Unifying Leakage Models on a Rényi Day

Dahmun Goudarzi[2] Ange Martinelli[3]
Alain Passelègue[1] Thomas Prest[2]

ENS DE LYON
PQSHIELD
THALES

Side-channel attacks in cryptography

Power analysis attacks [KJJ99]

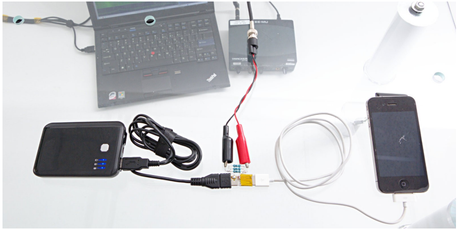

Electromagnetic attacks [Eck85, GMO01]

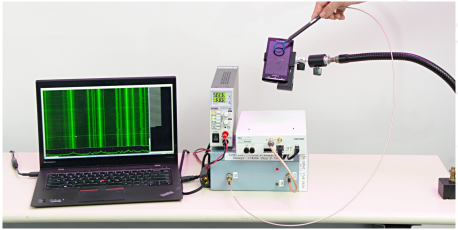

Timing attacks [Koc96, BB03]

Acoustic attacks [AA04, GST14]

How do we modelize a leakage trace?

[GPP+16]

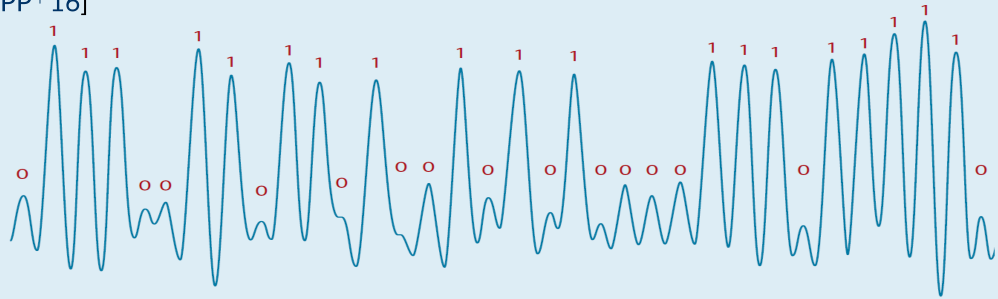

How do we modelize a leakage trace?

Each node of interest follows a distribution X. Its leakage Y is a randomized function $f(X)$.

[GPP+16]

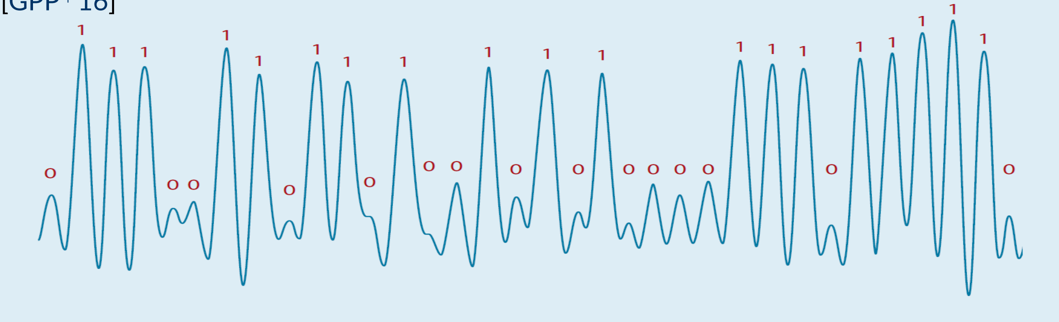

Concrete modelization of leakage

→ Popular one is “Hamming weight + Gaussian” [BCO04]: $f(X) = \mathrm{HW}(X) + \mathcal{N}(0, \sigma)$

Very realistic. Very hard to work with.

How do we modelize a leakage trace?

Each node of interest follows a distribution X. Its leakage Y is a randomized function $f(X)$.

[GPP$^+$16]

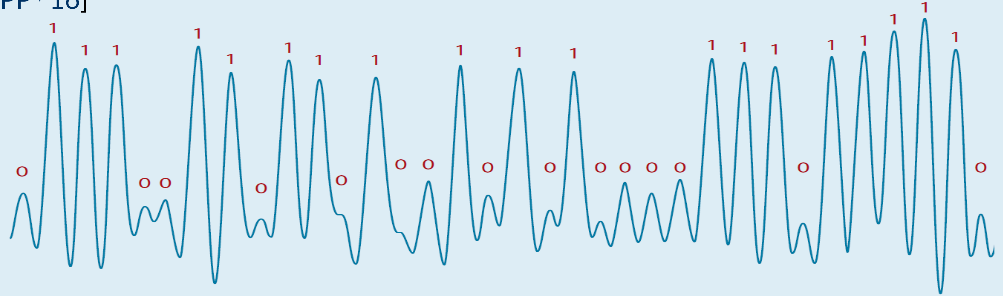

Noisy leakage models: “The leakage Y bias the expected distribution of X”

→ [PR13]: bias *metric* is $\mathrm{EN}(X|Y) = \mathbb{E}_Y \, \|X - (X|Y)\|_2$

→ [DDF14]: bias *metric* is $\mathrm{SD}(X|Y) = \frac{1}{2}\mathbb{E}_Y \, \|X - (X|Y)\|_1$

Realistic but unwieldy. Definition implicitly depends of X.

How do we modelize a leakage trace?

Each node of interest follows a distribution X. Its leakage Y is a randomized function $f(X)$.

[GPP+16]

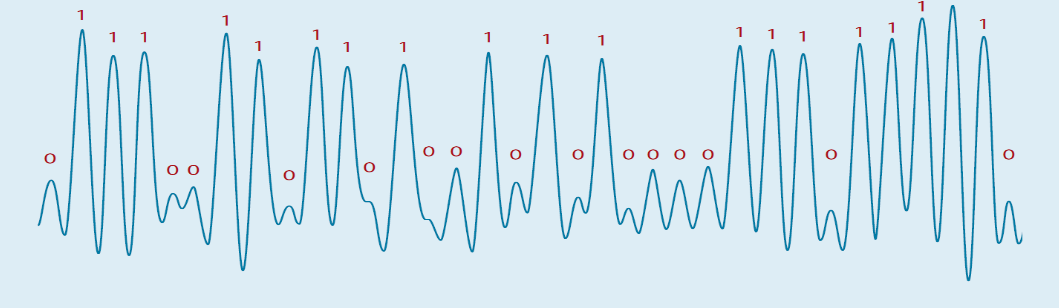

Probing models: "The adversary may know *exactly* some nodes"

- → Threshold [ISW03]: adv. chooses *exactly* t nodes to probe
- → Random [ISW03]: adv. probes each node with prob. ε

Idealized but easy to use.

The cryptographer's problem

People propose secure compilers to protect circuits.
We have circuit compilers and several shades of leakage models...

Concrete leakage modelizations

Noisy leakage models

Probing models

Circuit compilers

The cryptographer's problem

People propose secure compilers to protect circuits.
We have circuit compilers and several shades of leakage models...

Concrete leakage modelizations

↓

Noisy leakage models

↓

Probing models

↓

Circuit compilers

... and we want to show *in the most efficient way* that a circuit compiler is secure for a concrete modelization of leakage.

Previous works

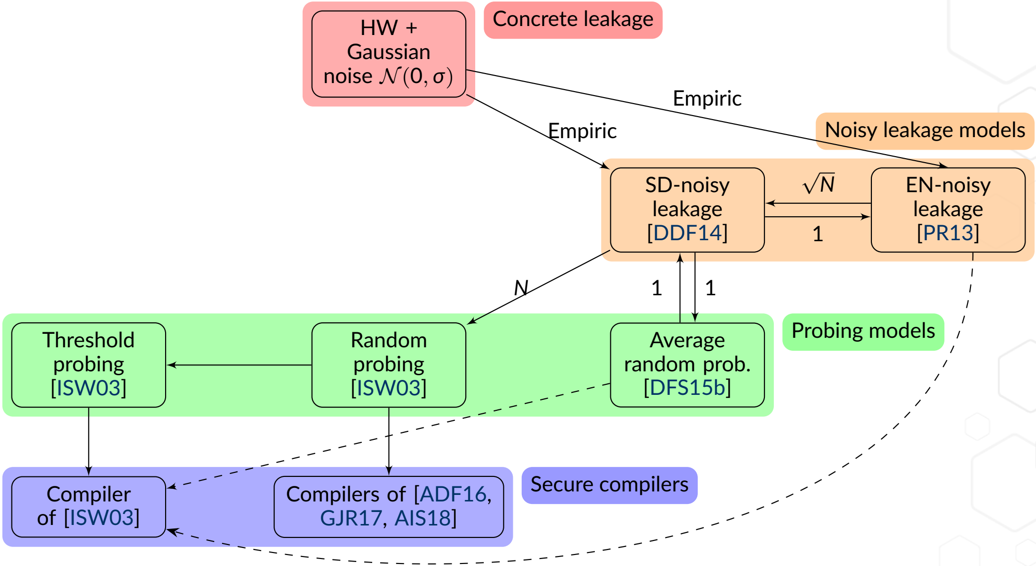

Previous and current works

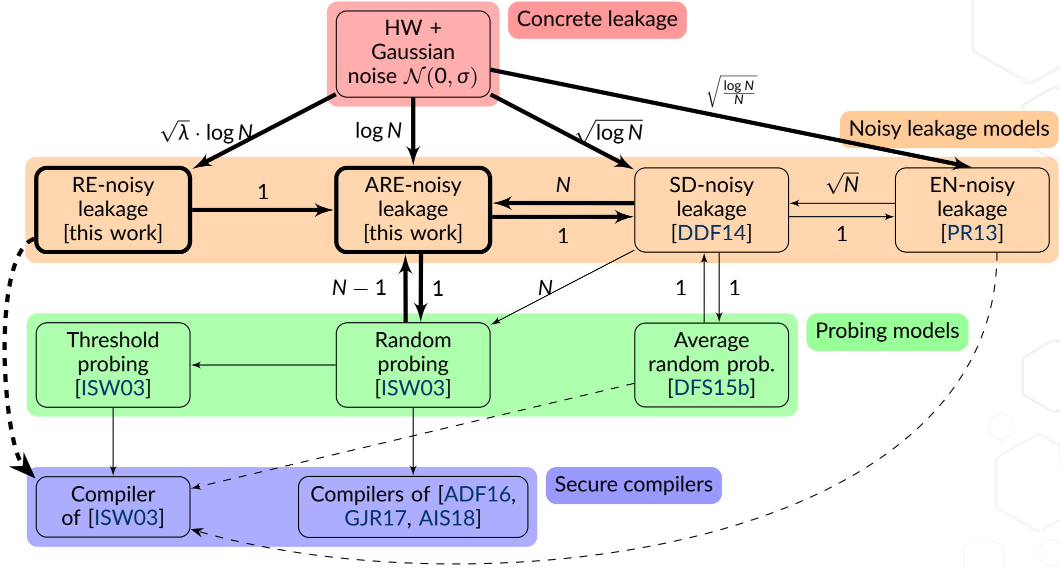

Roadmap

1. Unify the noisy leakage models and propose new ones
2. Link the noisy leakage models to a concrete modelization of leakage
3. Link the noisy leakage models to probing models
4. Prove compilers directly in a noisy leakage model

The Pointwise Mutual Information

Definition (Pointwise mutual information)

Let X, Y be random variables over $\mathcal{X}$. We note:

$$\mathrm{pmi}_{X,Y}(x, y) = \log\left(\frac{\Pr[X = x, Y = y]}{\Pr[X = x]\Pr[Y = y]}\right) .$$

$$\mathrm{PMI}_{X,Y}(x, y) = e^{\mathrm{pmi}_{X,Y}(x,y)} - 1 = \frac{\Pr[X = x, Y = y]}{\Pr[X = x]\Pr[Y = y]} - 1 .$$

Common tool in computational linguistics [CH89] as an association measure:

1. $\mathrm{pmi}(\text{“}\textit{Sean}\text{”}, \text{“}\textit{Penn}\text{”}) \gg 0$;
2. $\mathrm{pmi}(\text{“}\textit{Banana}\text{”}, \text{“}\textit{Bag}\text{”}) \approx 0$;
3. $\mathrm{pmi}(\text{“}\textit{Pineapple}\text{”}, \text{“}\textit{Italian Pizza}\text{”}) \ll 0$.

The mutual information verifies $\mathrm{MI}(X; Y) = \mathbb{E}_{(X,Y)}\left[\mathrm{pmi}_{X,Y}\right]$.

Unifying Leakage Metrics from the PMI

(Re)defining leakage metrics

- $\mathrm{EN}(X|Y) := \mathbb{E}_Y \sqrt{\mathbb{E}_X \left[\mathbb{P}[X]\, \mathrm{PMI}^2 \right]}$ [PR13]
- $\mathrm{SD}(X|Y) := \frac{1}{2} \cdot \mathbb{E}_X \mathbb{E}_Y \left[|\mathrm{PMI}| \right]$ [DDF14]
- $\mathrm{ARE}(X|Y) := \mathbb{E}_Y \left[\max_x |\mathrm{PMI}| \right]$ [this work, *average relative error*]
- $\mathrm{RE}(X|Y) := \max_{x,y} |\mathrm{PMI}|$ [this work, *relative error*]

Unifying Leakage Metrics from the PMI

(Re)defining leakage metrics

- $\mathrm{EN}(X|Y) := \mathbb{E}_Y \sqrt{\mathbb{E}_X \left[\mathbb{P}[X]\, \mathrm{PMI}^2\right]}$ [PR13]
- $\mathrm{SD}(X|Y) := \frac{1}{2} \cdot \mathbb{E}_X \mathbb{E}_Y \left[|\mathrm{PMI}|\right]$ [DDF14]
- $\mathrm{ARE}(X|Y) := \mathbb{E}_Y \left[\max_x |\mathrm{PMI}|\right]$ [this work, *average relative error*]
- $\mathrm{RE}(X|Y) := \max_{x,y} |\mathrm{PMI}|$ [this work, *relative error*]

We note that:

- ARE and RE are worst-case metrics;
- EN and SD are average-case metrics.

This raises three questions:

- Does this change their properties?
- Does it imply stronger conditions?
- Does it provide better proofs?

Some Nice Properties

Relations with other metrics

1. $2 \cdot \mathrm{SD}(X|Y) \leq \mathrm{ARE}(X|Y) \leq 2N \cdot \mathrm{SD}(X|Y)$;
2. $2 \cdot \mathrm{SD}(X|Y)^2 \leq \mathrm{MI}(X;Y) \leq 2 \cdot \mathrm{RE}(X|Y) \cdot \mathrm{SD}(X|Y)$.

→ The ARE- and SD-noisy leakage models are equivalent.
→ Bounds on MI simpler/tighter than previous ones [DFS15a, DDF14].

Self-reducibility

Let $f : X \to Y$ be a randomized leakage function.

1. If f is δ-RE-noisy for some X, then it is $\frac{2\delta}{1-\delta}$-RE-noisy for any X'.
2. If f is δ-ARE-noisy for some X, then it is $\frac{2\delta(1+\delta_{\mathrm{RE}})}{1-\delta}$-ARE-noisy for any X'.

→ **Consequence:** we don't care about the underlying distribution.
→ [DFS16] has a similar theorem for SD, but with a $O(N)$ blow-up, and only for X uniform.

From Concrete Leakage to Noisy Leakage

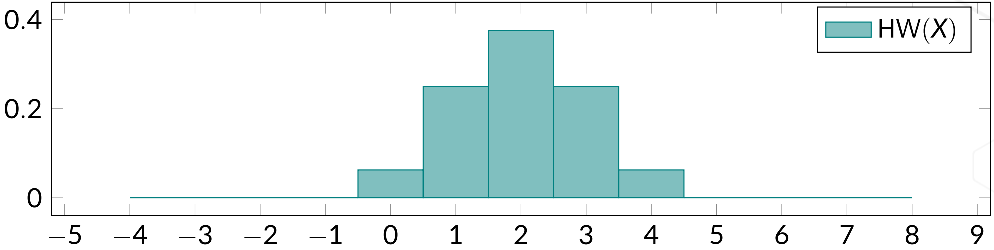

Figure 1: Distribution of HW(X) for X uniform in $\{0, \ldots, 2^4 - 1\}$

From Concrete Leakage to Noisy Leakage

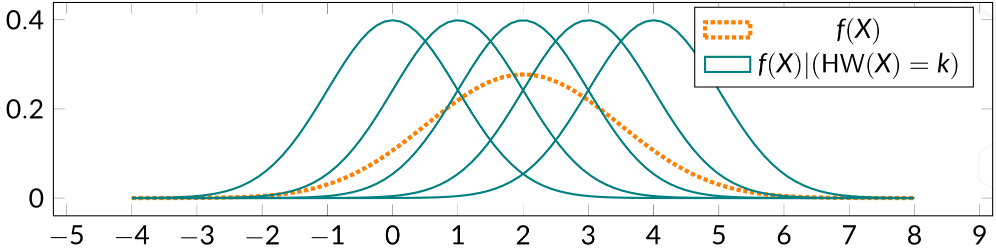

Figure 1: Distribution of $f(X) = \mathrm{HW}(X) + \mathcal{N}(0, \sigma)$ and $f(X)|(\mathrm{HW}(X) = k)$

From Concrete Leakage to Noisy Leakage

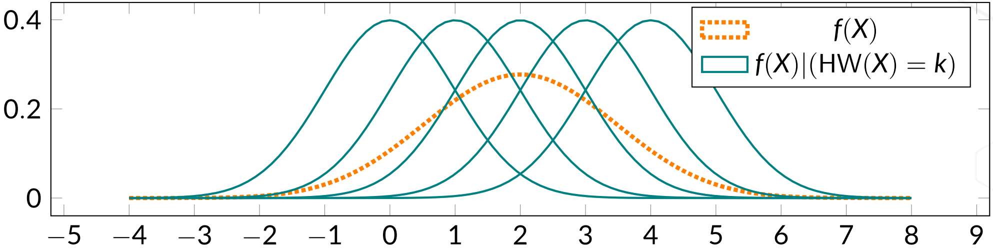

Figure 1: Distribution of $f(X) = \mathrm{HW}(X) + \mathcal{N}(0, \sigma)$ and $f(X)|(\mathrm{HW}(X) = k)$

Each metric (EN, SD, ARE, RE) can be interpreted as the average/max/... of:

$$\left| \frac{f(X)|(\mathrm{HW}(X) = k)}{f(X)} - 1 \right| .$$

From Concrete Leakage to Noisy Leakage

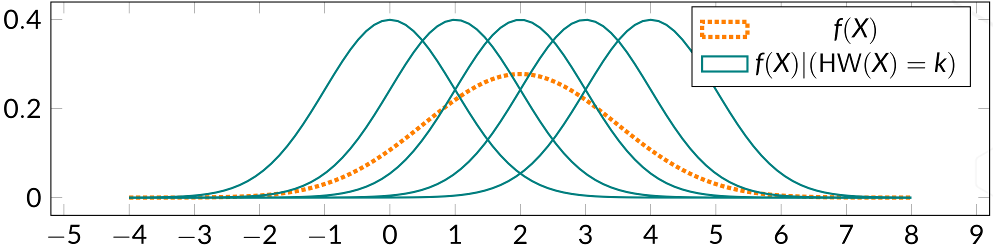

Figure 1: Distribution of $f(X) = \mathrm{HW}(X) + \mathcal{N}(0, \sigma)$ and $f(X)|(\mathrm{HW}(X) = k)$

We show that (omitting constant factors):

- $\mathrm{EN}(X|f(X)) \sim \frac{1}{\sigma}\sqrt{\frac{\log N}{N}}$
- $\mathrm{ARE}(X|f(X)) \sim \frac{\log N}{\sigma}$
- $\mathrm{SD}(X|f(X)) \sim \frac{\sqrt{\log N}}{\sigma}$
- $\mathrm{RE}(X|f(X)) \sim \frac{\tau \log N}{\sigma}$

Key takeaway: SD, RE and ARE essentially scale at the same speed.

Noisy Leakage $\Leftrightarrow$ Random Probing

Simulating a noisy adversary with a random probing adversary

- [DDF14]: a $(N \cdot \delta)$-random probing adv. can simulate a δ-SD-noisy adv.
- [this work]: a δ-random probing adv. can simulate a δ-ARE-noisy adv.

- Critical step is expressing $\varepsilon = 1 - \sum_y \min_x \mathbb{P}[f(x) = y]$ from δ:
  - if $\delta = \mathrm{SD}(X|f(X))$, a factor N is lost because "sum $\leq N \times$ max"
  - if $\delta = \mathrm{ARE}(X|f(X))$, no loss because "max $\leq$ max"
- We believe a fundamental reason is that random probing and ARE-noisy are "worst-case", whereas SD-noisy is "average-case".

Noisy Leakage $\Leftrightarrow$ Random Probing

Simulating a noisy adversary with a random probing adversary

- [DDF14]: a $(N \cdot \delta)$-random probing adv. can simulate a δ-SD-noisy adv.
- [this work]: a δ-random probing adv. can simulate a δ-ARE-noisy adv.

- Critical step is expressing $\varepsilon = 1 - \sum_y \min_x \mathbb{P}[f(x) = y]$ from δ:
  - if $\delta = \mathrm{SD}(X|f(X))$, a factor N is lost because "sum $\leq N \times$ max"
  - if $\delta = \mathrm{ARE}(X|f(X))$, no loss because "max $\leq$ max"
- We believe a fundamental reason is that random probing and ARE-noisy are "worst-case", whereas SD-noisy is "average-case".

We also show that an ARE-noisy adv. can simulate a random probing adv.
Consequence:

$$\text{ARE-noisy} \Leftrightarrow \text{SD-noisy} \Leftrightarrow \text{random probing} \Leftrightarrow \text{average random probing}$$

What we did

1. Unify existing noisy leakage metrics , propose new ones
   - **Tool:** pointwise mutual information
   - **New metrics:** RE and ARE
2. We link noisy leakage models to a concrete modelization of leakage
3. We reduce the ARE-noisy model to the random probing model :
   - No loss of a factor $O(N)$ as in [DDF14]
   - We show (leakage models) $\Leftrightarrow$ (probing models)
4. We prove compilers directly in the RE-noisy model
   - Hardness amplification
   - **Tool:** Rényi divergence
   - Parameters scale with #leakages (say 2^{30}), rather than security level (say 2^{256})
   - Not in this talk :-(

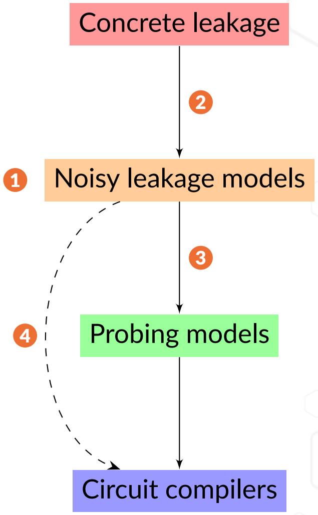

Thanks!

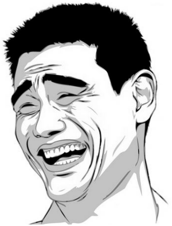

https://ia.cr/2019/138

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.

Marcin Andrychowicz, Stefan Dziembowski, and Sebastian Faust.
Circuit compilers with $O(1/\log(n))$ leakage rate.
In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 586–615. Springer, Heidelberg, May 2016.

Prabhanjan Ananth, Yuval Ishai, and Amit Sahai.
Private circuits: A modular approach.
In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 427–455. Springer, Heidelberg, August 2018.

David Brumley and Dan Boneh.
Remote timing attacks are practical.
In *USENIX Security 2003*. USENIX Association, August 2003.

Eric Brier, Christophe Clavier, and Francis Olivier.
Correlation power analysis with a leakage model.

In Marc Joye and Jean-Jacques Quisquater, editors, *CHES 2004*, volume 3156 of *LNCS*, pages 16–29. Springer, Heidelberg, August 2004.

Kenneth Ward Church and Patrick Hanks.
Word association norms, mutual information and lexicography.
In *ACL*, pages 76–83. ACL, 1989.

Alexandre Duc, Stefan Dziembowski, and Sebastian Faust.
Unifying leakage models: From probing attacks to noisy leakage.
In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 423–440. Springer, Heidelberg, May 2014.

Alexandre Duc, Sebastian Faust, and François-Xavier Standaert.
Making masking security proofs concrete - or how to evaluate the security of any leaking device.
In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 401–429. Springer, Heidelberg, April 2015.

Stefan Dziembowski, Sebastian Faust, and Maciej Skorski.
Noisy leakage revisited.
In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 159–188. Springer, Heidelberg, April 2015.

Stefan Dziembowski, Sebastian Faust, and Maciej Skórski.
Optimal amplification of noisy leakages.
In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 291–318. Springer, Heidelberg, January 2016.

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

Dahmun Goudarzi, Antoine Joux, and Matthieu Rivain.
How to securely compute with noisy leakage in quasilinear complexity.
Cryptology ePrint Archive, Report 2017/929, 2017.
http://eprint.iacr.org/2017/929.

Karine Gandolfi, Christophe Mourtel, and Francis Olivier.
Electromagnetic analysis: Concrete results.
In Çetin Kaya Koç, David Naccache, and Christof Paar, editors, *CHES 2001*, volume 2162 of *LNCS*, pages 251–261. Springer, Heidelberg, May 2001.

Daniel Genkin, Lev Pachmanov, Itamar Pipman, Adi Shamir, and Eran Tromer.
Physical key extraction attacks on pcs.
Commun. ACM, 59(6):70–79, 2016.

Daniel Genkin, Adi Shamir, and Eran Tromer.
RSA key extraction via low-bandwidth acoustic cryptanalysis.
In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 444–461. Springer, Heidelberg, August 2014.

Yuval Ishai, Amit Sahai, and David Wagner.
Private circuits: Securing hardware against probing attacks.
In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 463–481. Springer, Heidelberg, August 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.

Emmanuel Prouff and Matthieu Rivain.
Masking against side-channel attacks: A formal security proof.

In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 142–159. Springer, Heidelberg, May 2013.