A Unified Framework for the Analysis of Side-Channel Key-Recovery Attacks

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

UCL Crypto Group, Université catholique de Louvain
Dept. of Computer Science, Columbia University
Google Inc.

Eurocrypt 2009 - Cologne, Germany

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Introduction

- Side-channel attacks (the story made short)
  - Exploit the power consumption, electromagnetic radiation, . . . of a cryptographic implementation
  - Most of the times to recover keys
  - Powerful but device-specific ($\Rightarrow$ hard to evaluate)
  - Hard to prevent
  - Only a part of the physical reality
- Practical issues
  - "How to compare two implementations?"
  - "How to compare two adversaries?"

Goal of this framework: determine the extent to which these questions can be answered in a fair manner.

Example

- Evaluation and comparison of two implementations of the AES Rijndael (AES-CMOS and AES-WDDL)
- Tool: adversary A := { standard DPA, Hamming weight leakage model, target: one key byte }
  - $\mathbf{Succ}^{\text{sc-kr}}_{\mathrm{A}_{\text{AES-CMOS}}}(q, \ldots) = 0.9$ for $q = 10$
  - $\mathbf{Succ}^{\text{sc-kr}}_{\mathrm{A}_{\text{AES-WDDL}}}(q, \ldots) = 0.9$ for $q = 10000$

Is the lower success rate caused by a "secure implementation" or a "weak adversary"?

Introduction

- Limitations of previous (practical) works:
  - Mainly rely on heuristics
  - Use device-dependent metrics (*e.g.* variance)
  - Use adversary dependent metrics (*e.g.* correlation)

$\Rightarrow$ Separate the evaluation of the implementations from the evaluation of the side-channel adversaries

- Limitations of previous (theoretical) works
  - Hardly apply to actual implementations
  - Quantitative rather than qualitative

$\Rightarrow$ Propose a concrete evaluation methodology

A more friendly introduction

- Practice oriented provable security requires computational assumptions
  - *e.g.* the AES Rijndael is indistinguishable from a PRP for any polynomial-time adversary
- Leakage-resilient cryptography requires physical assumptions (*i.e.* bounded leakage, typically)
- This work attempts to provide foundations in order to determine what is a "reasonable physical assumption"
- Started from Micali & Reyzin (TCC)
- Ongoing research since 2004

A more friendly introduction

computational cryptography (since the 1980s)

$\Updownarrow$

computational cryptography + leakage (since $\sim$ 1994)

$\Updownarrow$

middleware between theory and practice

$\Updownarrow$

implementation issues (since the late 1990s)

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Terminology

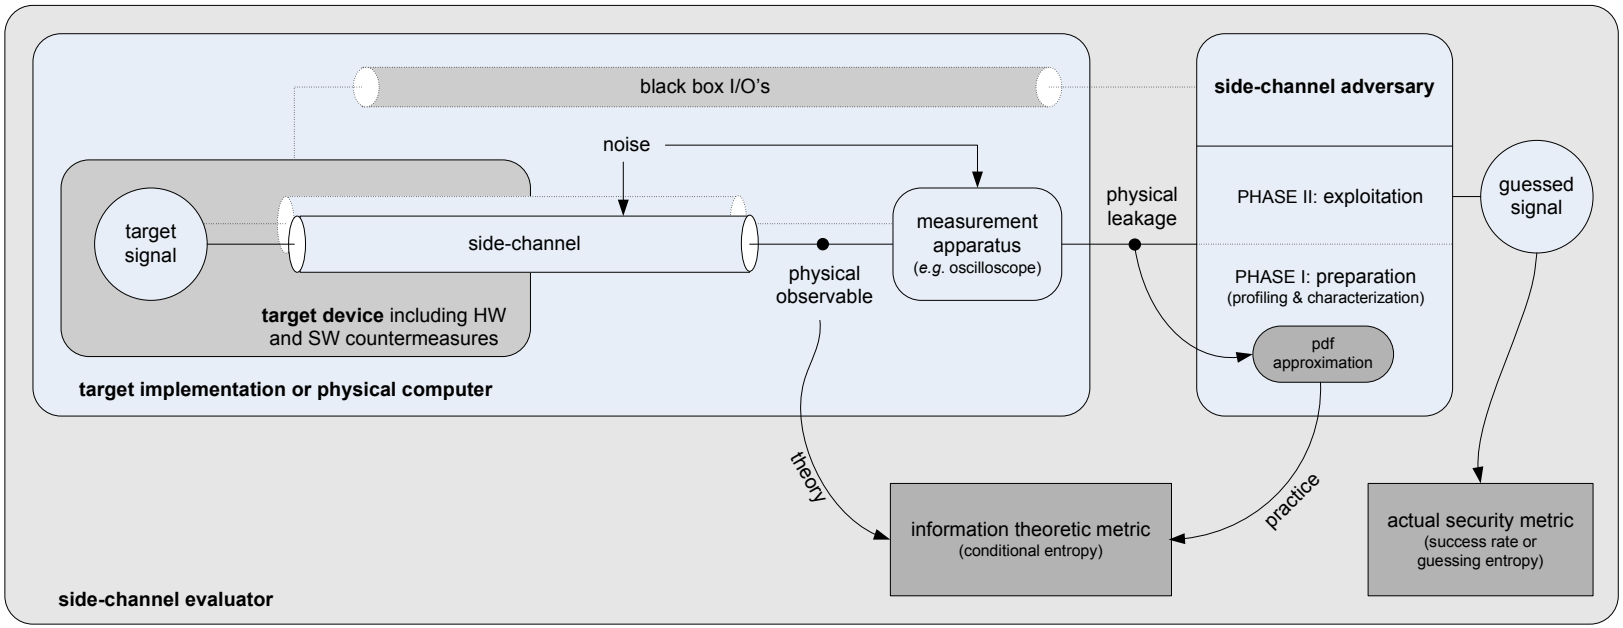

- primitive $\longrightarrow$ device
- device + side-channel + meas. setup = implementation
- (optional) preparation + exploitation = adversary

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Security metric # 1

o^{th} order success rate of the side-channel key recovery adversary $\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}$ against a key class variable S

Experiment $\mathbf{Exp}^{\text{sc-kr-}o}_{\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}}$

$k \xleftarrow{R} \mathcal{K}$;
$s = \delta(k)$;
$\mathbf{g} \leftarrow \mathsf{A}_{\mathsf{E}_k,\mathsf{L}}$;
if $s \in [g_1, \ldots, g_o]$ **then** return 1;
else return 0;

$$\mathbf{Succ}^{\text{sc-kr-}o,S}_{\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}}(\tau, m, q) = \Pr\ [\mathbf{Exp}^{\text{sc-kr-}o}_{\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}} = 1]$$

Security metric # 2

Guessing entropy of the side-channel key recovery adversary $\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}$ against a key class variable S

Experiment $\mathbf{Exp}^{\text{sc-kg}}_{\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}}$

$k \xleftarrow{R} \mathcal{K}$;

$s = \delta(k)$;

$\mathbf{g} \leftarrow \mathsf{A}_{\mathsf{E}_k,\mathsf{L}}$;

return i such that $g_i = s$;

$$\mathbf{GE}^{\text{sc-kr-}S}_{\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}}(\tau, m, q) = \mathbf{E}\left(\mathbf{Exp}^{\text{sc-kg}}_{\mathsf{A}_{\mathsf{E}_K,\mathsf{L}}}\right)$$

Information theoretic metric

Conditional entropy matrix

$$\mathbf{H}^{q}_{s,s^*} = -\sum_{\mathbf{l}_q} \Pr[\mathbf{l}_q|s] \cdot \log_2 \Pr[s^*|\mathbf{l}_q],$$

Shannon's conditional entropy

$$\mathrm{H}[S|\mathbf{L}_q] = -\sum_{s} \Pr[s] \sum_{\mathbf{l}_q} \Pr[\mathbf{l}_q|s] \cdot \log_2 \Pr[s|\mathbf{l}_q] = \underset{s}{\mathbf{E}} \ \mathbf{H}^{q}_{s,s}$$

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Practical limitations

- Computing $\mathrm{H}[S|\mathbf{L}_q]$ requires the knowledge of $\Pr[\mathbf{L}_q|S]$
  - Issue 1: the leakage distribution is generally unknown

  $\Rightarrow$ The IT metric has to be approximated
  - Issue 2: leakages generally have lots of samples

  $\Rightarrow$ We have to consider the approximated leakage distribution of a reduced set of samples

- In other words, we need to use generic template attacks (*e.g.* PCA-based, using a Gaussian assumption, stochastic models, ...)

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

1. Asymptotic meaning of $H[S|\mathbf{L}_q]$

"Can I approximate the leakage probability distribution?"

Definition 1. Asymptotic success rate of a side-channel key recovery adversary: $\mathbf{Succ}^{\text{sc-kr-}o,S}_{\mathsf{A}_{\mathsf{E}_K},\mathsf{L}}(q \rightarrow \infty)$

Definition 2. Bayesian side-channel key recovery adversary: selects $\tilde{s} = \mathit{argmax}_{s^*} \Pr[s^*|\mathbf{l}_q]$

Definition 3. Sound leakage probability distribution $\Pr[\mathbf{L}_q|S]$ or approximation $\hat{\Pr}[\tilde{\mathbf{L}}_q|S]$: if the first-order asymptotic success rate $\mathbf{Succ}^{\text{sc-kr-1},S}_{\mathsf{A}_{\mathsf{E}_K},\mathsf{L}}(q \rightarrow \infty) = 1$

1. Asymptotic meaning of $H[S|\boldsymbol{L}_q]$

Bounded preparation / unbounded exploitation:

$$\mathbf{H}^{q}_{s,s^*} = \begin{pmatrix} h_{1,1} & h_{1,2} & \dots & h_{1,|\mathcal{S}|} \\ h_{2,2} & h_{2,2} & \dots & h_{2,|\mathcal{S}|} \\ \dots & \dots & \dots & \dots \\ h_{|\mathcal{S}|,1} & h_{|\mathcal{S}|,2} & \dots & h_{|\mathcal{S}|,|\mathcal{S}|} \end{pmatrix}$$

Theorem 1. *(...) a leakage probability distribution is sound if and only if* $\text{argmin}_{s^*} \mathbf{H}^{1}_{s,s^*} = s, \forall s \in \mathcal{S}$

Intuitively: the diagonal elements $h_{s,s}$'s are minimum

Example (AES Rijndael)

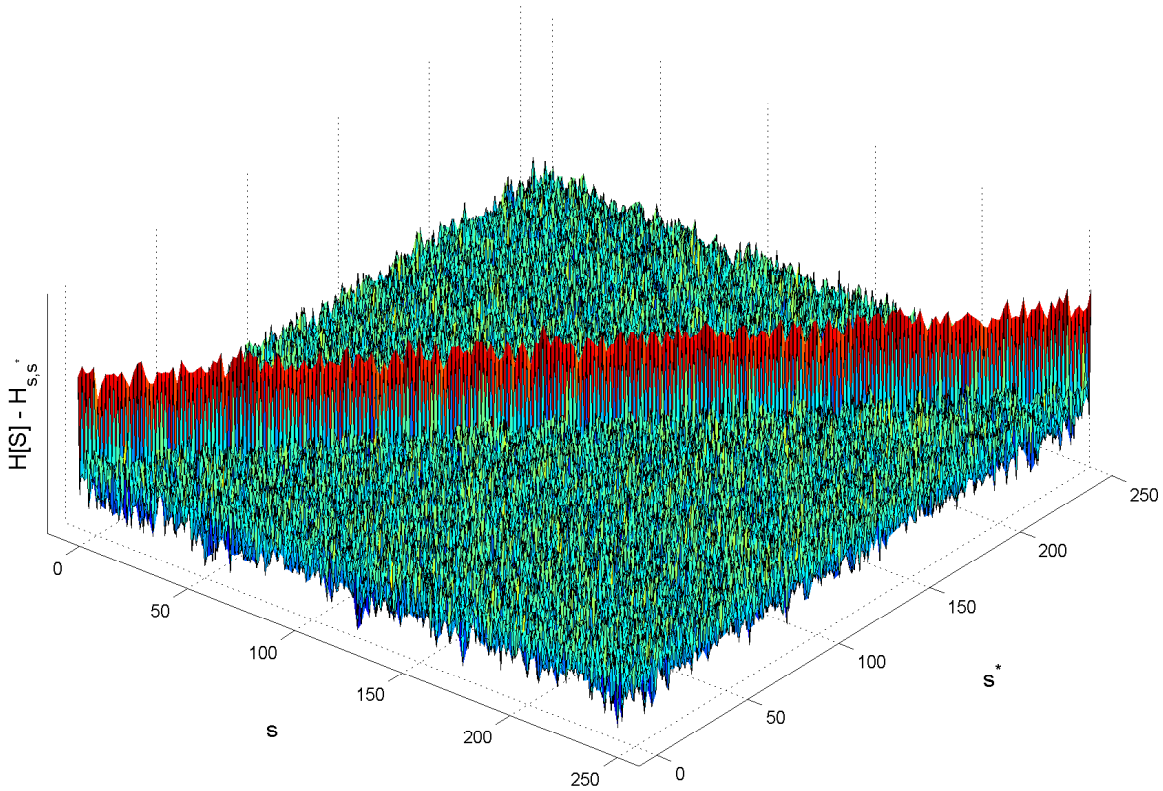

2. Comparative meaning of $H[S|\mathbf{L}_q]$

"Does more entropy imply more security?"

$$\mathbf{H}^q_{s,s^*} = \begin{pmatrix} h_{1,1} & h_{1,2} & \dots & h_{1,|\mathcal{S}|} \\ h_{2,2} & h_{2,2} & \dots & h_{2,|\mathcal{S}|} \\ \dots & \dots & \dots & \dots \\ h_{|\mathcal{S}|,1} & h_{|\mathcal{S}|,2} & \dots & h_{|\mathcal{S}|,|\mathcal{S}|} \end{pmatrix}$$

$h_{s,s}$: residual entropy of a key class s

$\mathrm{H}[S|\mathbf{L}_q] = \mathbf{E}_s \ \mathbf{H}^q_{s,s}$ (averaged diagonal of $\mathbf{H}^q_{s,s^*}$)

2. Comparative meaning of $H[S|\boldsymbol{L}_q]$

Definition 4. $|\mathcal{S}|$-target side-channel attack: tries to identify one key candidate out of $|\mathcal{S}|$

Definition 5. Gaussian leakage distribution: such that $\mathrm{L}(C_\alpha, M, R) = \mathrm{L'}(C_\alpha, M) + \mathrm{L''}(R)$, $\mathrm{L''}(R) =$ gaussian noise.

Definition 6. Ideal side-channel attack: Bayesian attack in which the leakages are perfectly predicted by the adversary's approximated probability density function.

2. Comparative meaning of $H[S|L_q]$

Unbounded preparation / bounded exploitation

- Does more entropy imply more security?
  - Ideal 2-target attacks with Gaussian leakages: yes
  - Ideal $|\mathcal{S}|$-target attacks with "perfect" leakages: yes
  - In general: no
    (a pdf cannot be summarized in a scalar value)

  - In practice?

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Implementation dependencies

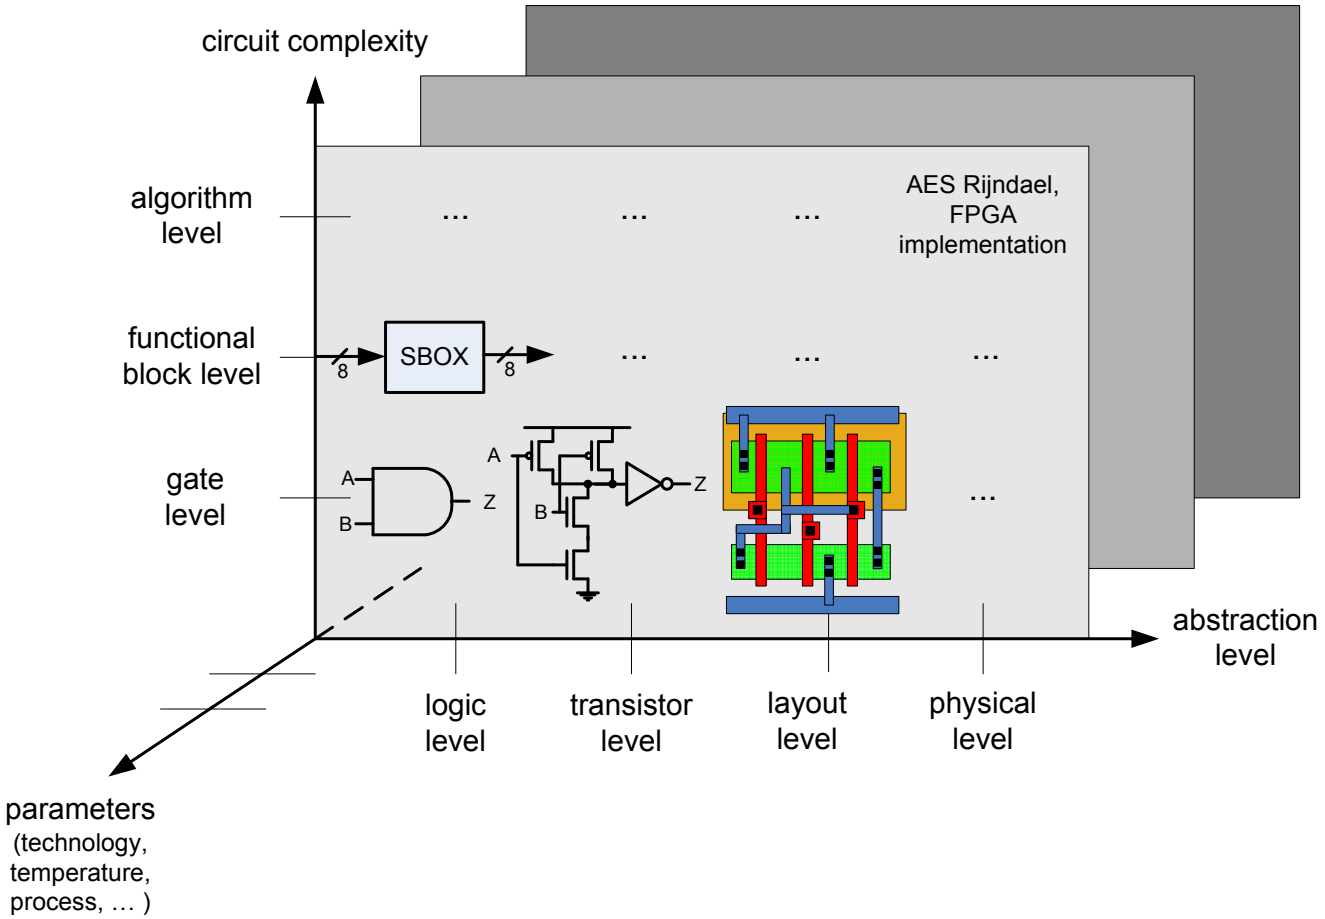

Comparing masking schemes (CHES 2006)

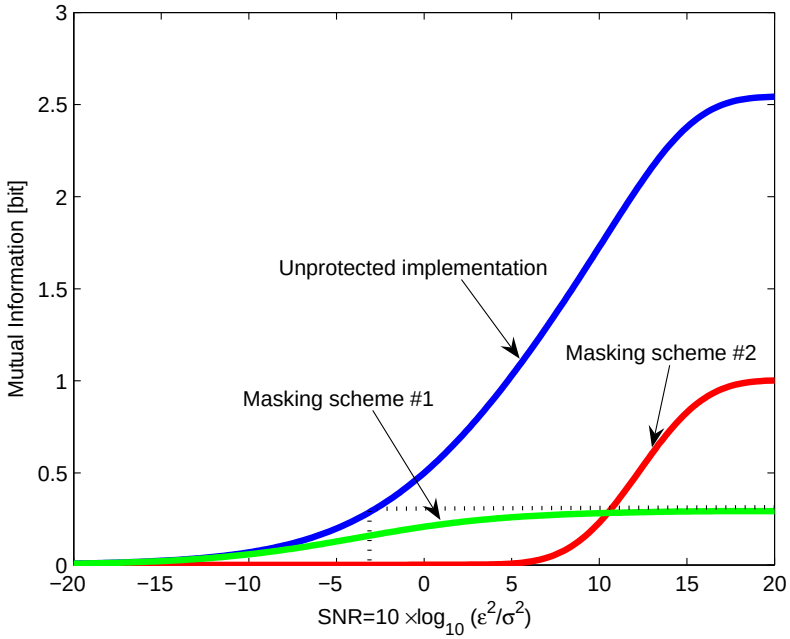

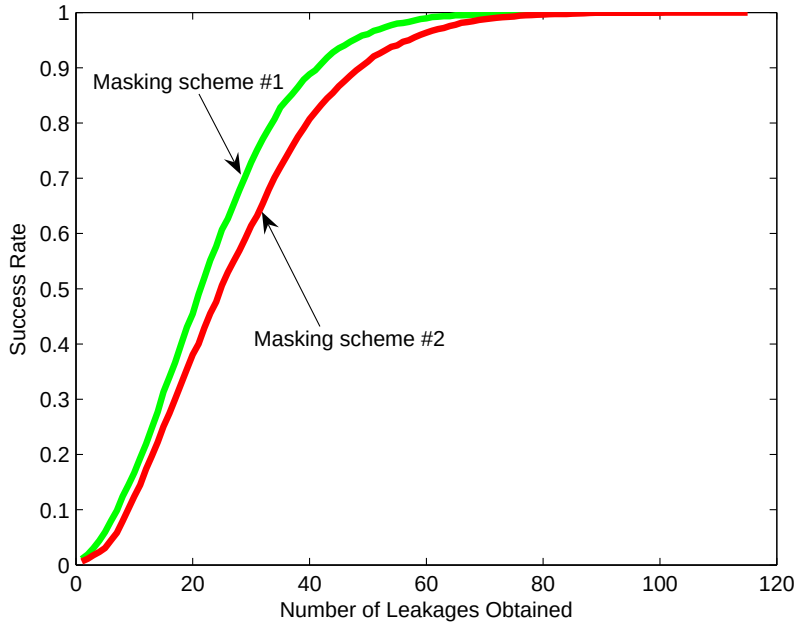

- SNR=10 -

Comparing masking schemes (CHES 2006)

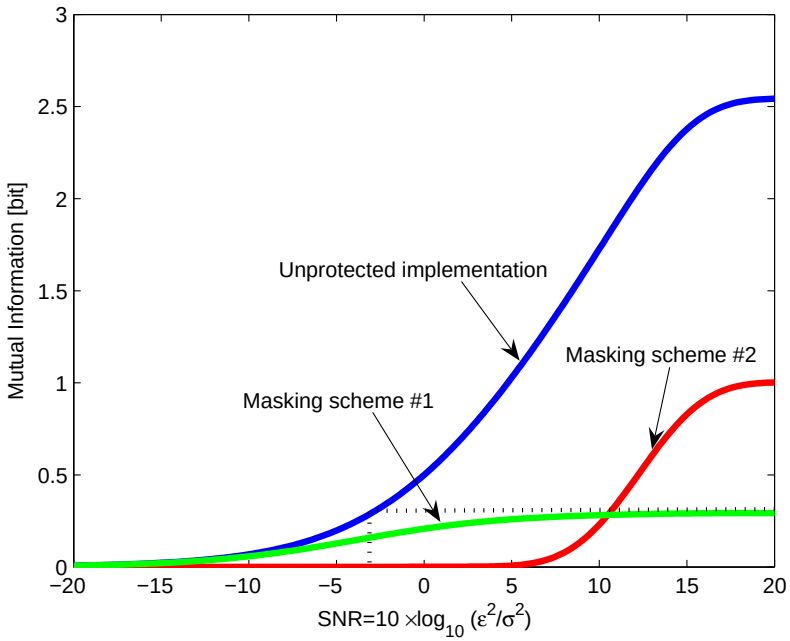

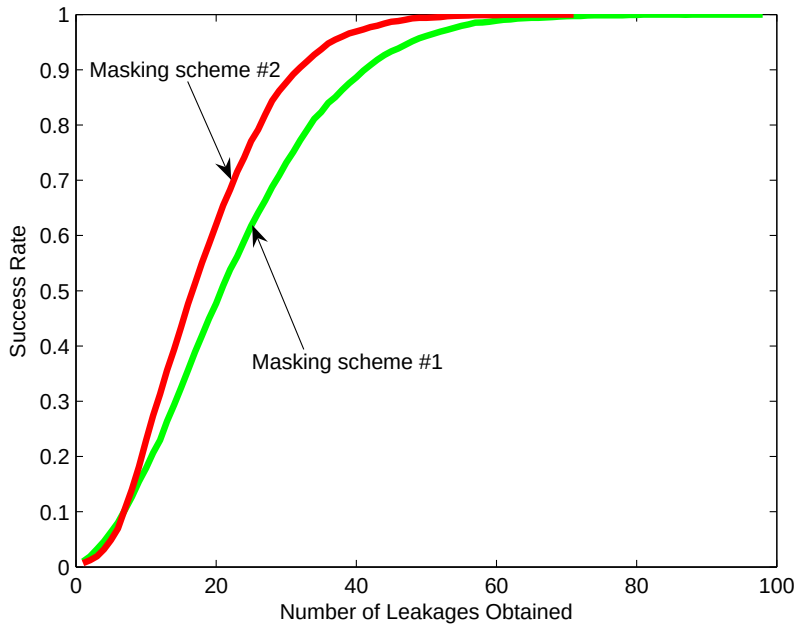

- SNR=11 -

Other experimental validations

- Comparison of different side-channel resistant logic styles from SPICE simulations (CHES 2007)
- Comparison of power and EM leakages using PCA/LDA from real measurements (CHES 2008)
- Experimental evaluation of various side-channel distinguishers in two microcontrollers (ICISC 2008)
- Evaluation of the profiling efficiency of template attacks and stochastic models (ACNS 2009)

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Evaluation methodology

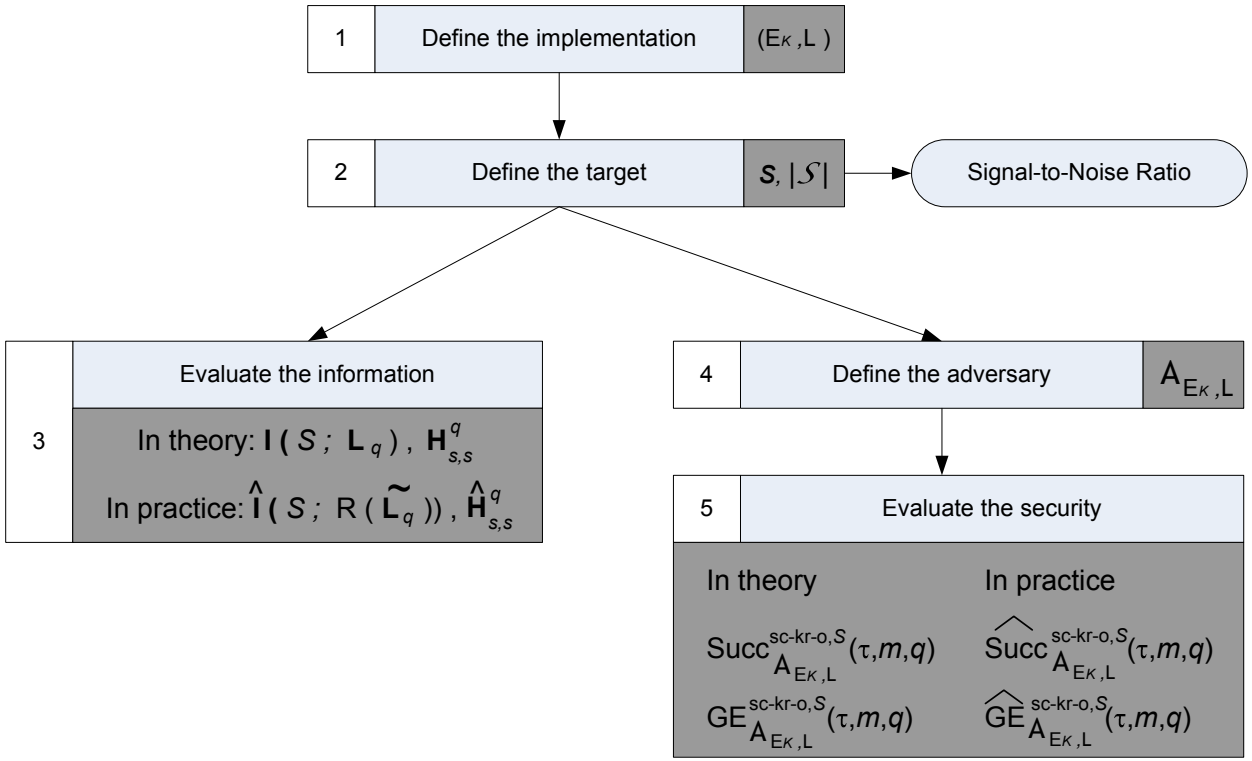

- Side-channel attacks ≈ statistical sampling problem

A last remark

- Side-channel attacks are an implementation problem
- Performances (and constants) are important !
- It is easy to build provably secure (but expensive) implementations, *e.g.* the AES as a 2^{128} table

- We need to trade efficiency for security on a fair basis
- We hope this work can be used as a fair basis

Outline

1. Introduction
2. Terminology
3. Formal definitions
4. Practical limitations
5. Relations between the metrics
6. Applications of the model
7. Evaluation methodology
8. Conclusions

Conclusions

leakage resilient PRGs (*e.g.* ASIACCS08, EUROCRYPT09)

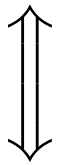

↑ sound assumptions

middleware between theory and practice

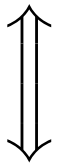

↓ fair evaluation

side-channel attacks & countermeasures

(Kocher's DPA, masking schemes, dual-rail logic styles,...)

- Side-channel attacks $\approx$ cryptanalytic problem
  - Having provably secure encryption modes do not remove the need of block cipher cryptanalysis !

THANKS

Questions?

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