A Formal Practice-Oriented Model for the Analysis of Side-Channel Attacks

F.-X. Standaert, T.G. Malkin, M. Yung

http://eprint.iacr.org/2006/139

Dept. of Computer Science, Columbia University.
UCL Crypto Group, Université Catholique de Louvain.
RSA Laboratories.

January 2007

Part I:

An Intuitive view of the model

- General context
- Examples of unanswered questions => 1st observation
- Comparison with black box attacks => 2nd observation
- Conclusion

General Context

Black box cryptanalysis…

only uses the primitive's inputs and outputs, *e.g*, the plaintexts, ciphertexts for block ciphers

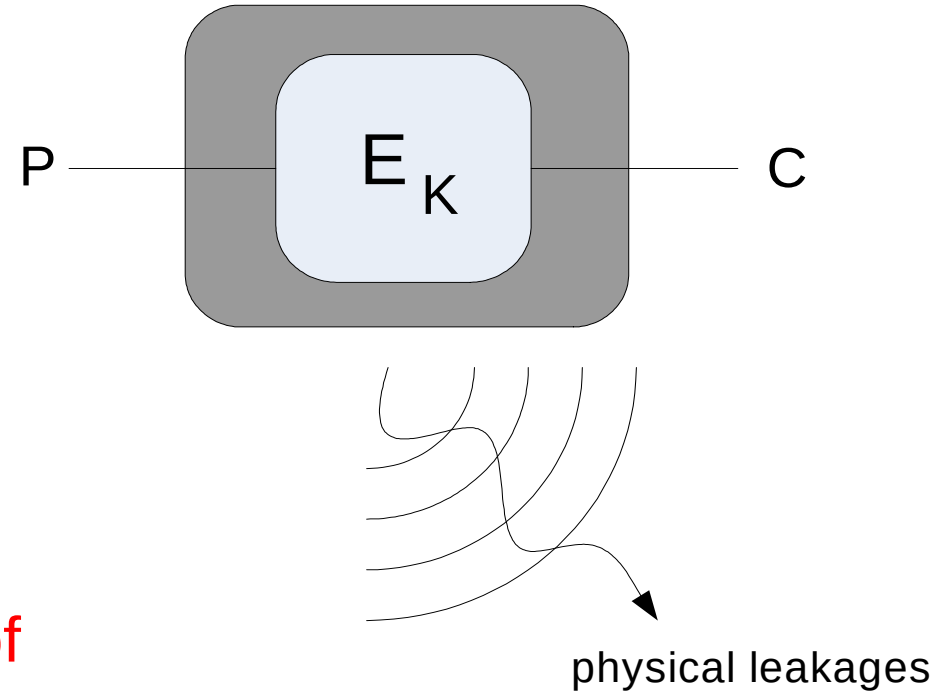

Side-channel attacks…

additionally takes advantage of physical leakages, *e.g.* power consumption, timing information, electromagnetic radiation

- Side-channel attacks (the story made short):
  - Powerful but specific (usually target a particular implementation rather than an abstract algorithm)
  - Still generic (most devices can be targeted)
  - Hard to evaluate (many different circuit technologies and statistical distinguishers)
  - Hard to prevent (actual countermeasures usually only make the attack more difficult)

  $\Rightarrow$ Many open questions

Example of (simple but) unanswered questions

AES

Given one implementation, how to we compare two side-channel adversaries?

DPA

template

Another example of unanswered questions

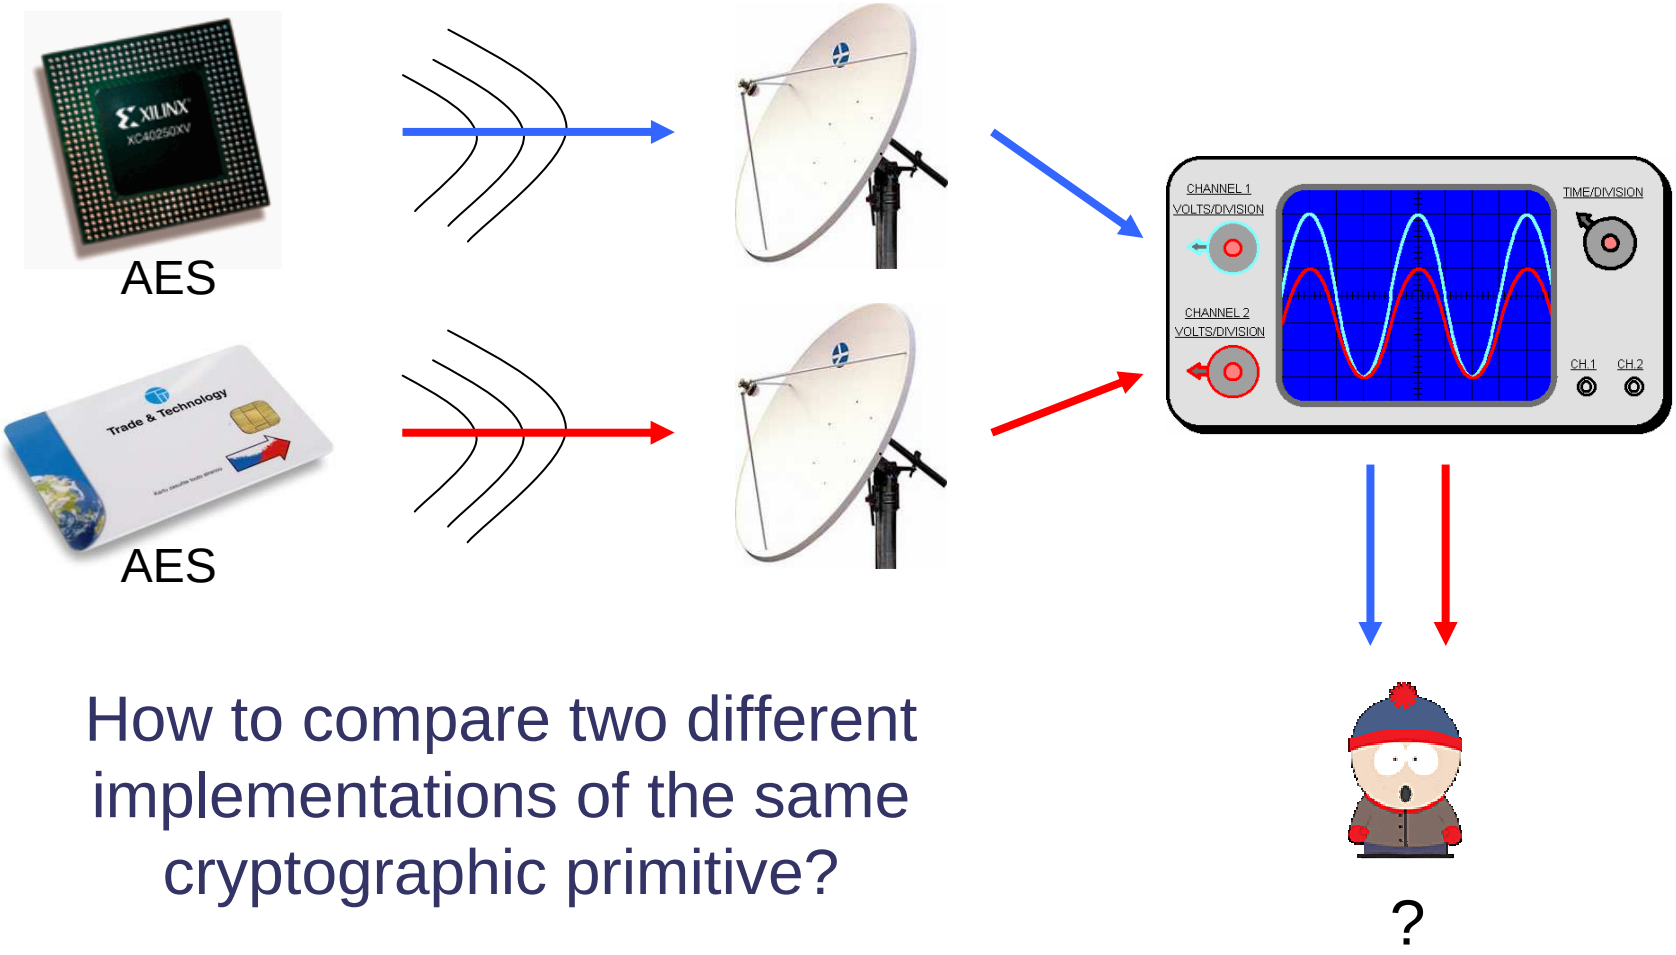

How to compare two different implementations of the same cryptographic primitive?

First observation

- It is important to distinguish the quality of an implementation from the strength of a side-channel adversary

Comparison with black box attacks, *e.g.* linear cryptanalysis

query

response

linear cryptanalysis

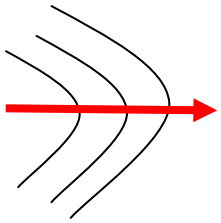

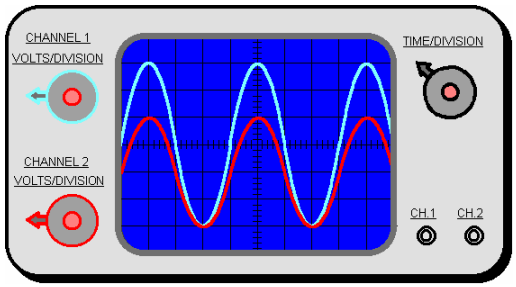

response

query

DPA

- Both attacks work in two steps:

  (1) Query the target and get information

  (2) Exploit the information with a distinguisher

- Both attacks yield two questions:

  (1) What is the amount of information provided by the adversarial queries?

  (2) How successfully can the adversary turn this information into a successful attack?

Second observation

- Black box attacks mainly care about security. Side-channel attacks don't, because the same amount of queries to different implementations of the same primitive give rise to different amounts of information

$\Rightarrow$ Information discriminates different implementations

Conclusion

A formal model for the analysis of side-channel attacks should consider:

implementations ↑ **information**

adversaries ↑ **security**

Intuitive expectations for the model

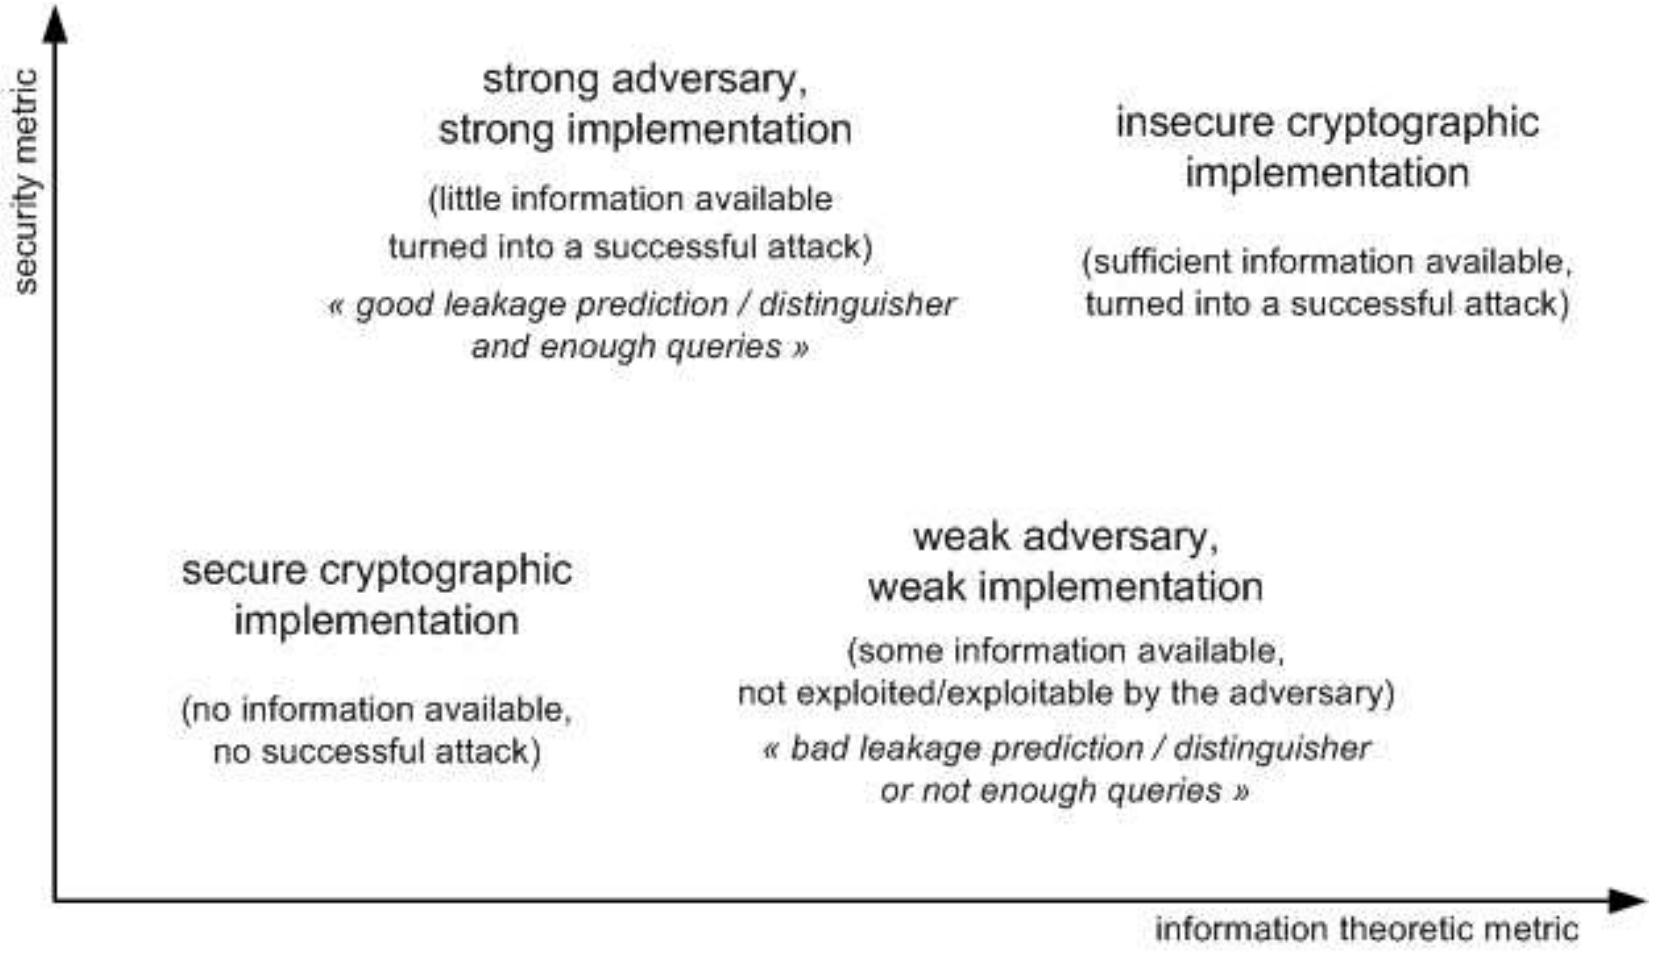

Part II:
Precise description of the model

Motivations and objectives

- A formal practice oriented model:
  - To understand the underlying mechanisms of physically observable cryptography
  - To evaluate (*i.e*. analyze and compare) actual implementations and adversaries
  - …to construct provably secure designs…

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Background: Micali & Reyzin model

- TCC 2004:

Physically Observable Cryptography

- Very general model (but few positive results)
- An open question was to meaningfully restrict the model to realistic adversaries

MR's Informal axioms

1. Computation and only computation leaks information

2. The same computation leaks differently on different computers

3. The information leakage depend on the chosen measurement

4. The information leakage is local (i.e. the maximum amount of information leaked by a device is the same in any execution of the algorithm with the same inputs)

5. All information leaked through physical observations can be efficiently computed from a computer's internal configuration

MR's Definitions

- Abstract computer: $\alpha := \{\alpha_1, \alpha_2, ..., \alpha_n\}$ with virtual memory Turing machines α_i's (VTMs)
- Physical computer: $\varphi = (\alpha, \mathcal{L})$, with $\varphi_i = (\mathcal{L}_i, \alpha_i)$
- Leakage functions: $\mathcal{L}(C_\alpha, M, R)$
  - *Cα:* abstract computer internal configuration
  - *M:* specification of the measurement
  - *R:* random string

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Target circuit

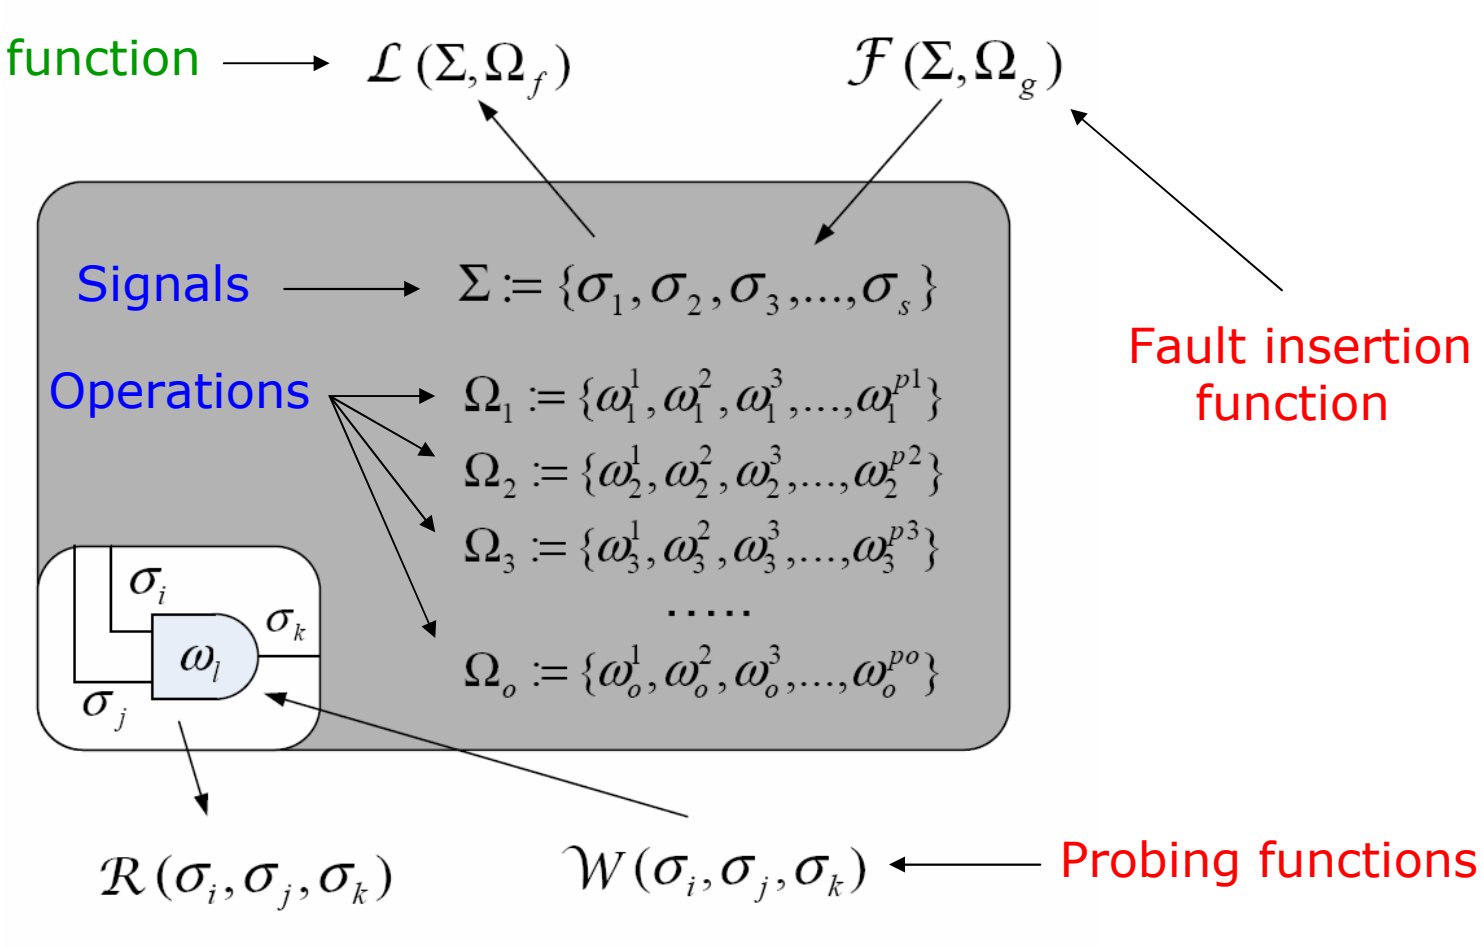

- With respect to MR:
  - Oracles and operations can be translated into abstract computers and VTMs
  - Signals are the VTM's inputs/outputs
- Abstract computers = cryptographic primitives
- Physical computers = implementations
- Implementations may differ by:
  - Their leakage function
  - Their division into elementary operations

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Target block cipher

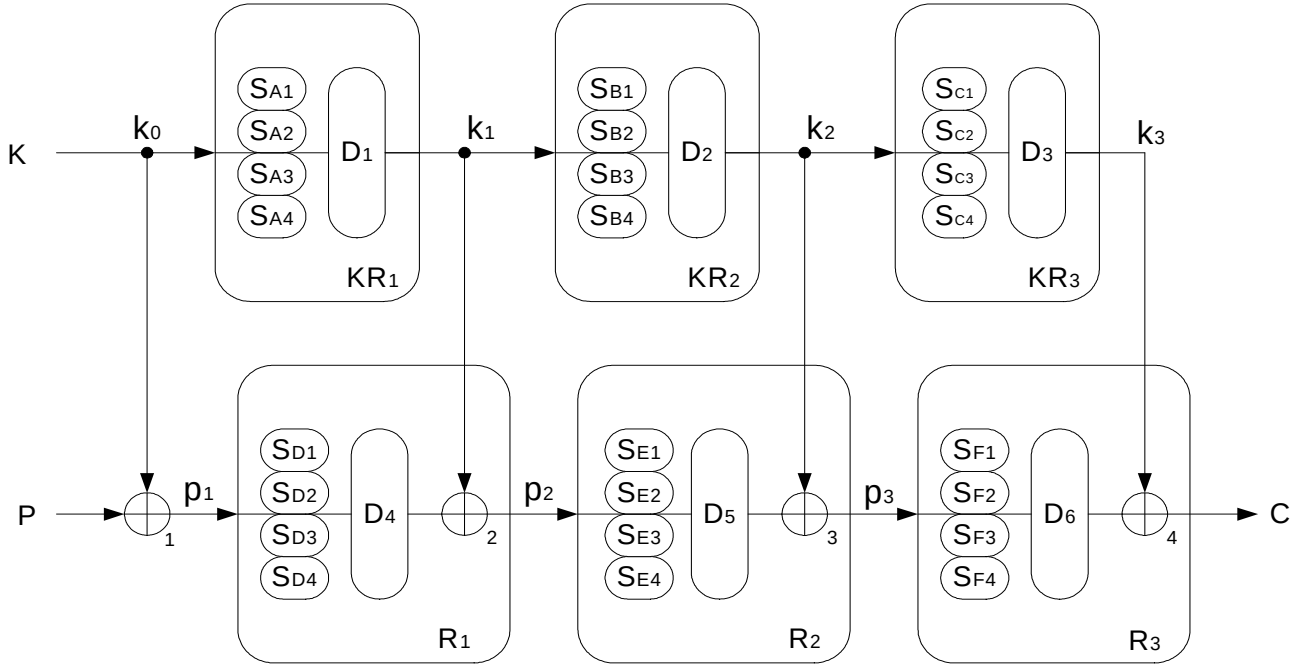

$$\Omega_1 := \{R_1, R_2, R_3, KR_1, KR_2, KR_3\}$$

$$\Omega_1 := \{\oplus_1, \ldots, \oplus_4, S_A, \ldots, S_F, D_1, \ldots, D_6\}$$

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Adversarial context

- Non adaptive known plaintext
- Non adaptive known ciphertext
- Non adaptive chosen plaintext
- Non adaptive chosen ciphertext
- Adaptive chosen plaintext
- Adaptive chosen ciphertext

& Strategy

- Hard strategy:

Given some physical observations and
a resulting classification of the key candidates,
select the best classified key only

- Soft strategy

Given some physical observations and
a resulting classification of the key candidates,
select the h best classified keys

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Generic side-channel attack

- According to some approximation of the leakage function, an adversary *predicts* (a part/function of) the key dependent leakages emanated from a device
- It then *measures* the actual leakages from the target physical implementation
- It finally *compares* the actual leakages with the key dependent predictions. If the attack is successful, it is expected that the correct key candidate gives rise to the best leakage prediction which can be detected with a side-channel distinguisher

Leakage predictions

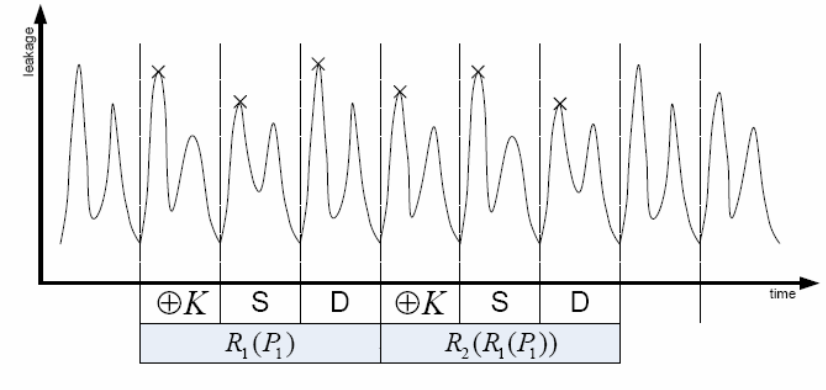

Univaritate

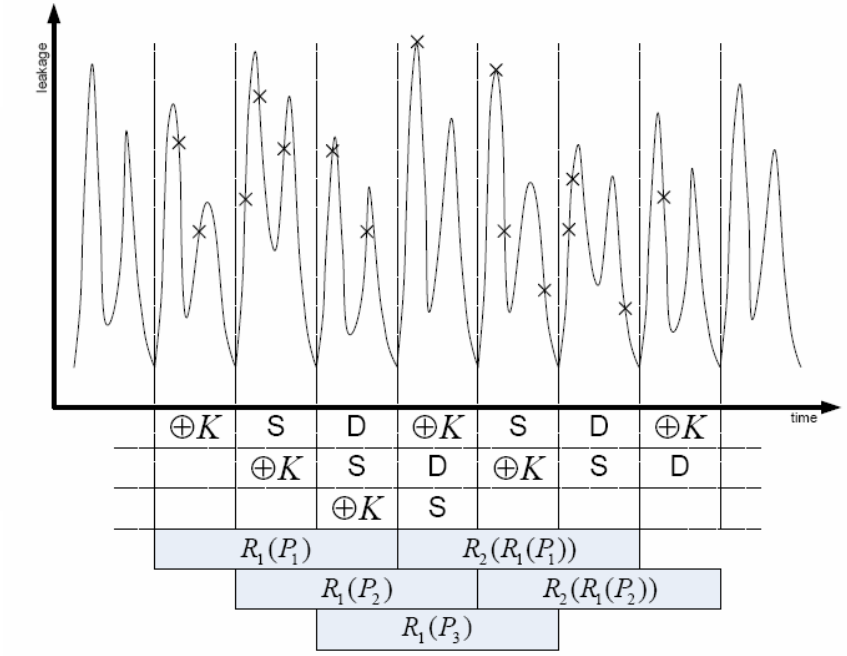

Multivaritate

Classification

- Non profiled
- Device profiled
- Key profiled

Adversary definition

$$A^{C}_{f_K,\mathcal{L}}(\tau, q)$$

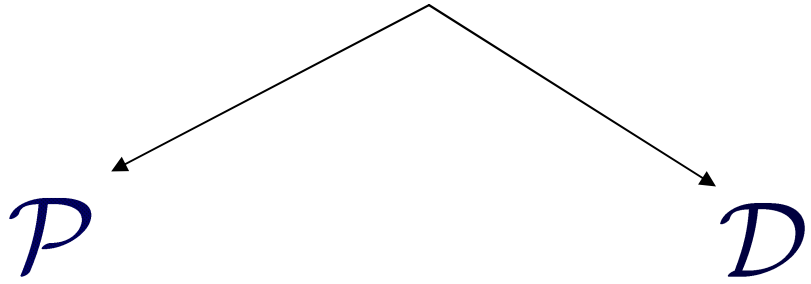

$\mathcal{P}$ prediction function

$\mathcal{D}$ distinguisher

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

- Memory complexity:
  - Size of the key guess: G

- Time and data complexity: τ, q

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Definition of security: Side-channel key recovery

Experiment $\mathbf{Exp}^{\text{sc-kr}}_{f_K,\mathcal{L}}$

$K \xleftarrow{R} \{0,1,2,\ldots,G-1\};$

$K^* \leftarrow A^{C}_{f_K,\mathcal{L}}(\tau,q);$

if $K = K^*$ **then** return 1;

else return 0;

- Key recovery advantage of the adversary

$$\mathbf{Adv}^{\text{sc-kr}}_{A}(\tau,q) = \mathbf{P}\ [\mathbf{Exp}^{\text{sc-kr}}_{f_K,\mathcal{L}} = 1]$$

- Key recovery advantage of the implementation

$$\mathbf{Adv}^{\text{sc-kr}}_{f_K,\mathcal{L}}(\tau,q) = \max_{A}\{\mathbf{Adv}^{\text{sc-kr}}_{A}(\tau,q)\}$$

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage function
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Evaluation of a side-channel attack

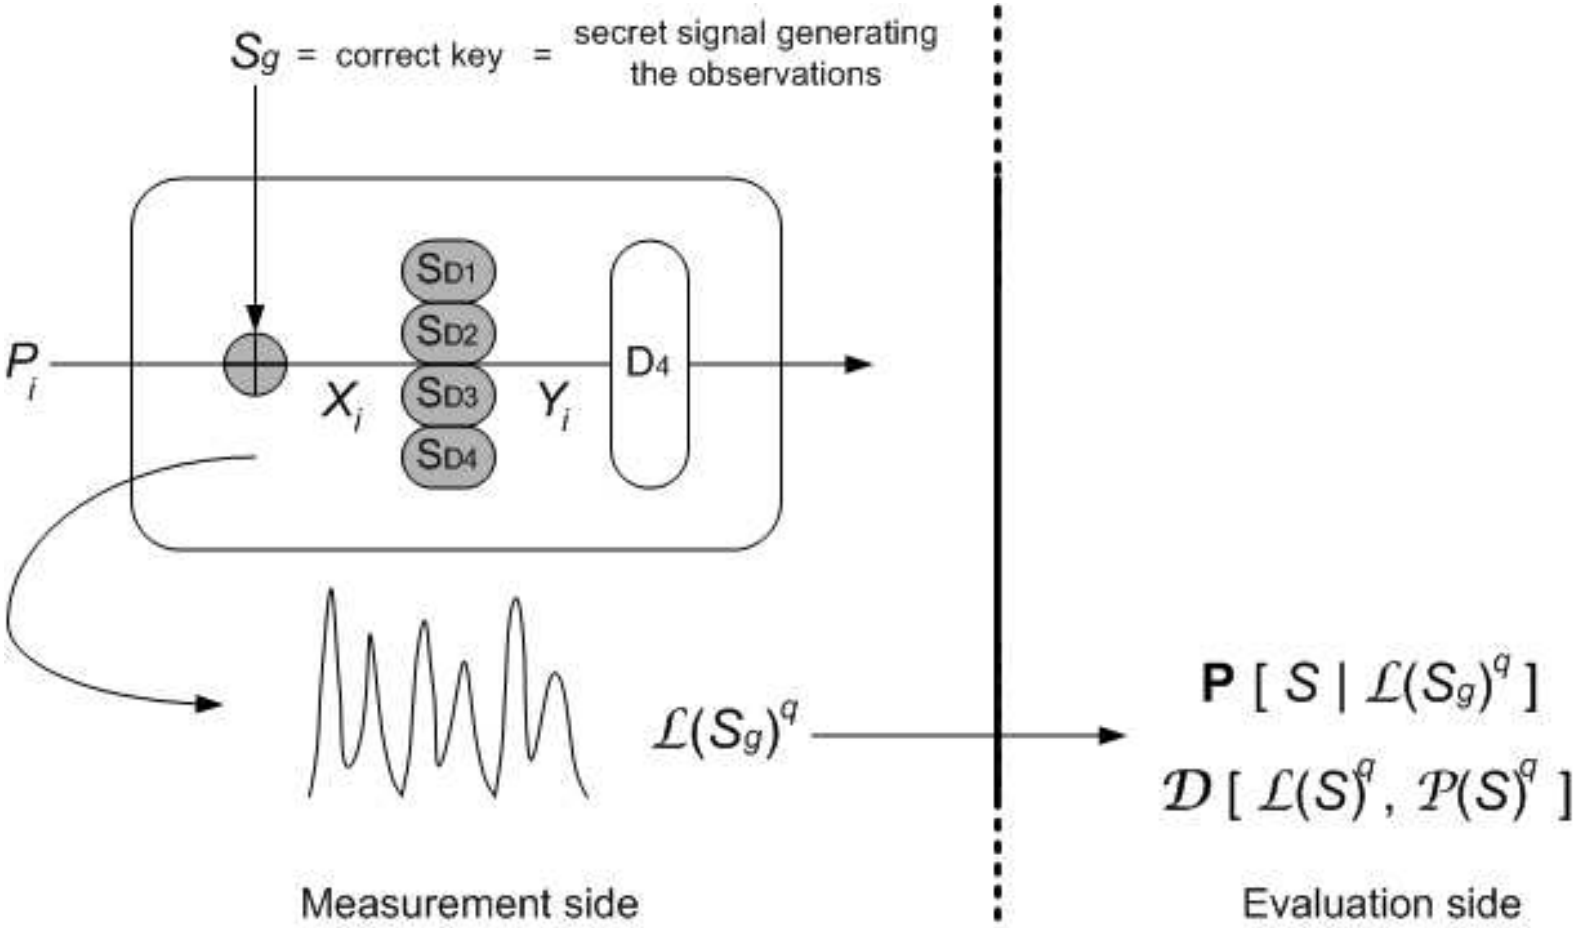

- Two questions to face:
  - What is the amount of information provided by a given leakage function?
  - How successfully can an adversary turn this information into a successful attack?

Definitions

- $L_{S_g}^q = \mathcal{L}(S_g)^q$: an observation generated by a secret S_g and q queries to the target device
- $P_S^q = \mathcal{P}(S)^q$: the adversary's predictions
- $\mathcal{D}(L_{S_g}^q, P_S^q)$: the distinguisher used by the adversary to compare an actual observation of a leaking device with its key dependent predictions

Security metric: average success rate

- Keys selected by the adversary (hard strategy):

$$M^q_{S_g} = \{\hat{s} \mid \hat{s} = \underset{S}{argmax}\, \mathcal{D}(L^q_{S_g}, P^q_S)\},$$

- Index matrix:

$$\mathrm{I}^q_{S_g,S} = \quad \frac{1}{|M^q_{S_g}|} \textbf{ if } S \in M^q_{S_g}, \quad \textbf{else } 0$$

- Success rate:

$$\mathbf{S_R}(S_g, q) = \underset{L^q_{S_g}}{\mathbf{E}}\, \mathrm{I}^q_{S_g,S_g}, \qquad \overline{\mathbf{S_R}}(q) = \underset{S_g}{\mathbf{E}}\ \underset{L^q_{S_g}}{\mathbf{E}}\, \mathrm{I}^q_{S_g,S_g}$$

Example: Bayesian classifier

S=0	S=1	S_g=2	S=3	Index
1/9	1/9	2/3	1/9	1
1/3	1/3	1/3	0	1/3
1/8	1/2	1/4	1/8	0
1/5	1/5	2/5	1/5	1

$$\mathbf{S_R}(S_g = 2, q) \simeq 58\%$$

Information theoretic metric: mutual information

- Entropy matrix: $\mathrm{H}^q_{S_g,S} = \underset{L^q_{S_g}}{\mathbf{E}} - \log_2 \mathbf{P}[S|L^q_{S_g}]$
- Conditional entropy: $\mathbf{H}[S_g|L^q_{S_g}] = \underset{S_g}{\mathbf{E}} \ \mathrm{H}^q_{S_g,S_g}$
- Leakage matrix: $\Lambda^q_{S_g,S} = \mathbf{H}[S_g] - \mathrm{H}^q_{S_g,S}$
- Mutual information:

$$\mathbf{I}(S_g; L^q_{S_g}) = \mathbf{H}[S_g] - \mathbf{H}[S_g|L^q_{S_g}] = \underset{S_g}{\mathbf{E}} \Lambda^q_{S_g,S_g}$$

Example

$$\Lambda^{q}_{S_g,S} = 2 - \mathrm{H}^{q}_{S_g,S}$$

S=0	S=1	S_g=2	S=3
1/9	1/9	2/3	1/9
2/7	2/7	2/7	1/7
1/5	1/5	2/5	1/5
-0.43	-0.43	0.77	-0.76

- Definition: a leakage function is sound

$$\Longleftrightarrow \max_S \Lambda^q_{S_g,S} = \Lambda^q_{S_g,S_g},\ \forall\ S_g, q.$$

- If provided with a sound leakage function, a Bayesian adversary with unlimited queries to the target device will eventually be successful
  - Intuitive meaning: there is *enough* information in the side-channel observations

Combining security and IT metrics

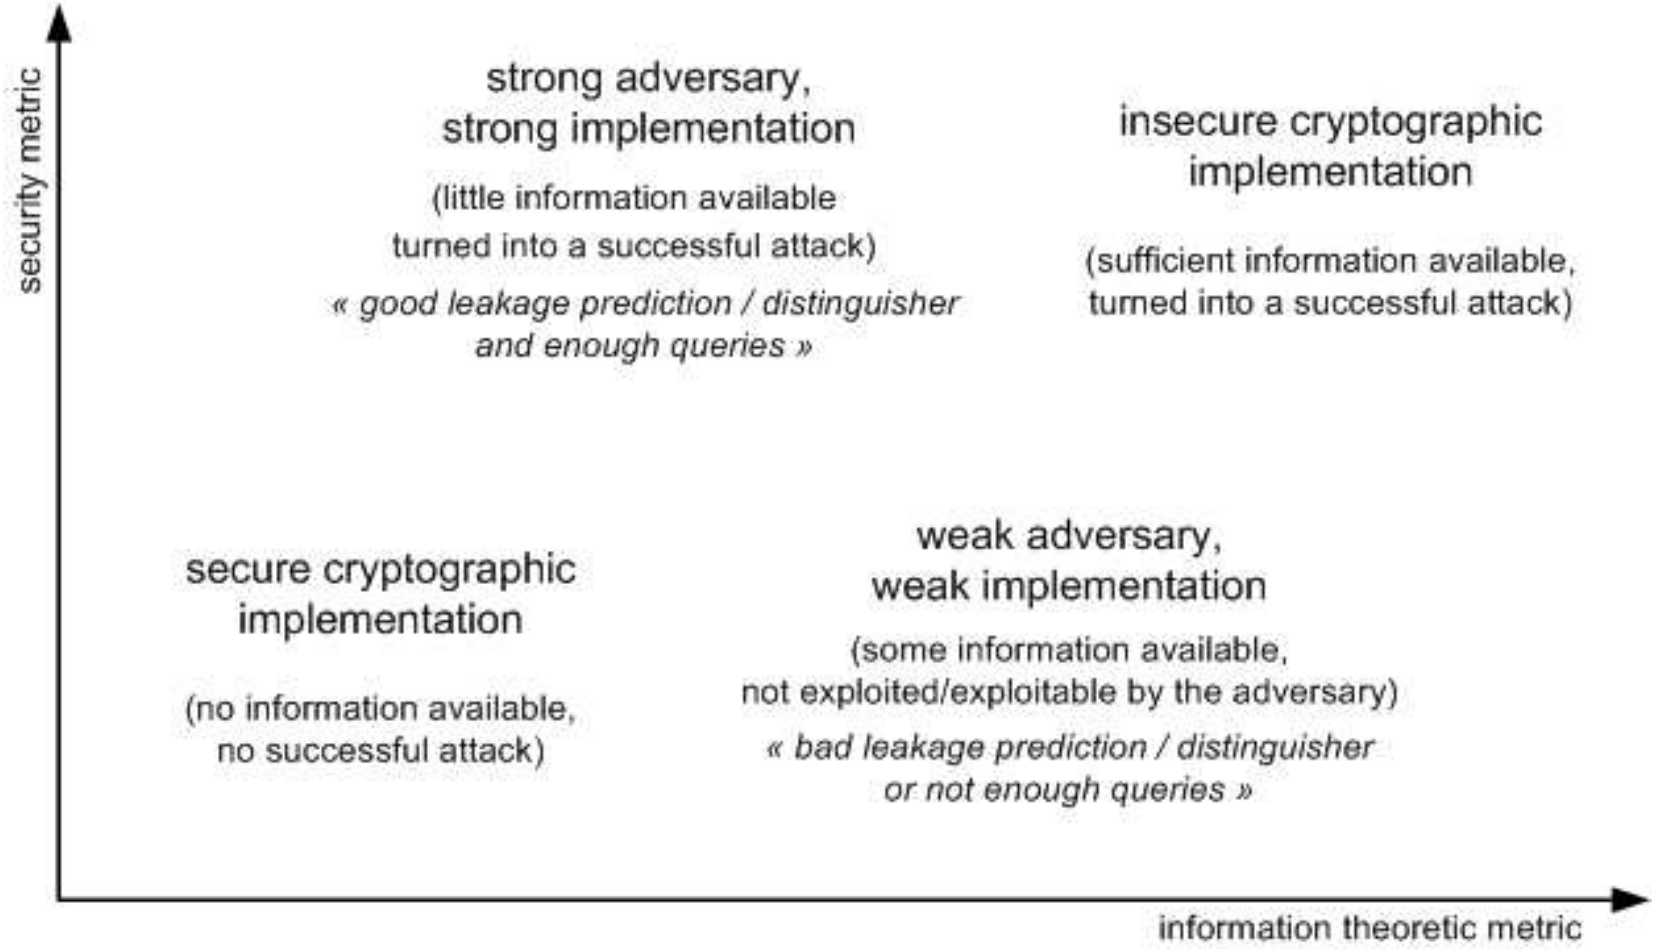

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

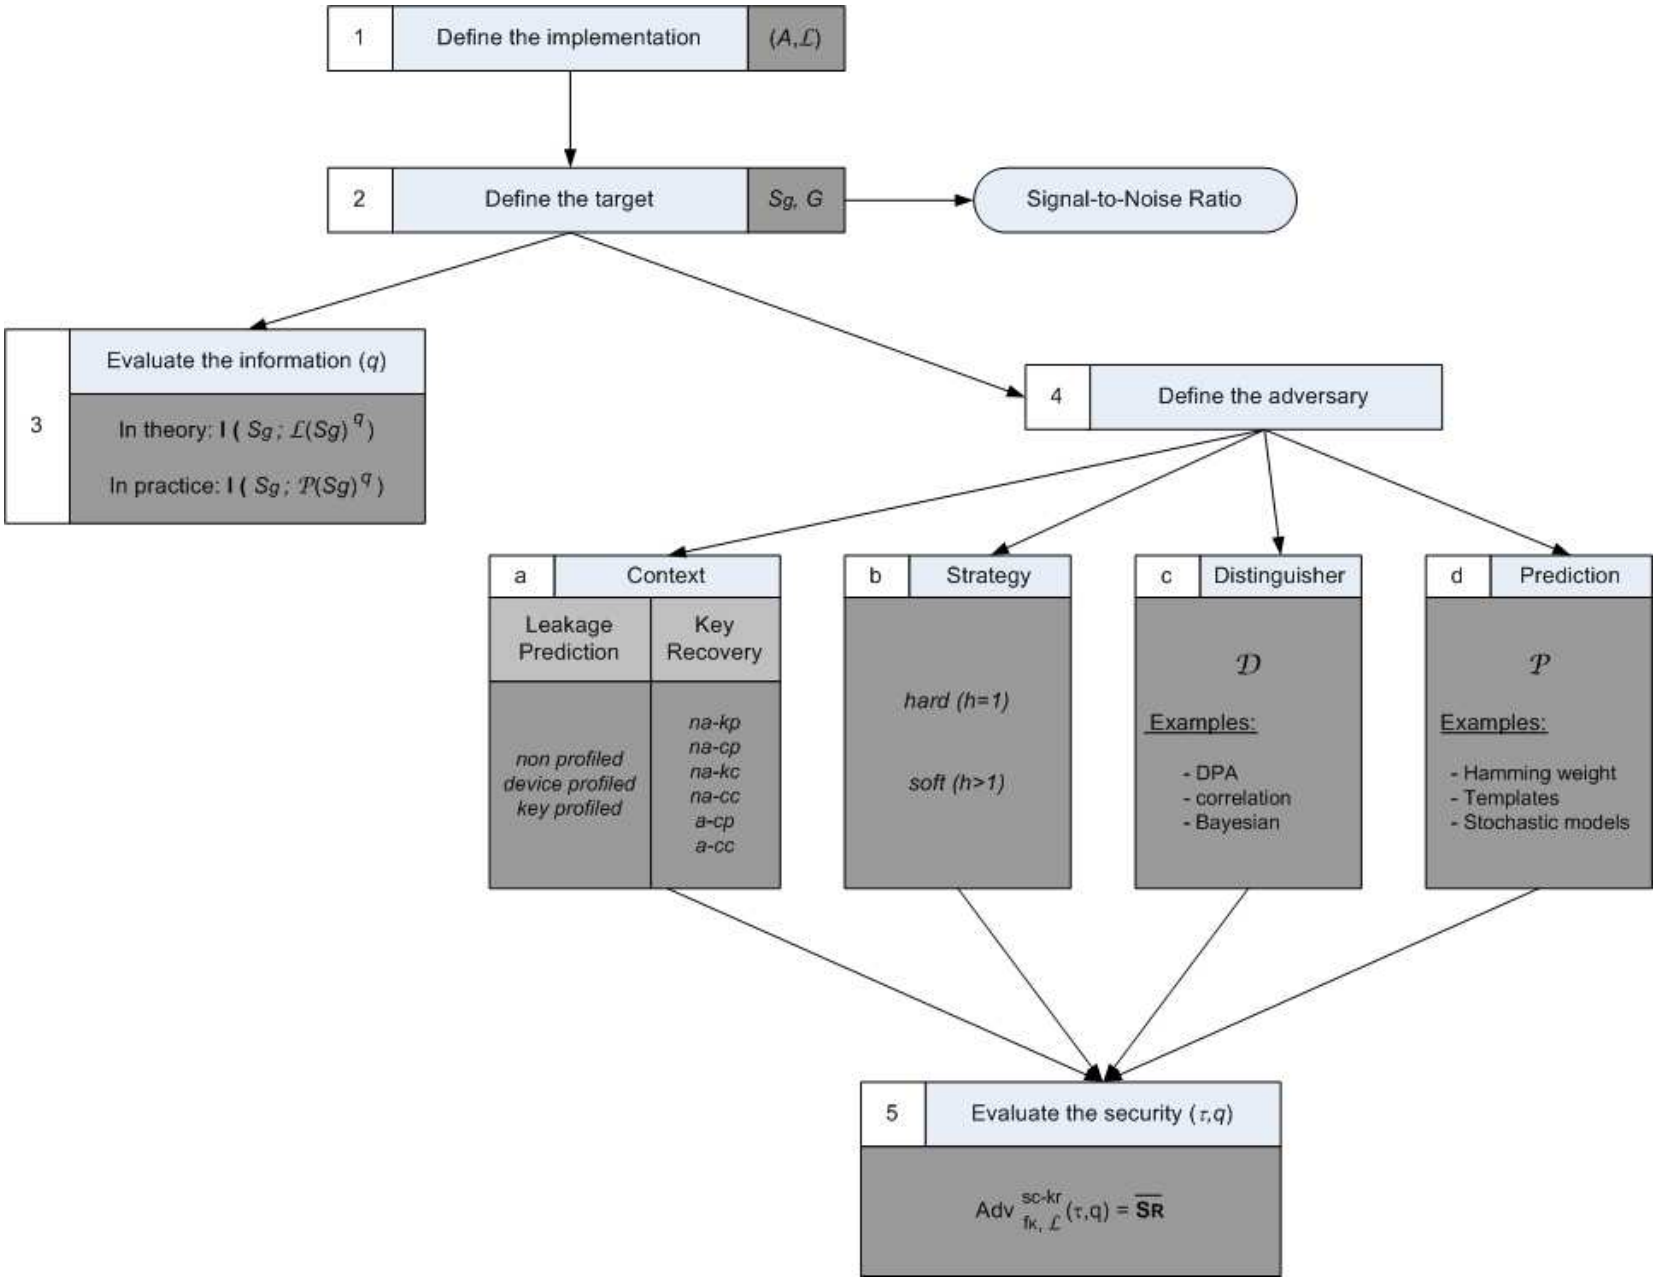
1
Define the implementation
(A,L)
2
Define the target
Sg, G
Signal-to-Noise Ratio
3
Evaluate the information (q)
In theory: I (Sg ; L(Sg)^q)
In practice: I (Sg ; P(Sg)^q)
4
Define the adversary
a
Context
Leakage Prediction
Key Recovery
non profiled
device profiled
key profiled
na-kp
na-cp
na-kc
na-cc
a-cp
a-cc
b
Strategy
hard (h=1)
soft (h>1)
c
Distinguisher
D
Examples:
- DPA
- correlation
- Bayesian
d
Prediction
P
Examples:
- Hamming weight
- Templates
- Stochastic models
5
Evaluate the security (τ,q)
Adv sc-kr fK, L (τ,q) = SR

Exemplary applications

- The security metric can be used to compare different adversaries

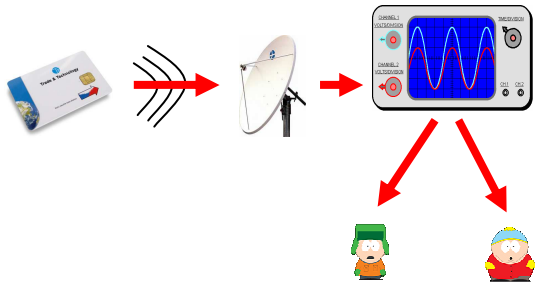

- The information theoretic metric can be used to compare different implementations

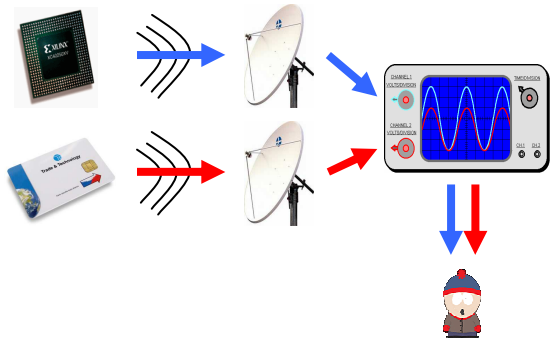

Outline

- Background: Micali & Reyzin model
- Target circuit
- Target block cipher
- Adversarial context & strategy
- Leakage predictions & adversary definition
- Computational restrictions
- Definition of security
- Evaluation criteria
- Evaluation methodology
- Side-channel tradeoffs
- Conclusions & open problems

Side-channel tradeoffs

- Flexibility *vs.* Efficiency

DPA … correlation … templates

- Information *vs.* Computation

hard strategy … soft strategy

Conclusions and open problems

- Introduction of a formal model for the analysis of cryptographic primitives against side-channel attacks as a specialization of Micali and Reyzin
- Both theoretical and practical contribution
- Open problems:
  - More analyzes of practical contexts
  - Design of provably secure implementations
  - Optimal leakage prediction functions / distinguishers

-THANKS-

Send comments to:

fstandae@uclouvain.be

More information on:

http://www.dice.ucl.ac.be/~fstandae/tsca/