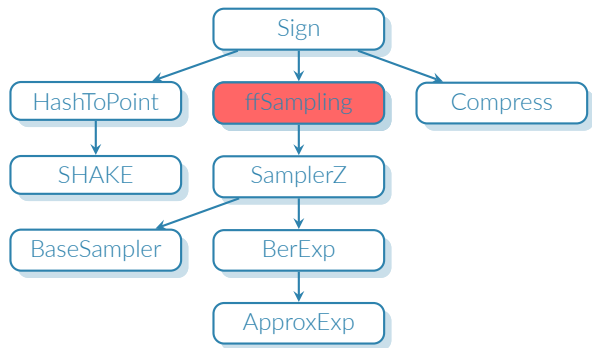
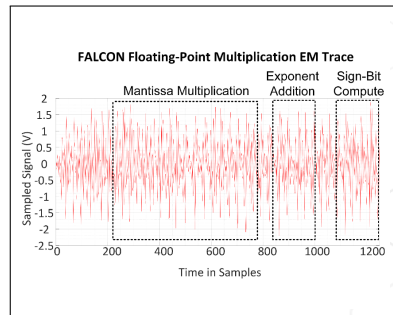


Figure 1: Flowchart of the signature



Learning **sk** directly

In Falcon, a signature **sig** is distributed as a Gaussian.

The signing key **sk** should remain private.

The power consumption leaks information about the dot product $\langle \mathbf{sig}, \mathbf{sk} \rangle$, or **sk** itself.

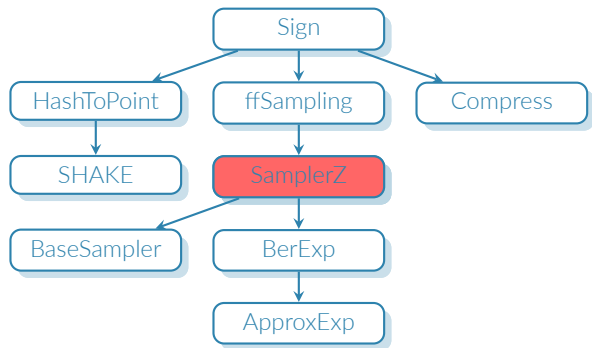
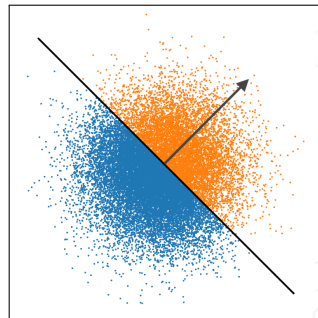




Figure 1: Flowchart of the signature



Filtering $\langle \mathbf{sig}, \mathbf{sk} \rangle > 0$

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:

$$\llbracket x \rrbracket = (x_0, x_1, \dots, x_{d-1}) \quad (1)$$

such that

$$x_0 + x_1 + \dots + x_{d-1} = x \pmod q \quad (\text{additive})$$

$$\text{or } x_0 \oplus x_1 \oplus \dots \oplus x_{d-1} = x \quad (\text{boolean})$$

-  In t -probing model, ideally 0 leakage if $d > t$
-  In “real life”, security is exponential in d
-  What about computations?



How difficult are operations to mask?

😊 Addition ($\llbracket c \rrbracket = \llbracket a + b \rrbracket$)?

- Compute $\llbracket c \rrbracket = (a_0 + b_0, \dots, a_{d-1} + b_{d-1})$, simple and fast: $\Theta(d)$ operations

😞 Multiplication ($\llbracket c \rrbracket = \llbracket a \cdot b \rrbracket$)?

- Complex and slower: $\Theta(d^2)$ operations

🤖 More complex operations?

- Use so-called *mask conversions* to convert between additive and boolean masking, very slow: $\gg \Theta(d^2)$ operations



Masking Dilithium

Dilithium follows the Fiat-Shamir **with aborts** paradigm.

$\text{Sign}(sk = \mathbf{s}, vk = (\mathbf{A}, \mathbf{t}), \text{msg}) \rightarrow \text{sig}$

- 1 Generate a short ephemeral secret \mathbf{r}
- 2 Compute the commitment $\mathbf{w} = \mathbf{A} \cdot \mathbf{r}$
- 3 Compute challenge $c = \text{H}(\mathbf{w}, \text{msg}, vk)$
- 4 Compute the response $\mathbf{z} = \mathbf{s} \cdot c + \mathbf{r}$
- 5 Check that \mathbf{z} is in a given interval. If not, restart.
- 6 Signature is $\text{sig} = (c, \mathbf{z})$

$\text{Verify}(vk, \text{msg}, \text{sig} = (c, \mathbf{z}))$

- 1 Verify that \mathbf{z} is small.
- 2 Recover $\mathbf{w} = \mathbf{A} \cdot \mathbf{z} - c \cdot \mathbf{t}$
- 3 Verify that $c = \text{H}(\mathbf{w}, \text{msg}, vk)$

Dilithium follows the Fiat-Shamir **with aborts** paradigm.

$\text{Sign}(sk = \mathbf{s}, vk = (\mathbf{A}, t), \text{msg}) \rightarrow \text{sig}$

- 1 Generate a short ephemeral secret \mathbf{r} ▷ Hard
- 2 Compute the commitment $\mathbf{w} = \mathbf{A} \cdot \mathbf{r}$ ▷ Easy
- 3 Compute challenge $c = H(\mathbf{w}, \text{msg}, vk)$ ▷ No mask
- 4 Compute the response $\mathbf{z} = \mathbf{s} \cdot c + \mathbf{r}$ ▷ Easy
- 5 Check that \mathbf{z} is in a given interval. If not, restart. ▷ Hard
- 6 Signature is $\text{sig} = (c, \mathbf{z})$

Masking bottlenecks:

- ⊗ Short secret generation (1) requires B2A.
- ⊗ Rejection sampling (5) requires A2B and B2A.


Total masking overhead: $\Theta(d^2 \log q)$

Sign($sk = \llbracket \mathbf{s} \rrbracket, vk = (\mathbf{A}, \mathbf{t}), msg$) \rightarrow sig


- 1 Generate a masked short ephemeral secret $\llbracket \mathbf{r} \rrbracket$ using “AddRepNoise” ▷ Easy
- 2 Compute the commitment $\llbracket \mathbf{w} \rrbracket = \mathbf{A} \cdot \llbracket \mathbf{r} \rrbracket$ ▷ Easy
- 3 Unmask $\llbracket \mathbf{w} \rrbracket$ to obtain \mathbf{w} ▷ Easy
- 4 Compute the challenge $c = H(\mathbf{w}, msg, vk)$ ▷ No mask
- 5 Compute the response $\llbracket \mathbf{z} \rrbracket = \llbracket \mathbf{s} \rrbracket \cdot c + \llbracket \mathbf{r} \rrbracket$ ▷ Easy
- 6 Unmask $\llbracket \mathbf{z} \rrbracket$ to obtain \mathbf{z} ▷ Easy
- 7 (No more rejection sampling!)
- 8 Signature is $\mathbf{sig} = (c, \mathbf{z})$

Total masking overhead: $O(d \log d)$


What happens inside AddRepNoise?



What happens inside AddRepNoise?



Problem: a probing adversary can learn the sum of T random in 2 probes.



Solution: add refresh gadgets to separate the algorithm in independent layers
Now a probing adversary learns at most (the sum of) t short noises.

- Vanilla Raccoon,
 - Output of **AddRepNoise** looks like a gaussian.
 - No rejection sampling: signatures leak part of the secret

- Vanilla Raccoon,
 - Output of **AddRepNoise** looks like a gaussian.
 - No rejection sampling: signatures leak part of the secret

Definition 1 (Hint-MLWE)

It is hard to distinguish $(\mathbf{A}, \mathbf{u}, (c_i \cdot \mathbf{s} + \mathbf{r}_i)_i)$ with (c_i) small,

- when \mathbf{u} is random
- or, when $\mathbf{u} = \mathbf{A} \cdot \mathbf{s}$ an MLWE sample

Assuming at most Q hints, Hint-MLWE is as hard as MLWE when taking \mathbf{r}_i of standard deviation $\approx \sqrt{Q} \|c\|$.


😊 Vanilla Raccoon is secure when taking large enough perturbations \mathbf{r}_i .

- Vanilla Raccoon,
 - Output of **AddRepNoise** looks like a gaussian.
 - No rejection sampling: signatures leak part of the secret
- Masked Raccoon:
 - Leak part of the perturbation $\mathbf{r} = \text{AddRepNoise}()$.

In t -probing model, write $\mathbf{r} = \mathbf{r}_{\text{safe}} + \mathbf{r}_{\text{leaked}}$.

If **rep** iterations in **AddRepNoise**, \mathbf{r}_{safe} has standard deviation $\sqrt{d \cdot \text{rep} - t} \cdot \sigma_{\mathbf{r}}$.

Security of **Masked Raccoon reduces to Vanilla Raccoon** with small loss.



💡 With some tricks [SR23], RAM consumption is < 128 kB

Masked Hash-and-Sign signatures

Keygen(1^λ)

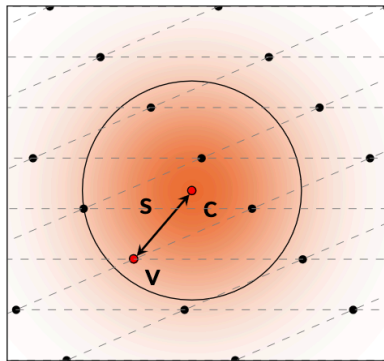
- 1 Gen. matrices \mathbf{A} , \mathbf{B} s.t.:
 - > \mathbf{A} is pseudo-random.
 - > $\mathbf{B} \cdot \mathbf{A} = \mathbf{0}$.
 - > \mathbf{B} has small coefficients.
- 2 $\text{vk} := \mathbf{A}$, $\text{sk} := \mathbf{B}$

Sign($\text{sk} = \mathbf{B}$, msg)

- 1 Compute \mathbf{c} such that $\mathbf{c} \cdot \mathbf{A} = H(\text{msg})$
- 2 $\mathbf{v} \leftarrow$ vector in $\mathcal{L}(\mathbf{B})$, close to \mathbf{c} .
- 3 $\text{sig} := \mathbf{s} = (\mathbf{c} - \mathbf{v})$

Verify($\text{vk} = \mathbf{A}$, msg , $\text{sig} = \mathbf{s}$)

Check that (\mathbf{s} is short) and ($\mathbf{s} \cdot \mathbf{A} = H(\text{msg})$)



⊗ But masking Gaussian sampling efficiently remains an open problem.

- In 2022, *Mitaka: a simpler, parallelizable, maskable variant of Falcon* [EFG+22]
 - But, *A Key-Recovery Attack against Mitaka in the t -Probing Model* [Pre23]

Mitaka cannot be masked efficiently with current techniques.

Eagle was recently introduced by Yu et al. in [YJW23].

Eagle.Keygen()

- 1 Generate matrices \mathbf{A}, \mathbf{T} s.t.:
 - > \mathbf{A} is pseudo-random
 - > $\mathbf{T} \cdot \mathbf{A} = \beta \cdot \mathbf{I}$
 - > \mathbf{T} has small coefficients
- 2 $\text{vk} := \mathbf{A}, \text{sk} := \mathbf{T}$

Eagle.Verify(msg, sig = \mathbf{z})

- 1 $u := H(\text{msg})$
- 2 Check that (\mathbf{z} is small) and $(\mathbf{A} \cdot \mathbf{z} = u)$

Eagle.Sign(sk, msg)

- 1 $\mathbf{u} := H(\text{msg})$
- 2 $\mathbf{p} \leftarrow D_{\mathcal{R}^\ell, \sqrt{s^2 \mathbf{I} - r^2 \mathbf{T} \mathbf{T}^*}}$
- 3 $\mathbf{c} := \mathbf{u} - \mathbf{A} \cdot \mathbf{p}$
- 4 Decompose \mathbf{c} as $\mathbf{c} = \beta \cdot \mathbf{c}_1 + \mathbf{c}_2$
- 5 $\mathbf{y} \leftarrow D_{\lfloor q/\beta \rfloor \cdot \mathcal{R}^\ell + \mathbf{c}_1, r}$
- 6 $\mathbf{z} := \mathbf{p} + \mathbf{T} \cdot \mathbf{y}$
- 7 return sig := \mathbf{z}

Eagle was recently introduced by Yu et al. in [YJW23].

Eagle.Sign(sk, msg)

- 1 $\mathbf{u} := H(\text{msg})$
- 2 $\mathbf{p} \leftarrow D_{\mathcal{R}^\ell, \sqrt{s^2 I - r^2 \mathbf{T}\mathbf{T}^*}}$
- 3 $\mathbf{w} := \mathbf{A} \cdot \mathbf{p}$
- 4 $\mathbf{c} := \mathbf{u} - \mathbf{w}$
- 5 Decompose \mathbf{c} as $\mathbf{c} = \beta \cdot \mathbf{c}_1 + \mathbf{c}_2$
- 6 $\mathbf{y} \leftarrow D_{\lfloor q/\beta \rfloor \cdot \mathcal{R}^\ell + \mathbf{c}_1, r}$
- 7 $\mathbf{z} := \mathbf{p} + \mathbf{T} \cdot \mathbf{y}$
- 8 return sig := \mathbf{z}

😊 Almost linear scheme, maybe we can do something with it!

Eagle was recently introduced by Yu et al. in [YJW23].

Eagle.Sign(sk, msg)

- 1 $\mathbf{u} := H(\text{msg})$ ▷ No mask
- 2 $\mathbf{p} \leftarrow D_{\mathcal{R}^\ell, \sqrt{s^2\mathbf{I} - r^2\mathbf{T}\mathbf{T}^*}}$ ▷ **Hard**
- 3 $\mathbf{w} := \mathbf{A} \cdot \mathbf{p}$ ▷ **Easy**
- 4 $\mathbf{c} := \mathbf{u} - \mathbf{w}$ ▷ No mask
- 5 Decompose \mathbf{c} as $\mathbf{c} = \beta \cdot \mathbf{c}_1 + \mathbf{c}_2$
- 6 $\mathbf{y} \leftarrow D_{\lfloor q/\beta \rfloor \cdot \mathcal{R}^\ell + \mathbf{c}_1, r}$ ▷ **Hard**
- 7 $\mathbf{z} := \mathbf{p} + \mathbf{T} \cdot \mathbf{y}$ ▷ **Easy**
- 8 return sig := \mathbf{z}

😊 Almost linear scheme, maybe we can do something with it!

Eagle was recently introduced by Yu et al. in [YJW23].

Eagle.Sign(sk, msg)

- 1 $\mathbf{u} := H(\text{msg})$ ▷ No mask
- 2 $\mathbf{p} \leftarrow D_{\mathcal{R}^\ell, \sqrt{s^2\mathbf{I} - r^2\mathbf{T}\mathbf{T}^*}}$ ▷ **Hard**
- 3 $\mathbf{w} := \mathbf{A} \cdot \mathbf{p}$ ▷ **Easy**
- 4 $\mathbf{c} := \mathbf{u} - \mathbf{w}$ ▷ No mask
- 5 Decompose \mathbf{c} as $\mathbf{c} = \beta \cdot \mathbf{c}_1 + \mathbf{c}_2$
- 6 $\mathbf{y} \leftarrow D_{\lfloor q/\beta \rfloor \cdot \mathcal{R}^\ell + \mathbf{c}_1, r}$ ▷ **Hard**
- 7 $\mathbf{z} := \mathbf{p} + \mathbf{T} \cdot \mathbf{y}$ ▷ **Easy**
- 8 return sig := \mathbf{z}

Plover.Sign(sk, msg)

- 1 $\mathbf{u} := H(\text{msg})$
- 2 $\llbracket \mathbf{p} \rrbracket \leftarrow \text{AddRepNoise}()$
- 3 $\mathbf{w} := \text{Unmask}(\mathbf{A} \cdot \llbracket \mathbf{p} \rrbracket)$
- 4 $\mathbf{c} := \mathbf{u} - \mathbf{w}$
- 5 Decompose \mathbf{c} as $\mathbf{c} = \beta \cdot \mathbf{c}_1 + \mathbf{c}_2$
- 6 $\mathbf{z} := \text{Unmask}(\llbracket \mathbf{p} \rrbracket + \llbracket \mathbf{T} \rrbracket \cdot \mathbf{c}_1)$
- 7 return sig := \mathbf{z}

😊 Almost linear scheme, maybe we can do something with it!

😄 Introducing Plover, the first hash-and-sign masking-friendly signature scheme.

→ Vanilla Plover

➤ Returns responses of the form $\mathbf{z} = \mathbf{p} + \mathbf{T} \cdot \mathbf{c}_1$: hints on the secret.

😊 Like Raccoon, rely on Hint-MLWE. Secure for large enough perturbation \mathbf{p} .

- Vanilla Plover
 - > Returns responses of the form $\mathbf{z} = \mathbf{p} + \mathbf{T} \cdot \mathbf{c}_1$: hints on the secret.
 - 😊 Like Raccoon, rely on Hint-MLWE. Secure for large enough perturbation \mathbf{p} .
- Masked Plover
 - > As in Raccoon, **AddRepNoise** leaks only a small part of the perturbation \mathbf{p} .
 - 😊 Unforgeability of Masked Plover in the t -probing model reduces to unforgeability of Vanilla Plover.

Plover introduces a very generic framework for masking friendly schemes:


- Replace non-linear operations with noise flooding. Leakage on the secret mitigated by taking large perturbations \mathbf{p} .
- Analyse leakage with Hint-MLWE problem.
- Use **AddRepNoise** to sample short vectors. New composable notion $t - \text{SN}l_u$ to prove security in the t -probing model.

Proofs in the t -probing model

t -probing model

 Adversary can probe t circuit values at runtime


Masking

 Each sensitive value x is split in d shares:

$$[[x]] = (x_0, x_1, \dots, x_{d-1})$$

such that

$$x_0 + x_1 + \dots + x_{d-1} = x \pmod{q} \quad (\text{additive})$$

 In t -probing model, ideally 0 leakage if $d > t$



t -probing model

 Adversary can probe t circuit values at runtime

Definition 1 (t -probing security)

A circuit C is t -probing secure, if there exists a simulator \mathcal{S} such that for any input \mathbf{x} , and set \mathcal{P} of up to t probes:

$$\mathcal{S}(\mathcal{P}, C_{\text{public}}(\llbracket \mathbf{x} \rrbracket)) = \underbrace{C_{\mathcal{P}}(\llbracket \mathbf{x} \rrbracket)}_{\text{Probes on } C \text{ executed with } \mathbf{x}}$$

i.e., probes are simulatable without knowledge of the circuit input \mathbf{x} , only from public output $C_{\text{public}}(\mathbf{x})$.

- The (strong) non-interference (or (S)NI) framework eases proofs in the t -probing model.
Composition of simple gadgets: masked additions, multiplications, etc.

- The (strong) non-interference (or (S)NI) framework eases proofs in the t -probing model.
- Composition of simple gadgets: masked additions, multiplications, etc.

Definition 2 (t -NI)

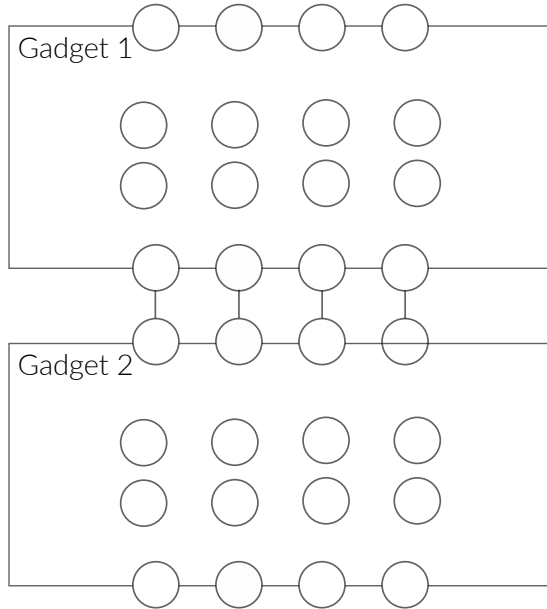
A circuit C is t -NI, if there exists simulators $\mathcal{S}_1, \mathcal{S}_2$ such that for any input $[[\mathbf{x}]]$, and any set \mathcal{P} of at most t probes:

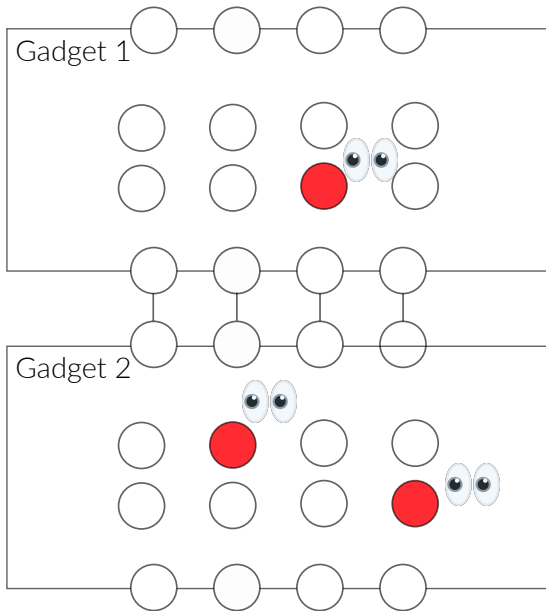
$$i_1, \dots, i_t := \mathcal{S}_1(\mathcal{P})$$
$$S_2(\mathcal{P}, [[\mathbf{x}]]_{i_1}, \dots, [[\mathbf{x}]]_{i_t}) = C_{\mathcal{P}}([[\mathbf{x}]])$$

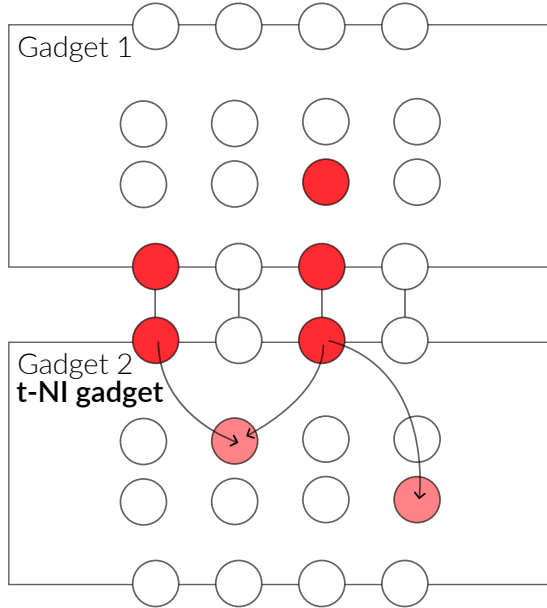
i.e. probes are simulatable from at most t shares of the input.

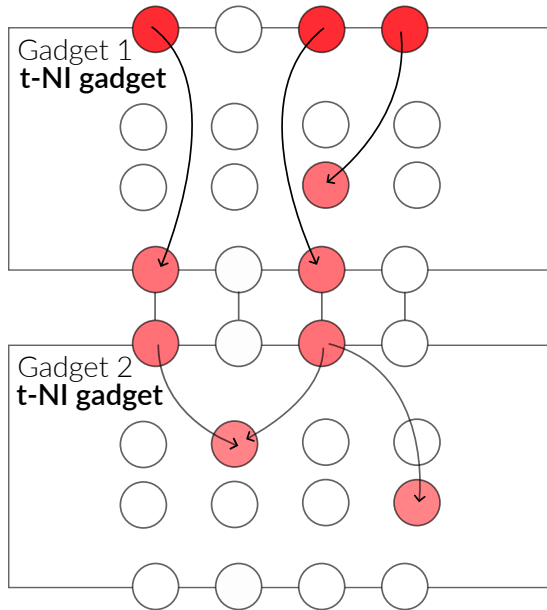
Definition 3 (t -SNI)

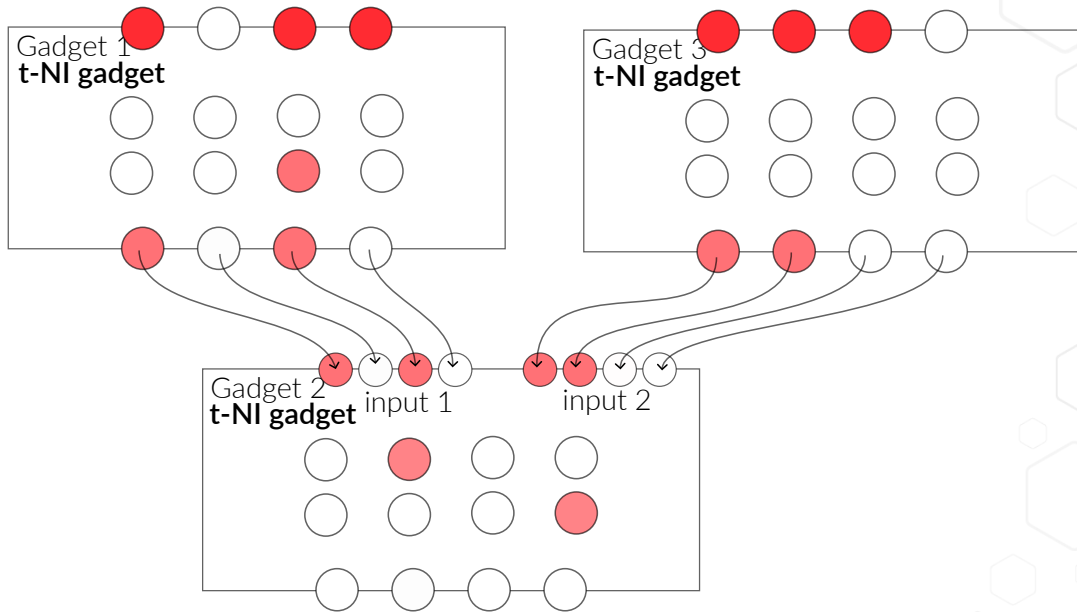
Same, but output probes are simulated from internal probes only. Formally, there exists an extra simulator S_3 for probes on output: $S_3(\mathcal{P}_{\text{out}}, S_2(\mathcal{P}_{\text{in}}, \dots)) = C_{\mathcal{P}_{\text{out}}}([[\mathbf{x}]])$.

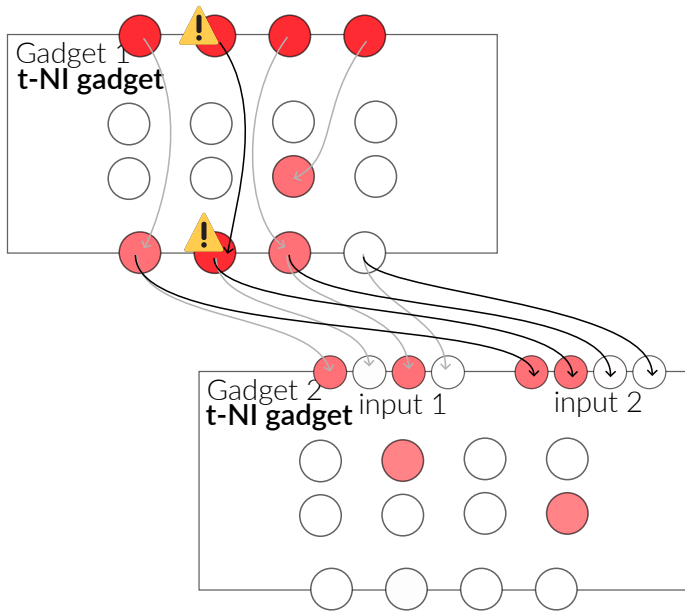


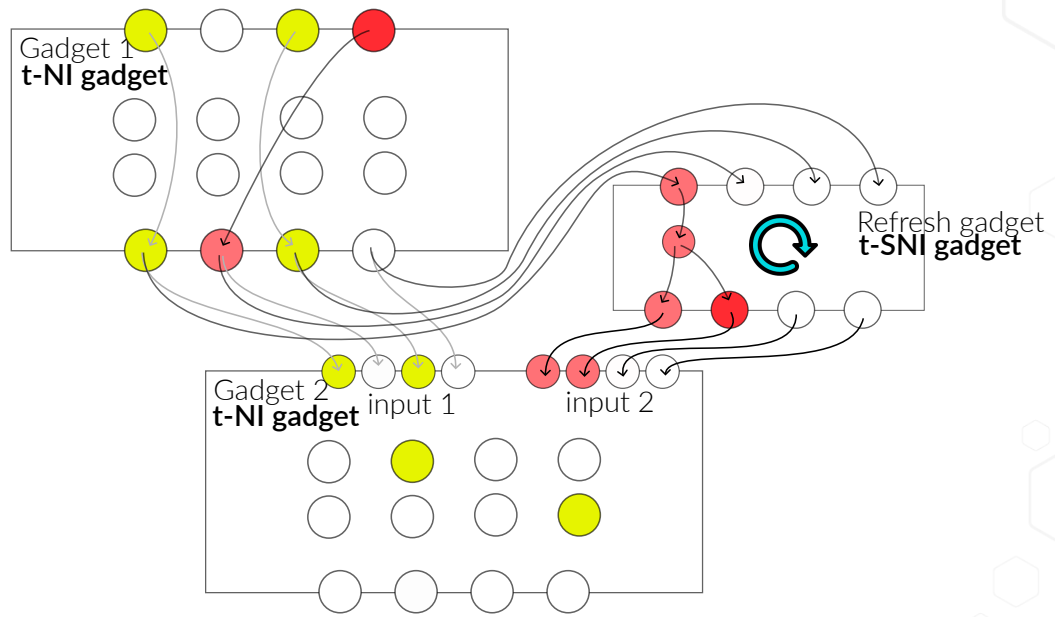












- The randomness added in **AddRepNoise** are secret inputs to the signature circuit, but some of them leak.
(S)NI model does not capture partial leakage of input.

- The randomness added in **AddRepNoise** are secret inputs to the signature circuit, but some of them leak.
(S)NI model does not capture partial leakage of input.
- New notion: t -SNlu, strong non-interference with unmasked inputs.

Definition 4 (t -SNlu)

A circuit C is t -SNlu, if there exists simulators $\mathcal{S}_1, \mathcal{S}_2$ such that for any input $[[\mathbf{x}]]$, unmasked values $(\mathbf{v}_i)_i$, and any set \mathcal{P} of at most t probes:

$$i_1, \dots, i_t, i'_1, \dots, i'_t := \mathcal{S}_1(\mathcal{P})$$
$$\mathcal{S}_2(\mathcal{P}, [[\mathbf{x}]]_{i_1}, \dots, [[\mathbf{x}]]_{i_t}, \mathbf{v}_{i'_1}, \dots, \mathbf{v}_{i'_t}) = C_{\mathcal{P}}([[\mathbf{x}]], (\mathbf{v}_i)_i)$$

i.e. probes are simulatable from at most t shares of the input \mathbf{x} , and t values (\mathbf{v}_i) .

- The randomness added in **AddRepNoise** are secret inputs to the signature circuit, but some of them leak.
(S)NI model does not capture partial leakage of input.
- New notion: t -SNlu, strong non-interference with unmasked inputs.
- We can show that **AddRepNoise** is t -SNlu secure for $t < d$.

- The randomness added in **AddRepNoise** are secret inputs to the signature circuit, but some of them leak.
(S)NI model does not capture partial leakage of input.
- New notion: t -SNlu, strong non-interference with unmasked inputs.
- We can show that **AddRepNoise** is t -SNlu secure for $t < d$.

- t -SNlu is composable: probes on Raccoon/Plover signing procedure can be simulated with at most t inputs shares, and t unmasked values.
 - t shares of masked input: independent from actual input
 - t unmasked values: remains $d \cdot \text{rep} - t$ safe values to ensure security

Raccoon and Plover are specific-purpose schemes aimed at high side-channel resistance:







- 😊 Standard assumptions: MLWE, MSIS
- 😊 Simpler
- 😊 Verification key size is similar
- 😞 Signatures are larger ($\approx 10\text{kB}$)
- 😊 **When masked, orders of magnitude faster than other schemes are**

General framework to create masking friendly schemes:

- ➔ Noise-flooding to replace non-linear operations
- ➔ Prove unmasked security with Hint-MLWE
- ➔ Sample short vectors with **AddRepNoise** and use t -SNlu notion to prove security in the t -probing model

Questions?



- 
-  Dmitri Asonov and Rakesh Agrawal.
Keyboard acoustic emanations.
pages 3–11, 2004.
 -  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.
pages 222–253, 2022.
 -  Emre Karabulut and Aydin Aysu.
FALCON down: Breaking FALCON post-quantum signature scheme through side-channel attacks.
In 58th ACM/IEEE Design Automation Conference, DAC 2021, San Francisco, CA, USA, December 5-9, 2021, pages 691–696. IEEE, 2021.
 -  Paul C. Kocher, Joshua Jaffe, and Benjamin Jun.
Differential power analysis.
pages 388–397, 1999.
- 

-  Paul C. Kocher.
Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems.
pages 104–113, 1996.
-  Thomas Prest.
A key-recovery attack against mitaka in the t -probing model.
pages 205–220, 2023.
-  Markku-Juhani O. Saarinen and Mélissa Rossi.
Mask compression: High-order masking on memory-constrained devices.
Cryptology ePrint Archive, Paper 2023/1117, 2023.
<https://eprint.iacr.org/2023/1117>.
-  Yang Yu, Huiwen Jia, and Xiaoyun Wang.
Compact lattice gadget and its applications to hash-and-sign signatures.
pages 390–420, 2023.
-  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>.

