Noisy Leakage Revisited

Stefan Dziembowski
University of Warsaw

Sebastian Faust
RUB
Ruhr-Universität Bochum

Maciej Skórski
University of Warsaw

INNOVATIVE ECONOMY
NATIONAL COHESION STRATEGY

FNP
Foundation for Polish Science

EUROPEAN UNION
EUROPEAN REGIONAL DEVELOPMENT FUND

This paper in a nutshell

We introduce a **new model for noisy leakage**.

This allows us to **improve the "noise parameters"** of a leakage-resilient construction from [Duc, D., Faust, Eurocrypt 2014].

The improvement is of factor $|\mathbb{F}|$, where $\mathbb{F}$ is the field over which the computation is made.

Cryptographic Implementations

very secure

- well-defined mathematical object
- often proof-driven security analysis

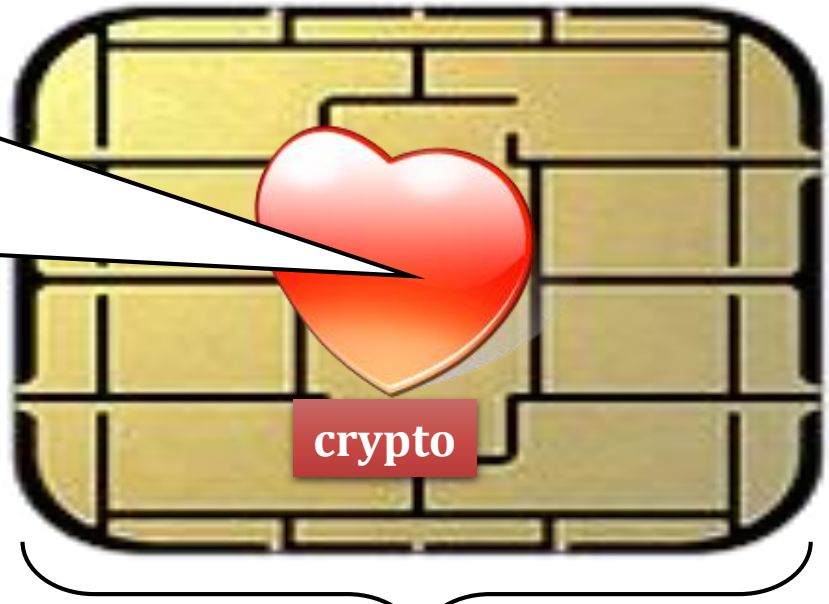

much less secure!

side-channel leakage devastating for security

Leakage-resilient cryptography

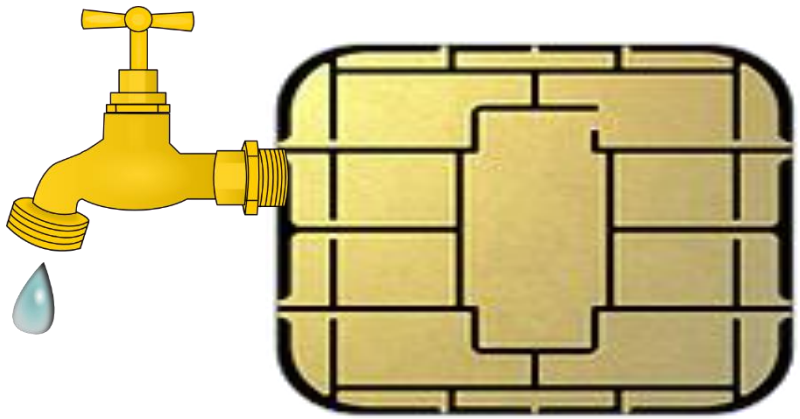

Goal: extend the proof-driven security analysis for implementations that leak information.

Many theoretical models

Probing

Bounded leakage

Bounded space

Low-depth circuits

Hard-to-invert leakage

Models do not match with my engineering experience

Leakages are not quantitatively bounded

Model preferred by practitioners: noisy leakage

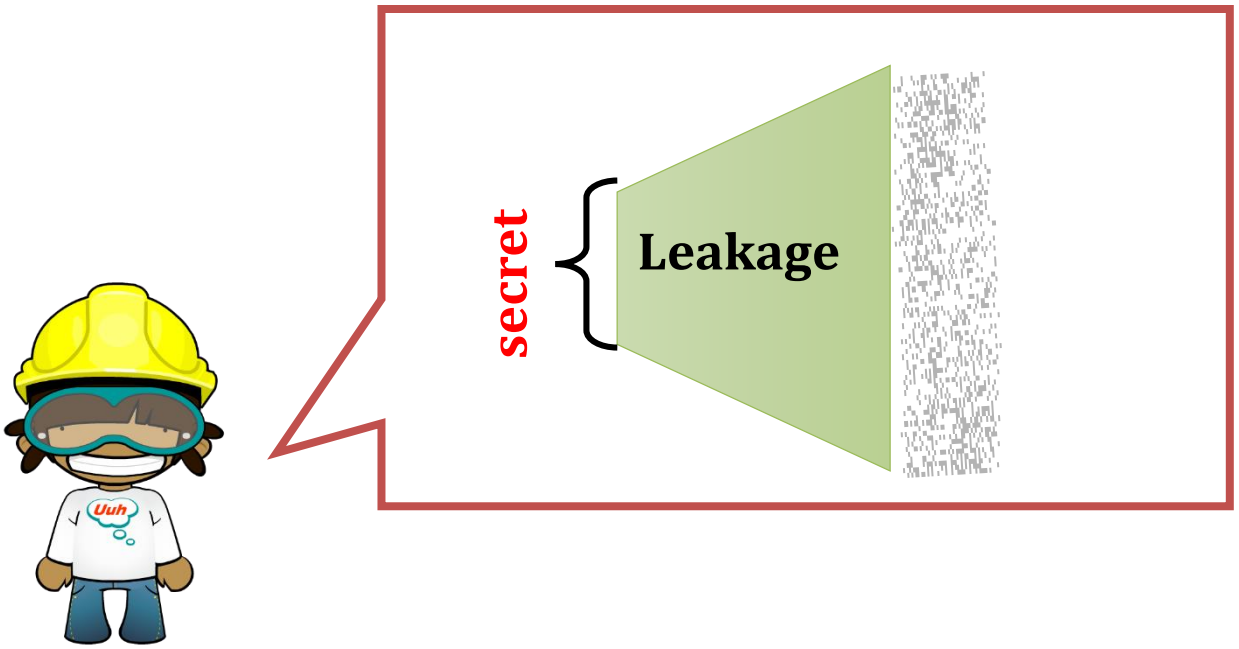

Formalized by Prouff and Rivain [Eurocrypt 2013]

Model of Prouff and Rivain

Computation is modelled as a circuit (over some field $\mathbb{F}$).

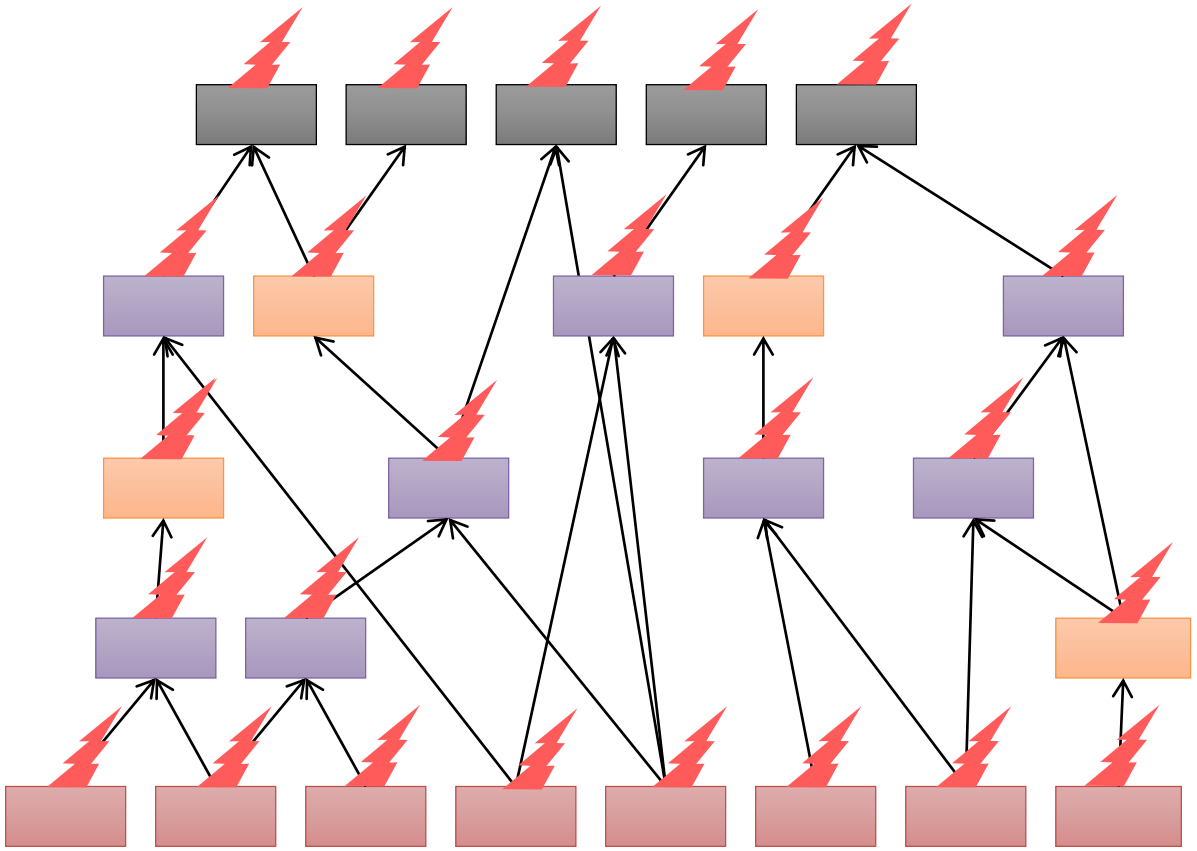

Every wire leaks information independently.

Noisy functions of Prouff and Rivain

a value on the wire

$C \in \mathbf{F}$

a "δ-noisy function"
$N: \mathbf{F} \to \{0,1\}^*$

$N(C)$

N is a δ-noisy function if
"the leakage from random elements is close in the sense of the statistical distance"

Statistical distance

$$\Delta(A\,;B) \coloneqq \frac{1}{2}\sum_{x} |P(A=x) - P(B=x)|$$

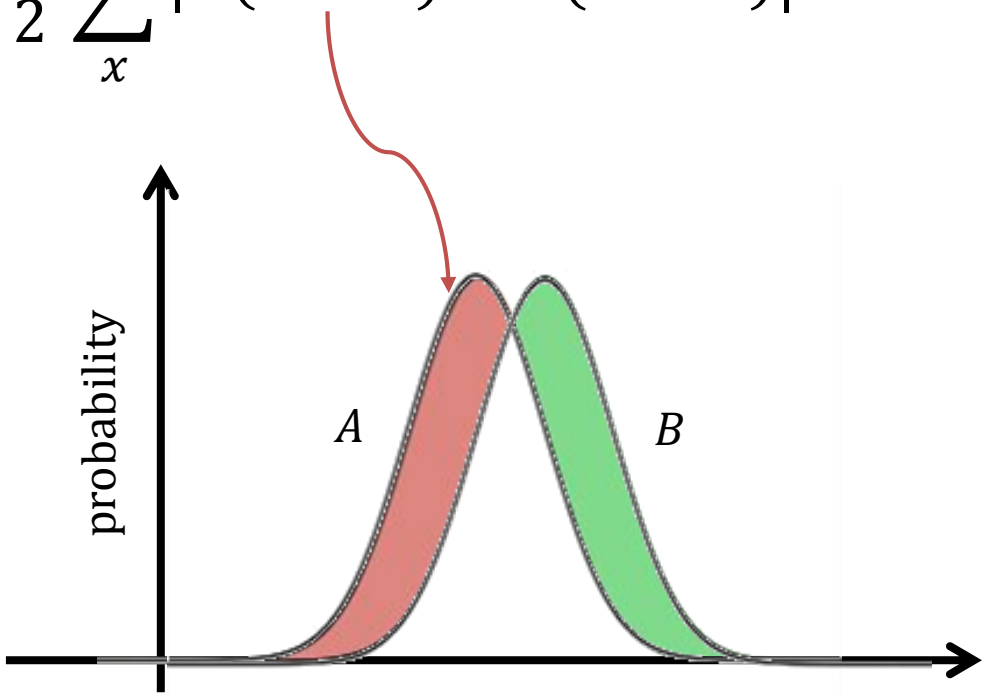

Notation: $\boldsymbol{d(A)}$ denotes distance of $\boldsymbol{A}$ from a uniform distribution $\boldsymbol{U}$ over the same set: $\boldsymbol{d(A) \coloneqq \Delta(A;U)}$

Conditional versions

$\boldsymbol{\Delta(A; B \mid C)}$ and $\boldsymbol{d(A|C)}$ are defined as expected values:

$$\boldsymbol{\Delta(A; B|C) := \sum_c P(C = c) \cdot \Delta(A; B|C = c)}$$

and

$$\boldsymbol{d(A|C) := \sum_c P(C = c) \cdot d(A|C = c)}$$

Noisy functions of Prouff and Rivain

a value on the wire

$C \in \mathbf{F}$

a "δ-noisy function"
$N: \mathbf{F} \to \{0,1\}^*$

$N(C)$

N is a δ-noisy function if "the leakage from random elements is close":

for uniformly random C_0 and C_1 from $\mathbf{F}$ we have

$$\Delta\big((N(C_0); N(C_1) \mid C_0, C_1)\big) \leq \delta$$

How to interpret the noise definition?

Fact:

for uniformly random C_0 and C_1 from $\mathbf{F}$ we have

$$\Delta\big((N(C_0); N(C_1) \mid C_0, C_1)\big) \leq \delta$$

No adversary A wins the following game with probability greater than $\frac{1}{2} + \delta$:

1. $C_0, C_1 \leftarrow \mathbf{F}$
2. $b \leftarrow \{0, 1\}$
3. The adversary A learns (C_0, C_1) and $N(C_b)$ and has to guess b.

Some examples ($\mathbb{F} = \mathbb{Z}_2$)

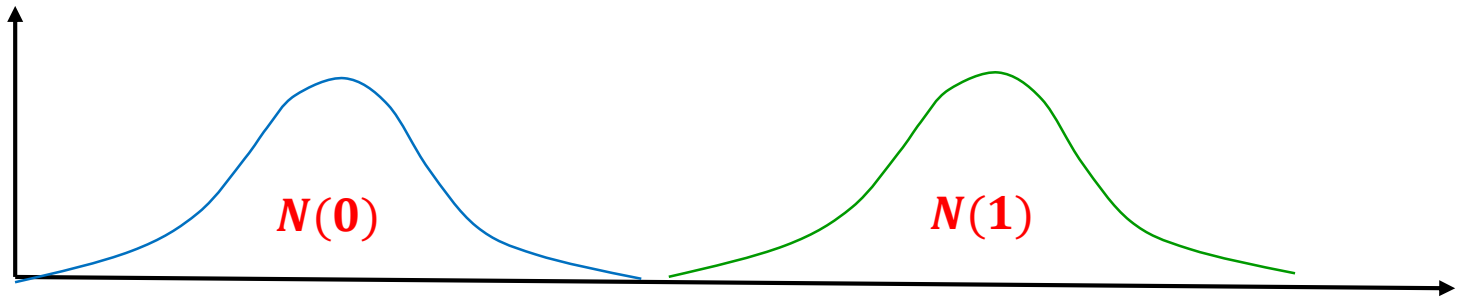

No noise $\delta = 1$: very informative leakage

$\Longrightarrow$ Adversary learns $N(C)$: full knowledge about C.

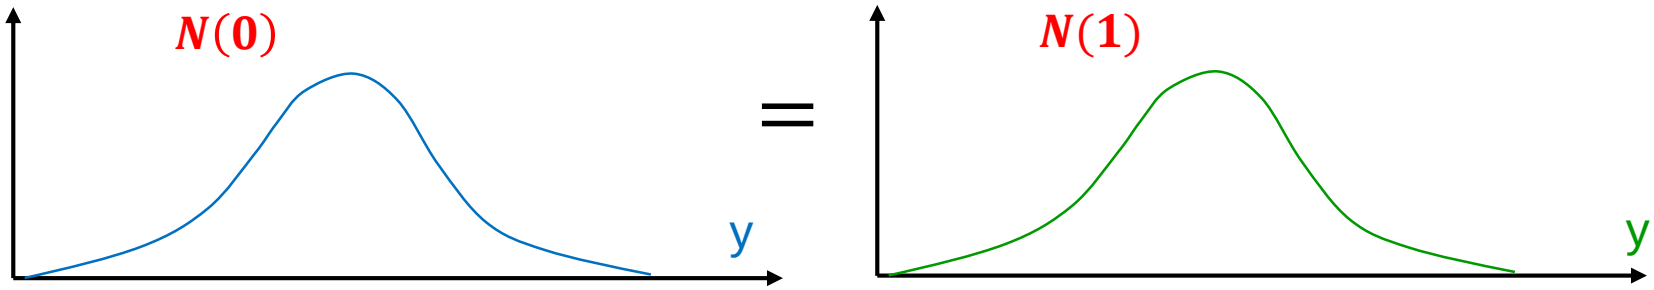

High noise $\delta = 0$: non-informative leakage

$\Longrightarrow$ Adversary learns $N(C)$: no knowledge about C.

The main result of [DDF, EC 2014]

A reduction of the [PR13]-model to a "random probing model"

Random probing model:

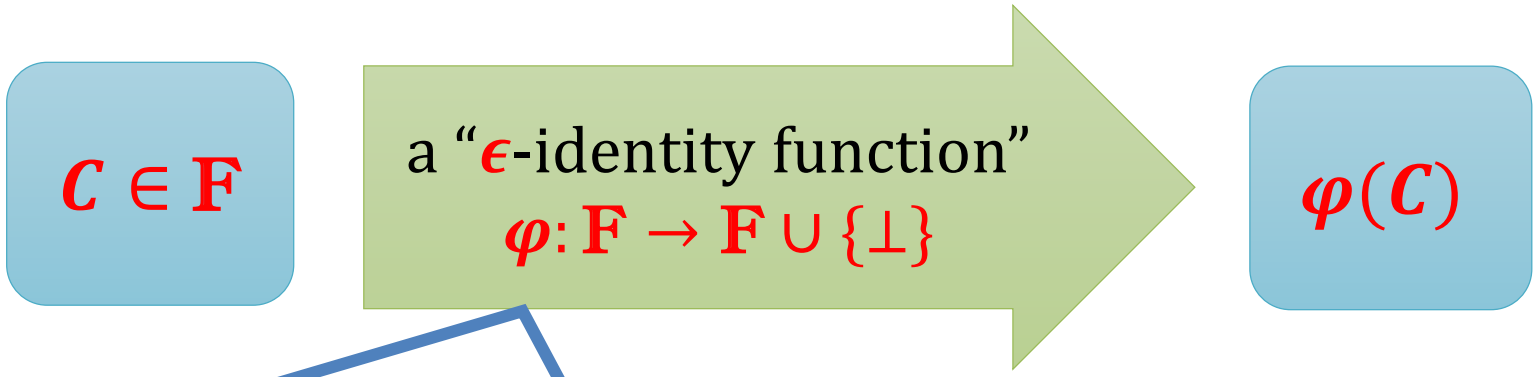

φ is a ϵ-identity function if every element leaks with probability $1 - \epsilon$.

More formally $\forall_c \quad \varphi(c) :=$

$\perp$ with probability ϵ

c otherwise

An observation

The random probing model is **much easier to analyze**:

- either an element "**leaks completely**",

- or it is **completely hidden** from the adversary

It is very similar to a classical "**probing**" model of [ISW03].

An example: additive encoding of a secret s.

Suppose $s_1, \ldots, s_k$ are random elements of $\mathbf{F}$ such that

$$s = s_1 + \cdots + s_k.$$

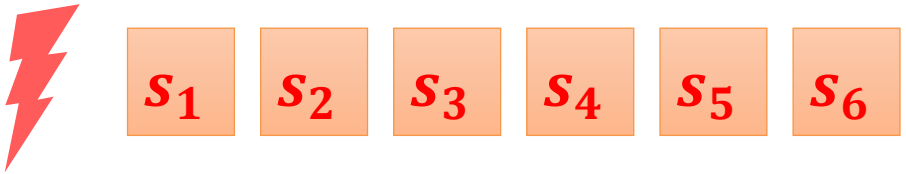

$\varphi(s_i) =$ s_1 s_2 $\perp$ s_4 s_5 s_6

It is enough that the adversary doesn't learn one s_i and s remains secret.

Note: $P(\exists_i\, \varphi(s_i) = \perp)$ is overwhelming in k.

The parameters of the [DDF14] reduction

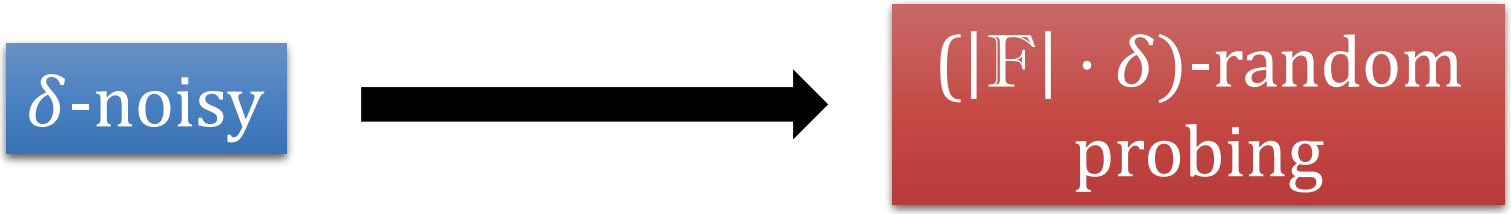

Every $\boldsymbol{\delta}$-noisy function can be "simulated" by an $\boldsymbol{\epsilon}$-random probing function $\boldsymbol{\varphi}$, where

$$\boldsymbol{\epsilon = |\mathbb{F}| \cdot \delta}$$

this "factor $|\mathbb{F}|$ loss" is a problem for some applications

Example of an application

Γ – a circuit

Γ is secure against ϵ-random probing → Γ is secure against δ-noisy leakage

where $\epsilon = |F| \cdot \delta$

Caveat:

to get meaningful results one needs to assume that

$$\delta < 1/|F|$$

Problem when we consider larger fields.
For example for **AES**: **F = GF(2^8)**).

Natural question

Can the "factor $|\mathbf{F}|$ loss" be avoided?

First observation:

The [DDF14] reduction from δ-noisy to ϵ-probing model is essentially optimal.

Our contribution

We propose

the ϵ-average probing model

that has the following properties:

1. it allows a reduction from δ-noisy leakage model **without** the "factor $|\mathbf{F}|$ loss",

   yet

2. it still permits to do **simple security proofs** in many cases.

The difference

φ is a ϵ-identity function if **every** element leaks with probability $\mathbf{1} - \epsilon$.

More formally: $\forall_c \ \varphi(c) \coloneqq \begin{cases} \perp & \text{with probability } \epsilon \\ c & \text{otherwise} \end{cases}$

φ is a ϵ-average identity function if a **random element** leaks with probability $\mathbf{1} - \epsilon$.

More formally: $\varphi(C) \coloneqq \begin{cases} \perp & \text{with probability } \epsilon \\ C & \text{otherwise} \end{cases}$

where $C \leftarrow \mathbf{F}$.

Our first main result

δ-noisy ⟵ δ-average probing

Every $\boldsymbol{\delta}$-noisy function can be "simulated" by a $\boldsymbol{\delta}$-average probing function $\boldsymbol{\varphi}$.

no "factor $|\mathbf{F}|$ loss"

This is essentially optimal since we can also show:

$(2 \cdot \delta)$-noisy ⟵ δ-average probing

A problem

In this model when the adversary gets $\perp$ she still knows something about the leaking element.

Example: suppose $\mathbf{F = Z_2}$ and:

$$\boldsymbol{P(\varphi(0) = \perp) = \frac{1}{4}}$$

$$\boldsymbol{P(\varphi(1) = \perp) = \frac{1}{2}}$$

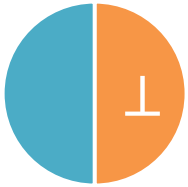

Then for a uniformly random $\boldsymbol{C \leftarrow F}$ we have

$$\boldsymbol{P(C = 1 \mid \varphi(C) = \perp) > P(C = 0 \mid \varphi(C) = \perp)}$$

Our second main contribution

A technique for dealing with this problem.

First step: make the randomness in φ explicit, i.e., assume φ is a deterministic function of a type

$$\varphi: \mathbf{F} \times \mathcal{R} \rightarrow \mathbf{F} \cup \{\perp\}$$

Lemma. Suppose φ is an ϵ-average probing function. Then for uniform variables $X \leftarrow \mathbf{F}$ and $R \leftarrow \mathcal{R}$ we have:

$$d(X \mid \varphi(X, R) = \perp, R) = \epsilon.$$

Suppose $d(X \mid \varphi(X, R) = \perp, R) = \epsilon$

Consider security of the additive encoding:

$s_1, \ldots, s_k$ are random elements of $\mathbf{F}$ such that

$$s = s_1 + \cdots + s_k.$$

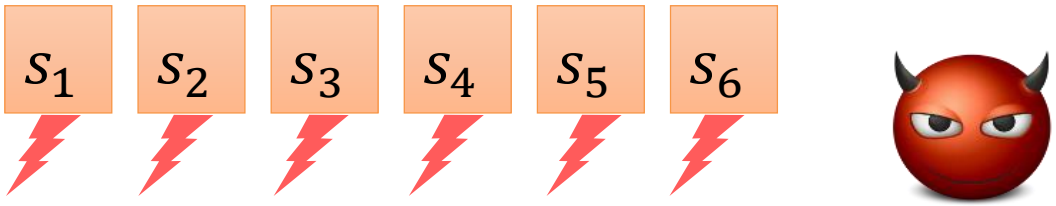

Using **Chernoff-** and **Markov-type** arguments we can show that:

there exists a linear (in k) number of positions i such that

$$d(s_i \mid \text{view of adversary}) \leq \hat{\epsilon}.$$

Call these positions "**good**".

$\hat{\epsilon} \approx \epsilon$

Why is it useful?

Suppose the "good" positions are at the beginning.

$d(s_1 \mid \text{view of } A) \le \hat{\epsilon}$

$\vdots$

$d(s_k \mid \text{view of } A) \le \hat{\epsilon}$

(*)

$d(s_1 + \cdots + s_k \mid \text{view of } A) \le (2 \cdot \hat{\epsilon})^k/2$

$d(s_1 + \cdots + s_n \mid \text{view of } A) \le (2 \cdot \hat{\epsilon})^k/2$

negligible for $\hat{\epsilon} < \frac{1}{2}$

How to prove (*)?

Lemma [Maurer, Pietrzak, and Renner, CRYPTO 2007]:

For every two independent random variables S_1 and S_2 over a finite group $(\mathbf{F}, +)$ such that

$$d(S_1), d(S_2) \leq \hat{\epsilon}$$

we have that

$$d(S_1 + S_2) \leq 2 \cdot \hat{\epsilon}^2$$

Applying this lemma inductively we get:

$$d(S_1 + \cdots + S_k) \leq (2 \cdot \hat{\epsilon})^k / 2$$

This also holds for the **conditional** statistical distance.

Therefore

$$d(s_1 \mid \text{view of } A) \le \hat{\epsilon}$$
$$\vdots$$
$$d(s_k \mid \text{view of } A) \le \hat{\epsilon}$$

$$\Rightarrow\ d(s_1 + \cdots + s_k \mid \text{view of } A) \le (2 \cdot \hat{\epsilon})^k / 2$$

and **(*)** is proven.

As a result we get:

The first proof of **additive encoding security** in the noisy leakage model that is secure for **noise independent of the field size**.

With similar techniques we also get a **compiler for any functionality** that is secure for **noise independent of the field size** (assuming a leak-free component).

Thank you!