


High (**Physical**) Security & Lightweight (**Symmetric**) Cryptography



François-Xavier Standaert








UCL Crypto Group, Belgium

HIGHLIGHT – LIGHTCRYPTO, November 2016

Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- Mitigating hardware defaults (is hard)
- Transparency is needed (open source)
- Summary and conclusions

Outline

- Preliminary questions / definitions 
- Side-channel basics (attack steps) 
- Noise (aka hardware) is not enough 
- Noise amplification (aka masking) 
- Reductions help (aka leakage resilience) 
- Mitigating hardware defaults (is hard) 
- Transparency is needed (open source) 
- Summary and conclusions

intro secure design eval.

Outline

- **Preliminary questions / definitions**
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- Mitigating hardware defaults (is hard)
- Transparency is needed (open source)
- Summary and conclusions

- For block ciphers, best attack has:
 - Time complexity 2^{n_1}
 - Data complexity 2^{n_2}
 - Memory complexity 2^{n_3}

- For block ciphers, best attack has:
 - Time complexity 2^{n_1}
 - Data complexity 2^{n_2}
 - Memory complexity 2^{n_3}
- With typical security parameters:
 - $n_L = 80, n_S = 128, n_{PQ} = 256$
 - <https://www.keylength.com/en/>
- Function of the algorithms' deployment time and the adversary's computational power

- When also considering physical attacks:
 - Measurement complexity 2^{m_1}
 - Fault complexity 2^{m_2}
- Typical security parameters: ???

- When also considering physical attacks:
 - Measurement complexity 2^{m_1}
 - Fault complexity 2^{m_2}
- Typical security parameters: ???

Q1. Deployment time of implementations?

- IMO not much less (minimum 5-10 years)
- So we can gain a factor 2 to 4 (i.e., 1,2 bits)

Yuanyuan Zhou, Yu Yu, François-Xavier Standaert, Jean-Jacques Quisquater: *On the Need of Physical Security for Small Embedded Devices: A Case Study with COMP128-1 Implementations in SIM Cards*. Financial Cryptography 2013: 230-238.
Junrong Liu, Yu Yu, François-Xavier Standaert, Zheng Guo, Dawu Gu, Wei Sun, Yijie Ge, Xinjun Xie: *Small Tweaks Do Not Help: Differential Power Analysis of MILENAGE Implementations in 3G/4G USIM Cards*. ESORICS (1) 2015: 468-480

Q2. What's the adversary's measurement power?

Q2. What's the adversary's measurement power?

- Generic answer:
 - “Cost” of collecting one side-channel meas. vs. cost of collecting one pt/ct pair?
 - Min. $\times 2^{10}$, avg. $\times 2^{20}$, opt. $\times 2^{30}$
 - Roughly assume \approx the same for faults
 - So we can gain 10 to 30 bits (roughly)
- $\Rightarrow m_L > 60$ and $m_S > 100$ (no clue about m_{PQ})

- Specific answer (if physical access is limited):
 - $m_{VL} > 40$ (\approx some days/weeks)
 - $m_{UL} > 20$ (\approx some hours)
 - Excluding network access (timing attacks)!

- Specific answer (if physical access is limited):
 - $m_{VL} > 40$ (\approx some days/weeks)
 - $m_{UL} > 20$ (\approx some hours)
 - Excluding network access (timing attacks)!
- **Fact 1. Currently, we mostly design for VL/UL physical security / Fact 2. Current physical security evaluations are limited to VL security**


- Specific answer (if physical access is limited):
 - $m_{VL} > 40$ (\approx some days/weeks)
 - $m_{UL} > 20$ (\approx some hours)
 - Excluding network access (timing attacks)!
- Fact 1. Currently, we mostly design for VL/UL physical security / Fact 2. Current physical security evaluations are limited to VL security





- **VL or UL security \approx no security**


Outline


- Preliminary questions / definitions
- **Side-channel basics (attack steps)**
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- Mitigating hardware defaults (is hard)
- Transparency is needed (open source)
- Summary and conclusions











- Noise reduction via good setups
- Filtering, averaging (FFT, SSA, ...)
- Detection of Points-Of-Interest (POI)
- Dimensionality reduction (PCA, LDA,...)
- ...

Victor Lomné, Emmanuel Prouff, Thomas Roche: *Behind the Scene of Side Channel Attacks*. ASIACRYPT (1) 2013: 506-525. Santos Merino Del Pozo, François-Xavier Standaert: *Blind Source Separation from Single Measurements Using Singular Spectrum Analysis*. CHES 2015: 42-59. Oscar Reparaz, Benedikt Gierlichs, Ingrid Verbauwhede: *Selecting Time Samples for Multivariate DPA Attacks*. CHES 2012: 155-174. François Durvaux, François-Xavier Standaert: *From Improved Leakage Detection to the Detection of Points of Interests in Leakage Traces*. EUROCRYPT (1) 2016: 240-262 - 50. Cédric Archambeau, Eric Peeters, François-Xavier Standaert, Jean-Jacques Quisquater: *Template Attacks in Principal Subspaces*. CHES 2006: 1-14. François-Xavier Standaert, Cédric Archambeau: *Using Subspace-Based Template Attacks to Compare and Combine Power and Electromagnetic Information Leakages*. CHES 2008: 411-425

- Noise reduction via good setups
- Filtering, averaging (FFT, SSA, ...)
- Detection of Points-Of-Interest (POI)
- Dimensionality reduction (PCA, LDA,...)
- ...
- Inherently heuristic (!) \Rightarrow hard to determine what is the optimal solution (\neq next steps)

Victor Lomné, Emmanuel Prouff, Thomas Roche: *Behind the Scene of Side Channel Attacks*. ASIACRYPT (1) 2013: 506-525. Santos Merino Del Pozo, François-Xavier Standaert: *Blind Source Separation from Single Measurements Using Singular Spectrum Analysis*. CHES 2015: 42-59. Oscar Reparaz, Benedikt Gierlichs, Ingrid Verbauwhede: *Selecting Time Samples for Multivariate DPA Attacks*. CHES 2012: 155-174. François Durvaux, François-Xavier Standaert: *From Improved Leakage Detection to the Detection of Points of Interests in Leakage Traces*. EUROCRYPT (1) 2016: 240-262 - 50. Cédric Archambeau, Eric Peeters, François-Xavier Standaert, Jean-Jacques Quisquater: *Template Attacks in Principal Subspaces*. CHES 2006: 1-14. François-Xavier Standaert, Cédric Archambeau: *Using Subspace-Based Template Attacks to Compare and Combine Power and Electromagnetic Information Leakages*. CHES 2008: 411-425

- General case: profiled DPA
 - Build “*templates*”, i.e., $\hat{f}(l_i|k, x_i)$
 - e.g. Gaussian, regression-based
 - Which directly leads to $\widehat{\Pr}[k|l_i, x_i]$

- General case: profiled DPA
 - Build “*templates*”, i.e., $\hat{f}(l_i|k, x_i)$
 - e.g. Gaussian, regression-based
 - Which directly leads to $\widehat{\Pr}[k|l_i, x_i]$
- “Simplified” case: non-profiled DPA
 - Just assumes some model
 - e.g., $m_i^{k^*} = \text{HW}(z_i)$

- General case: profiled DPA
 - Build “*templates*”, i.e., $\hat{f}(l_i | k, x_i)$
 - e.g. Gaussian, regression-based
 - Which directly leads to $\widehat{\Pr}[k | l_i, x_i]$
- “Simplified” case: non-profiled DPA
 - Just assumes some model
 - e.g., $m_i^{k^*} = \text{HW}(z_i)$
- Separation: only profiled DPA is guaranteed to succeed against any leaking device (!)

- Profiled case: maximum likelihood


- Profiled case: maximum likelihood
- Unprofiled case:
 - Difference-of-Means
 - Correlation (CPA)
 - « On-the-fly » regression
 - Mutual Information Analysis (MIA)

Omar Choudary, Markus G. Kuhn: *Efficient Template Attacks*. CARDIS 2013: 253-270. Paul C. Kocher, Joshua Jaffe, Benjamin Jun: *Differential Power Analysis*. CRYPTO 1999: 388-397. Eric Brier, Christophe Clavier, Francis Olivier: *Correlation Power Analysis with a Leakage Model*. CHES 2004: 16-29. Julien Doget, Emmanuel Prouff, Matthieu Rivain, François-Xavier Standaert: *Univariate side channel attacks and leakage modeling*. J. Cryptographic Engineering 1(2): 123-144 (2011). Lejla Batina, Benedikt Gierlichs, Emmanuel Prouff, Matthieu Rivain, François-Xavier Standaert, Nicolas Veyrat-Charvillon: *Mutual Information Analysis: a Comprehensive Study*. J. Cryptology 24(2): 269-291 (2011)

- Profiled case: maximum likelihood
- Unprofiled case:
 - Difference-of-Means
 - Correlation (CPA)
 - « On-the-fly » regression
 - Mutual Information Analysis (MIA)
- **Advanced topic: analytical (algebraic) attacks**

Omar Choudary, Markus G. Kuhn: *Efficient Template Attacks*. CARDIS 2013: 253-270. Paul C. Kocher, Joshua Jaffe, Benjamin Jun: *Differential Power Analysis*. CRYPTO 1999: 388-397. Eric Brier, Christophe Clavier, Francis Olivier: *Correlation Power Analysis with a Leakage Model*. CHES 2004: 16-29. Julien Doget, Emmanuel Prouff, Matthieu Rivain, François-Xavier Standaert: *Univariate side channel attacks and leakage modeling*. J. Cryptographic Engineering 1(2): 123-144 (2011). Lejla Batina, Benedikt Gierlichs, Emmanuel Prouff, Matthieu Rivain, François-Xavier Standaert, Nicolas Veyrat-Charvillon: *Mutual Information Analysis: a Comprehensive Study*. J. Cryptology 24(2): 269-291 (2011). Nicolas Veyrat-Charvillon, Benoît Gérard, François-Xavier Standaert: *Soft Analytical Side-Channel Attacks*. ASIACRYPT (1) 2014: 282-296

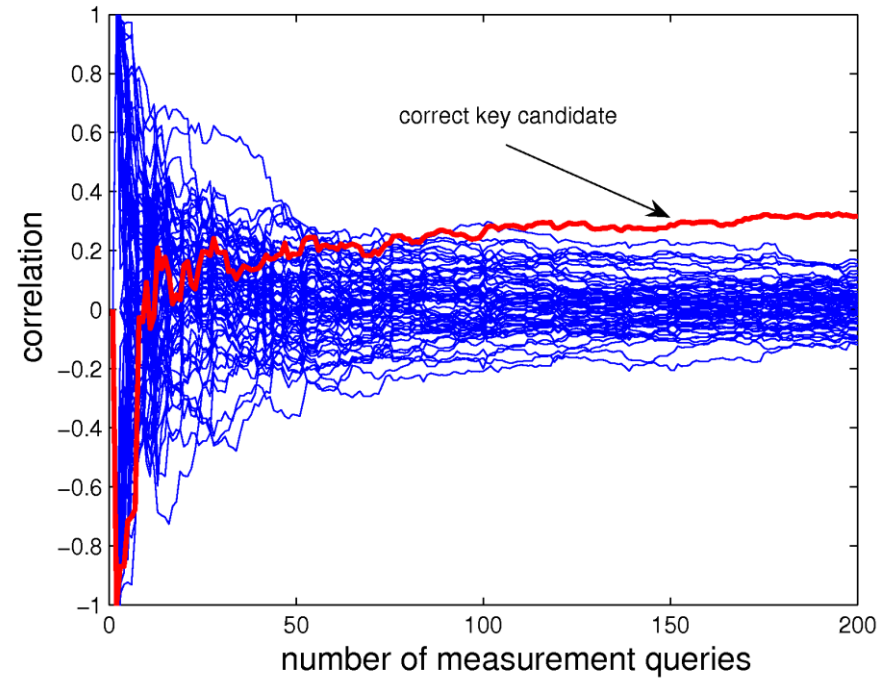
Gaussian templates



$$\tilde{k} = \operatorname{argmax}_{k^*} \prod_{i=1}^q \frac{1}{\sqrt{2 \cdot \pi \cdot \sigma(L)}} \cdot \exp\left(-\frac{1}{2} \cdot \left(\frac{l_i - m_i^{k^*}}{\sigma(L)}\right)^2\right)$$

- More efficient (**why?**)
- Outputs probabilities

CPA



$$\tilde{k} = \operatorname{argmax}_{k^*} \frac{E(L \cdot M^{k^*}) - E(L) \cdot E(M^{k^*})}{\sigma(L) \cdot \sigma(M^{k^*})}$$

- Less efficient (**why?**)
- Outputs scores

- CPA: $\tilde{k} = \operatorname{argmax}_{k^*} \frac{E(L \cdot M^{k^*}) - E(L) \cdot E(M^{k^*})}{\sigma(L) \cdot \sigma(M^{k^*})}$

- CPA: $\tilde{k} = \operatorname{argmax}_{k^*} \frac{E(L \cdot M^{k^*}) - E(L) \cdot E(M^{k^*})}{\sigma(L) \cdot \sigma(M^{k^*})}$ = 0 (normalization)

- CPA: $\tilde{k} = \operatorname{argmax}_{k^*} \frac{E(L \cdot M^{k^*}) - E(L) \cdot E(M^{k^*})}{\sigma(L) \cdot \sigma(M^{k^*})}$ = 0 (normalization)
independent of k^*

- CPA:
$$\tilde{k} = \operatorname{argmax}_{k^*} \frac{E(L \cdot M^{k^*}) - E(L) \cdot E(M^{k^*})}{\sigma(L) \cdot \sigma(M^{k^*})} = 0 \text{ (normalization)}$$

independent of k^* asymptotically independent of k^*

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
- Gaussian templates:

$$\tilde{k} = \operatorname{argmax}_{k^*} \prod_{i=1}^q \frac{1}{\sqrt{2 \cdot \pi \cdot \sigma(L)}} \cdot \exp\left(-\frac{1}{2} \cdot \left(\frac{l_i - m_i^{k^*}}{\sigma(L)}\right)^2\right)$$

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
- Gaussian templates:

$$\tilde{k} = \operatorname{argmax}_{k^*} \prod_{i=1}^q \frac{1}{\sqrt{2 \cdot \pi \cdot \sigma(L)}} \cdot \exp\left(-\frac{1}{2} \cdot \left(\frac{l_i - m_i^{k^*}}{\sigma(L)}\right)^2\right)$$

independent of k^*

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
- Gaussian templates:

$$\tilde{k} \propto \operatorname{argmax}_{k^*} \prod_{i=1}^q \exp \left(-\frac{1}{2} \cdot \left(\frac{l_i - m_i^{k^*}}{\sigma(L)} \right)^2 \right)$$

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$

- Gaussian templates:

$$\tilde{k} \propto \operatorname{argmin}_{k^*} E(L^2) - 2 \cdot E(L \cdot M^{k^*}) + E((M^{k^*})^2)$$

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$

- Gaussian templates:

$$\tilde{k} \propto \operatorname{argmin}_{k^*} E(L^2) - 2 \cdot E(L \cdot M^{k^*}) + E((M^{k^*})^2)$$

independent of k^*

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$

- Gaussian templates:

$$\tilde{k} \propto \operatorname{argmin}_{k^*} E(L^2) - 2 \cdot E(L \cdot M^{k^*}) + E((M^{k^*})^2)$$

independent of k^*

asymptotically
independent of k^*

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
- Gaussian templates: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
 - Gaussian templates: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
- ⇒ Both attacks are asymptotically equivalent
- For 1st-order leakages
 - i.e., unprotected implementations
 - Given they exploit the same model

- CPA: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
 - Gaussian templates: $\tilde{k} \propto \operatorname{argmax}_{k^*} E(L \cdot M^{k^*})$
- ⇒ Both attacks are asymptotically equivalent
- For 1st-order leakages
 - i.e., unprotected implementations
 - Given they exploit the same model
- ⇒ Gaussian templates outperforms CPA because it (usually) exploits a better (profiled) **model**

Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- **Noise (aka hardware) is not enough**
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- Mitigating hardware defaults (is hard)
- Transparency is needed (open source)
- Summary and conclusions

- **Lemma 1.** The mutual information between two normally distributed random variables X, Y with means μ_X, μ_Y and variances σ_X^2, σ_Y^2 equals:

$$\text{MI}(X; Y) = -\frac{1}{2} \log_2(1 - \rho(X, Y)^2)$$

- **Lemma 1.** The mutual information between two normally distributed random variables X, Y with means μ_X, μ_Y and variances σ_X^2, σ_Y^2 equals:

$$\text{MI}(X; Y) = -\frac{1}{2} \log_2(1 - \rho(X, Y)^2)$$

- **Lemma 2.** In a CPA, the number of samples required to distinguish the correct key with model M_k from the other key candidates with models M_{k^*} is $\propto \frac{c}{\rho(M_k, L)^2}$ (with c a small constant depending on the SR & # of key candidates)

- **Lemma 3.** Let X, Y and L be three random variables s.t. $Y = X + N_1$ and $L = Y + N_2$ with N_1 and N_2 two additive noise variables. Then:

$$\rho(X, L) = \rho(X, Y) \cdot \rho(Y, L)$$

- **Lemma 3.** Let X, Y and L be three random variables s.t. $Y = X + N_1$ and $L = Y + N_2$ with N_1 and N_2 two additive noise variables. Then:

$$\rho(X, L) = \rho(X, Y) \cdot \rho(Y, L)$$

- **Lemma 4.** The correlation coefficient between the sum of n independent and identically distributed random variables and the sum of the first $m < n$ of these equals $\sqrt{m/n}$

- FPGA implementation of the AES
- Adversary targeting the 1st byte of key
- Hamming weight leakage function/model
- 8-bit loop architecture broken in 10 traces

- FPGA implementation of the AES
- Adversary targeting the 1st byte of key
- Hamming weight leakage function/model
- 8-bit loop architecture broken in 10 traces

- How does the attack data complexity scale
 - For a 32-bit architecture?
 - i.e., with 24 bits of « algorithmic noise »
 - For a 128-bit architecture?
 - i.e., with 120 bits of « algorithmic noise »

- Hint: $L = M + N = (M_P + M_U) + N$

- Hint: $L = M + N = (M_P + M_U) + N$
- Lemma 3: $\rho(M_P, L) =$

- Hint: $L = M + N = (M_P + M_U) + N$
- Lemma 3: $\rho(M_P, L) = \rho(M_P, M) \cdot \rho(M, L)$
- Lemma 4: $\rho(M_P, M) = ?$
 - For the 8-bit architecture:
 - For the 32-bit architecture:
 - For the 128-bit architecture:

- Hint: $L = M + N = (M_P + M_U) + N$
- Lemma 3: $\rho(M_P, L) = \rho(M_P, M) \cdot \rho(M, L)$
- Lemma 4: $\rho(M_P, M) = ?$
 - For the 8-bit architecture: $\sqrt{8/8}$
 - For the 32-bit architecture:
 - For the 128-bit architecture:

- Hint: $L = M + N = (M_P + M_U) + N$
- Lemma 3: $\rho(M_P, L) = \rho(M_P, M) \cdot \rho(M, L)$
- Lemma 4: $\rho(M_P, M) = ?$
 - For the 8-bit architecture: $\sqrt{8/8}$
 - For the 32-bit architecture: $\sqrt{8/32}$
 - For the 128-bit architecture: $\sqrt{8/128}$

- Hint: $L = M + N = (M_P + M_U) + N$
- Lemma 3: $\rho(M_P, L) = \rho(M_P, M) \cdot \rho(M, L)$
- Lemma 4: $\rho(M_P, M) = ?$
 - For the 8-bit architecture: $\sqrt{8/8}$
 - For the 32-bit architecture: $\sqrt{8/32}$
 - For the 128-bit architecture: $\sqrt{8/128}$
- Lemma 2: $\frac{c}{(\sqrt{8/8} \cdot \rho(M, L))^2} = 10$

- Data complexity for the 32-bit case:
- Data complexity for the 128-bit case:

- Data complexity for the 32-bit case: 40
- Data complexity for the 128-bit case: 160
- Is noise an efficient countermeasure?

- Data complexity for the 32-bit case: 40
- Data complexity for the 128-bit case: 160
- Is noise an efficient countermeasure?
 - 32-bit case: security $\times 4$, **cost** $\times ?$

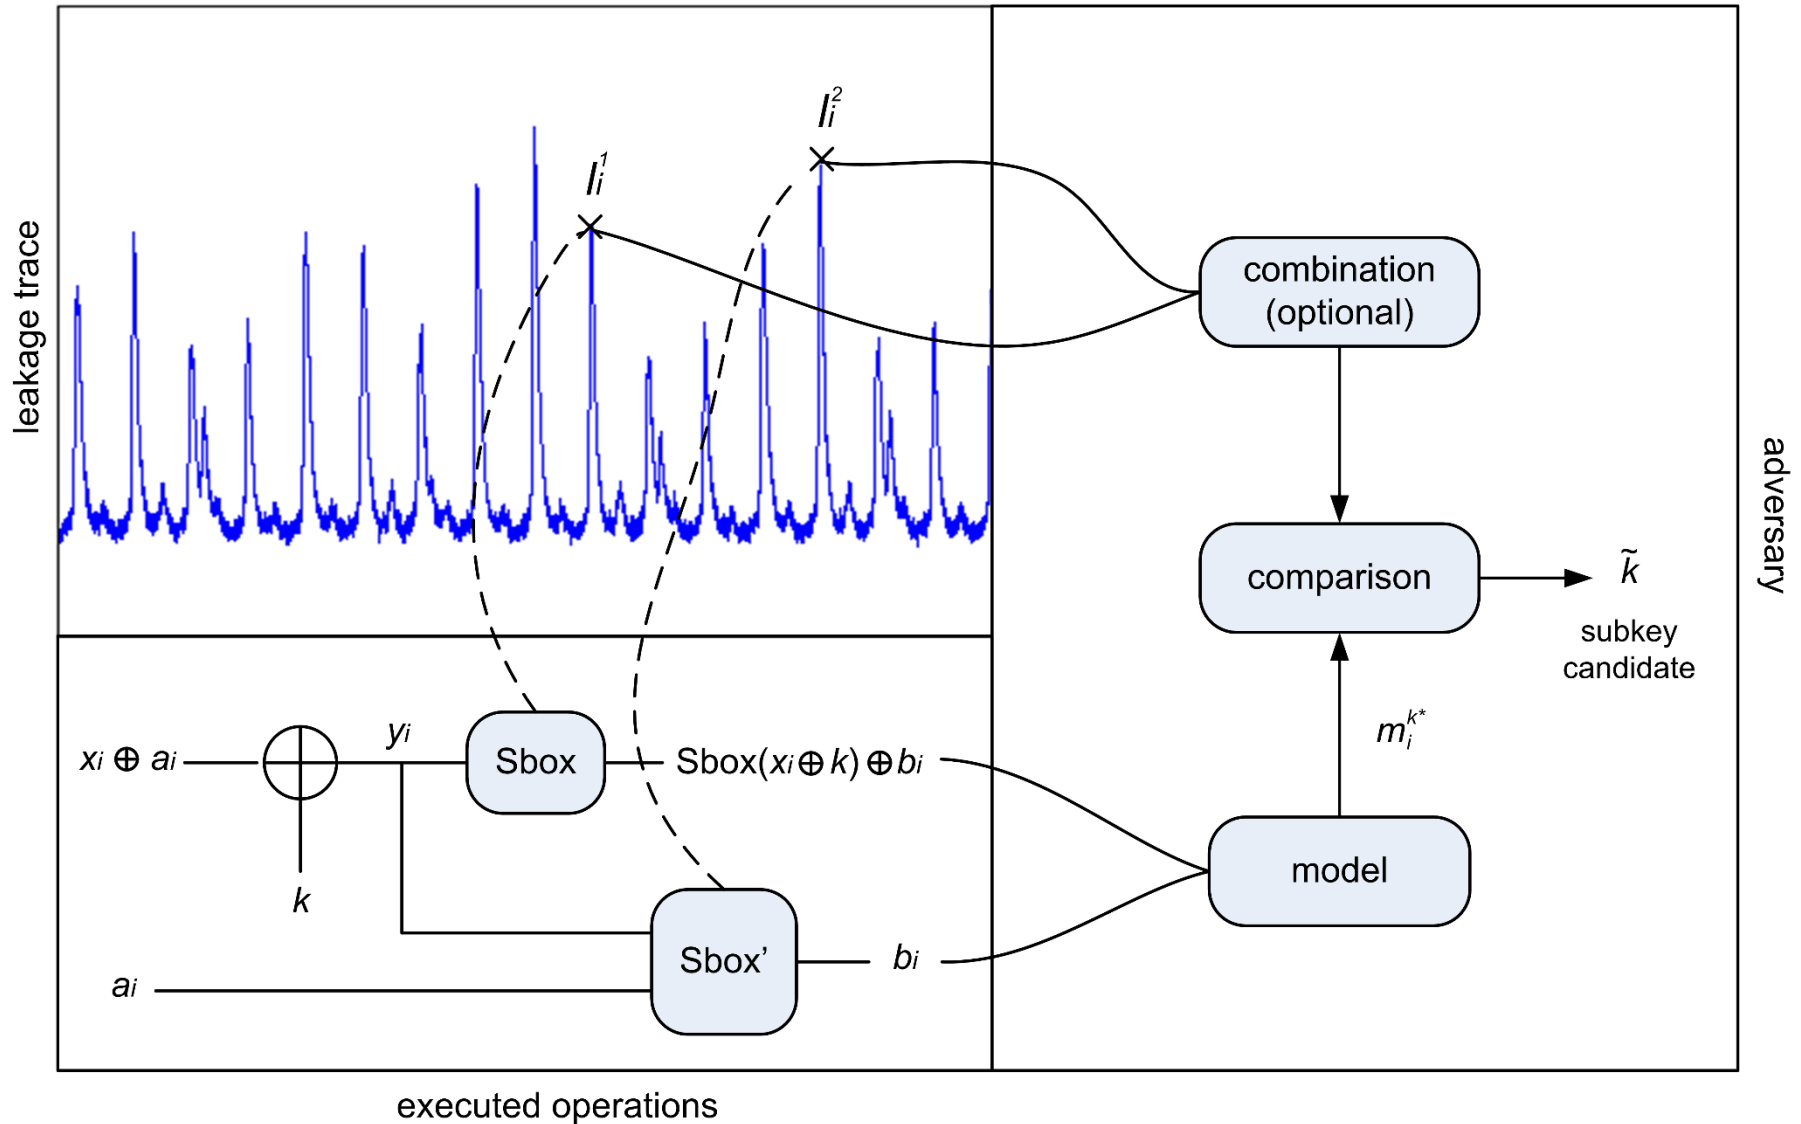
- Data complexity for the 32-bit case: 40
- Data complexity for the 128-bit case: 160
- Is noise an efficient countermeasure?
 - 32-bit case: security $\times 4$, **cost** $\times 4$
- How to trade data for time?

- Data complexity for the 32-bit case: 40
- Data complexity for the 128-bit case: 160
- Is noise an efficient countermeasure?
 - 32-bit case: security $\times 4$, **cost** $\times 4$
- How to trade data for time?
 - Target more than 8 bits at once
 - Cancels (a part of) the « algorithmic noise »
 - e.g., 32-bit architecture: $\rho(M_P, M) =$


- Data complexity for the 32-bit case: 40
- Data complexity for the 128-bit case: 160
- Is noise an efficient countermeasure?
 - 32-bit case: security $\times 4$, **cost** $\times 4$
- How to trade data for time?
 - Target more than 8 bits at once
 - Cancels (a part of) the « algorithmic noise »
 - e.g., 32-bit architecture: $\rho(M_P, M) = \sqrt{32/32}$
 - ($10 < \text{data complexity} < 40$ because of c)

Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- **Noise amplification (aka masking)**
- Reductions help (aka leakage resilience)
- Mitigating hardware defaults (is hard)
- Transparency is needed (open source)
- Summary and conclusions



- Let $z = S(x \oplus k) = S(y)$ be a leaking S-box
- Let $y = y_1 \oplus y_2 \oplus \dots \oplus y_d$ be a sharing of y



- Perform computations on “shared” variables

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \cdots \oplus f(a_d)$

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \dots \oplus f(a_d)$
- Multiplications: $c = a \times b$ in three steps

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \dots \oplus f(a_d)$
- Multiplications: $c = a \times b$ in three steps

$$\begin{bmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{bmatrix}$$

partial products

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \dots \oplus f(a_d)$
- Multiplications: $c = a \times b$ in three steps

$$\begin{bmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{bmatrix} + \begin{bmatrix} 0 & r_1 & r_2 \\ -r_1 & 0 & r_3 \\ -r_2 & r_3 & 0 \end{bmatrix}$$

partial products

refreshing

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \dots \oplus f(a_d)$
- Multiplications: $c = a \times b$ in three steps

$$\begin{bmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{bmatrix} + \begin{bmatrix} 0 & r_1 & r_2 \\ -r_1 & 0 & r_3 \\ -r_2 & r_3 & 0 \end{bmatrix} \Rightarrow \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$$

partial products

refreshing

compression

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \dots \oplus f(a_d)$
- Multiplications: $c = a \times b$ in three steps

$$\begin{bmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{bmatrix} + \begin{bmatrix} 0 & r_1 & r_2 \\ -r_1 & 0 & r_3 \\ -r_2 & r_3 & 0 \end{bmatrix} \Rightarrow \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$$

partial products

refreshing

compression

\Rightarrow Quadratic overheads & randomness

- Linear operations: $f(a) = f(a_1) \oplus f(a_2) \oplus \dots \oplus f(a_d)$
- Multiplications: $c = a \times b$ in three steps

$$\begin{bmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{bmatrix} + \begin{bmatrix} 0 & r_1 & r_2 \\ -r_1 & 0 & r_3 \\ -r_2 & r_3 & 0 \end{bmatrix} \Rightarrow \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$$

partial products

refreshing

compression

⇒ Quadratic overheads & randomness


⇒ Composable (from gadgets to circuits)

- Assume leakage variables $L_{Z_i} = \delta(Z_i) + N$ s.t.
 - $\text{MI}(Z_i; L_{Z_i}) \leq \frac{c}{d}$ (why d ? – or d^2 in proofs)
 - The leakages of the shares are independent
- For a masking scheme with d shares
- And an adversary using m measurements
- Then: $\text{SR} \leq 1 - \left(1 - \text{MI}(Z_i; L_{Z_i})\right)^d)^m$


- Assume leakage variables $L_{Z_i} = \delta(Z_i) + N$ s.t.
 - $\text{MI}(Z_i; L_{Z_i}) \leq \frac{c}{d}$ (multiplications)
 - The leakages of the shares are independent
- For a masking scheme with d shares
- And an adversary using m measurements
- Then: $\text{SR} \leq 1 - \left(1 - \text{MI}(Z_i; L_{Z_i})\right)^d)^m$

- Assume leakage variables $L_{Z_i} = \delta(Z_i) + N$ s.t.
 - $\text{MI}(Z_i; L_{Z_i}) \leq \frac{c}{d}$ (multiplications)
 - The leakages of the shares are independent
- For a masking scheme with d shares
- And an adversary using m measurements
- Then: $\text{SR} \leq 1 - (1 - \text{MI}(Z_i; L_{Z_i})^d)^m$
- For $m = 1$, $\text{SR} \leq \text{MI}(Z_i; L_{Z_i})^d \propto (\sigma_N^2)^d$
- (Intuitively \approx “noisy” piling up lemma)


- 1-bit, 2-shares example




(a) $Z = 0$, serial.



(b) $Z = 1$, serial.




(c) $Z = 0$, parallel.




(d) $Z = 1$, parallel.

- 1-bit, 2-shares example


key-independent means




(a) $Z = 0$, serial.



(b) $Z = 1$, serial.




(c) $Z = 0$, parallel.




(d) $Z = 1$, parallel.

- Slope of the IT curves = d (if independent leaks)




- Is masking an efficient countermeasure?
 - Security (data) is exponential in d
 - Cost is [...]

- Is masking an efficient countermeasure?
 - Security (data) is exponential in d
 - **Cost is [...]** quadratic in d



- Is masking an efficient countermeasure?
 - Security (data) is exponential in d
 - Cost is [...] quadratic in d



- If the leakages are noisy and independent (!)

- Is masking an efficient countermeasure?
 - Security (data) is exponential in d
 - Cost is [...] quadratic in d



- If the leakages are noisy and independent (!)
- How does the time complexity scale in d ?

- Is masking an efficient countermeasure?
 - Security (data) is exponential in d
 - Cost is [...] quadratic in d

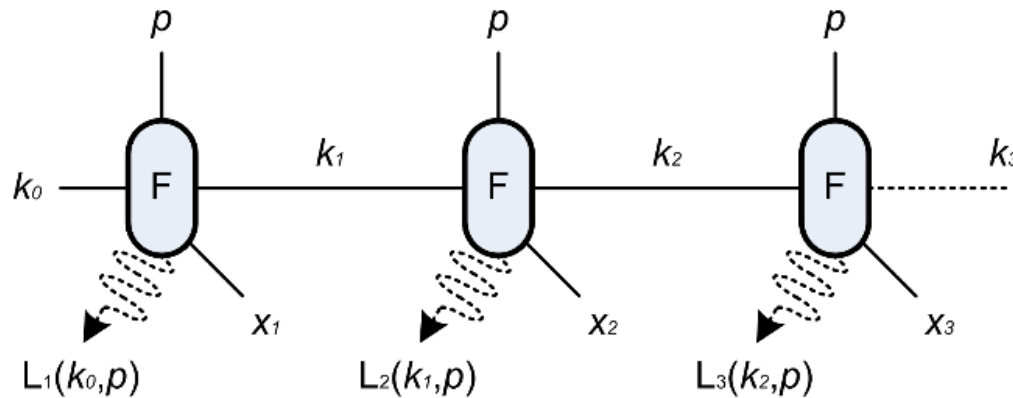


- If the leakages are noisy and independent (!)
- How does the time complexity scale in d ?
 - Depends on the implem. (e.g., serial or //)

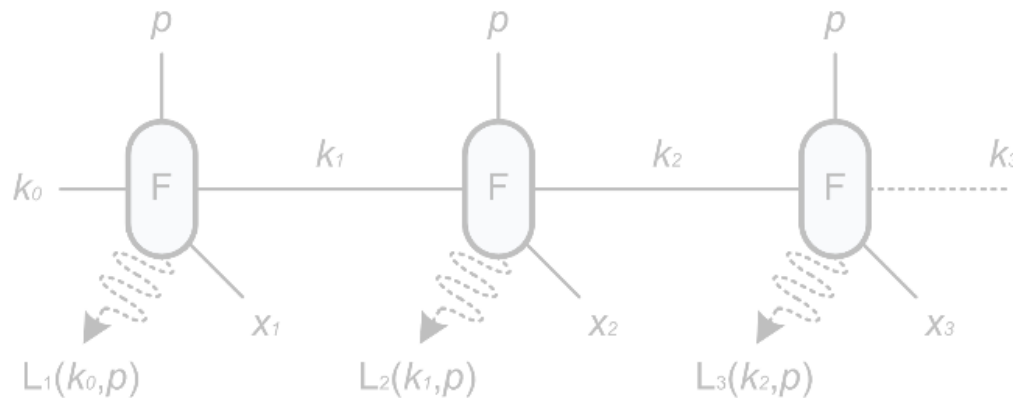
Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- **Reductions help (aka leakage resilience)**
- Mitigating hardware defaults (is hard)
- Transparency is needed (open source)
- Summary and conclusions

- Most natural construction: forward-secure PRG

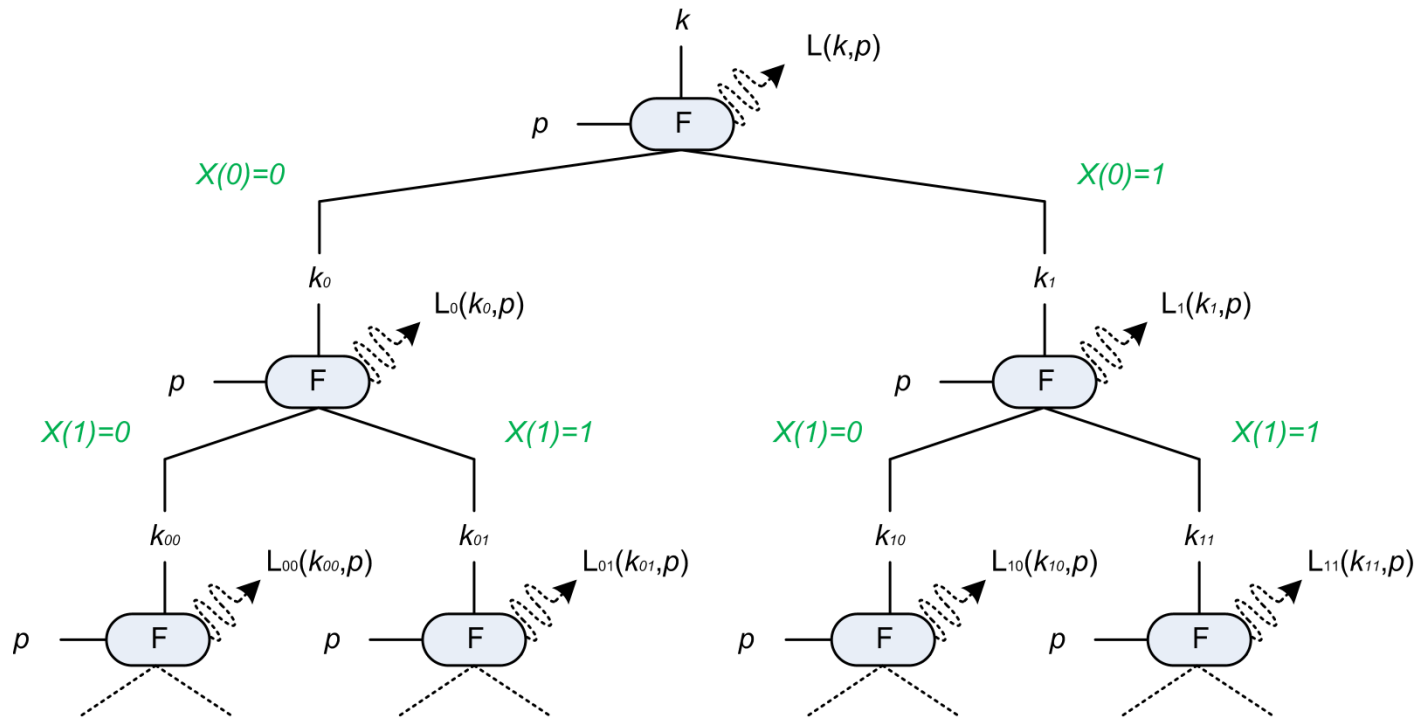


- Most natural construction: forward-secure PRG

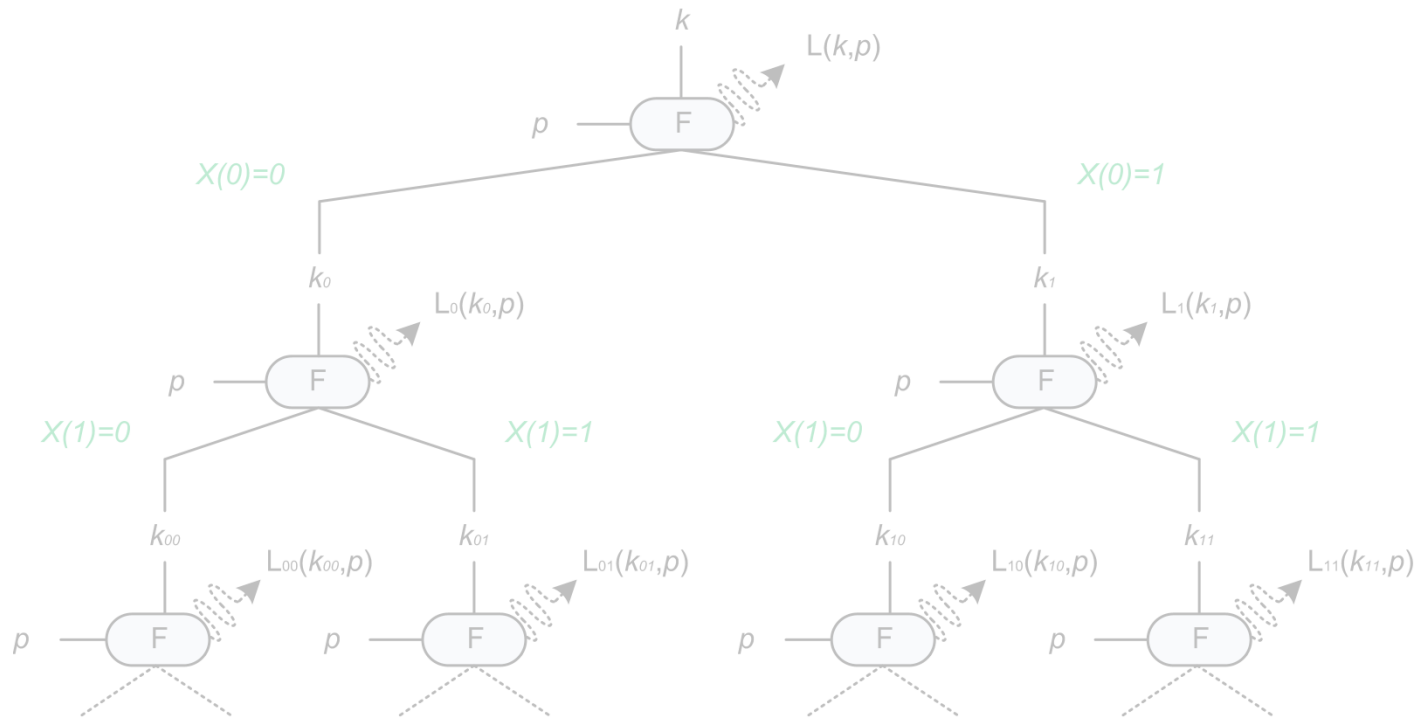


- Re-keying impact: bounds the number of (noisy) measurements per key (*prevents averaging*)

- Most natural construction: GGM tree



- Most natural construction: GGM tree





- Re-keying impact: bounds the number of noise-free observations per key (*allows averaging*)

- Key recovery security (standard DPA):

PRG

PRF




- « Bounded security » for the PRG only
 - (Analytical/algebraic attacks not considered)


- A call to a stateless primitive is always needed
 - For initialization / randomization
 - For authentication and encryption

- A call to a stateless primitive is always needed
 - For initialization / randomization
 - For authentication and encryption
- But we can try to encrypt large messages with a single call to this (more expensive) primitive

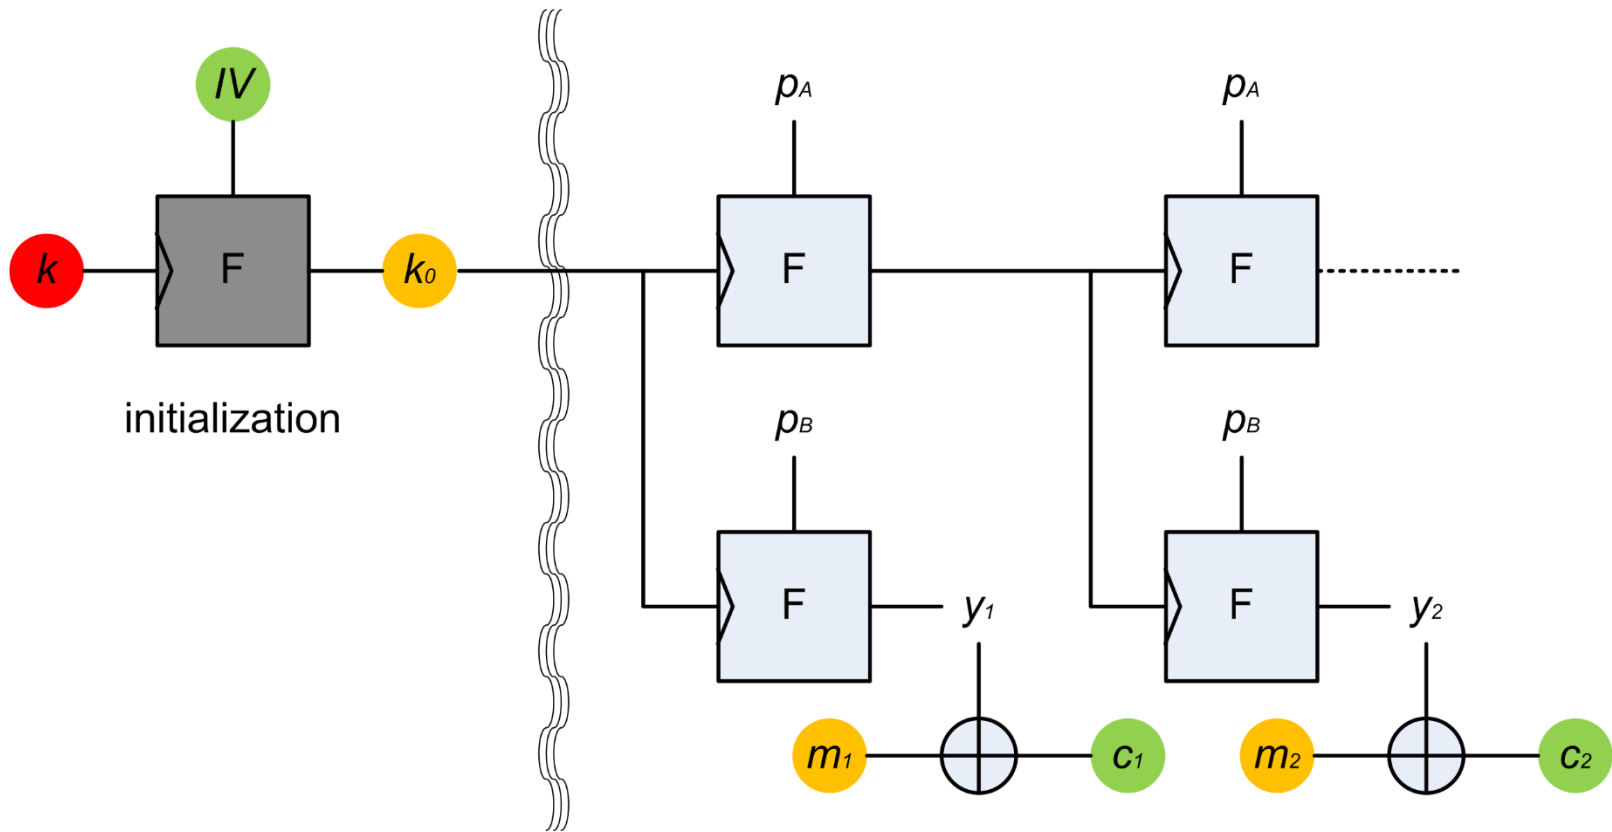
- A call to a stateless primitive is always needed
 - For initialization / randomization
 - For authentication and encryption
- But we can try to encrypt large messages with a single call to this (more expensive) primitive
- **And to use leakage-resilient PRGs otherwise**



- Green: public value, orange: ephemeral secret, red: long-term secret (protected with leak-free F^*)



- Green: public value, orange: ephemeral secret, red: long-term secret (protected with leak-free F^*)
- τ unforgeable even with leakage (during enc.)
- Security of 1-block \approx security of l -blocks
- & high-security levels expected
 - Because it is an unpredictability game!



- Similar reduction but lower security levels
 - Because it is an indistinguishability game!

- In theory, the proof challenge remains open

François-Xavier Standaert, Olivier Pereira, Yu Yu: *Leakage-Resilient Symmetric Cryptography under Empirically Verifiable Assumptions*. CRYPTO (1) 2013: 335-352. Jake Longo, Daniel P. Martin, Elisabeth Oswald, Daniel Page, Martijn Stam, Michael Tunstall: *Simulatable Leakage: Analysis, Pitfalls, and New Constructions*. ASIACRYPT (1) 2014: 223-242

- Yet, the pragmatic model seems sound

- In theory, the proof challenge remains open

François-Xavier Standaert, Olivier Pereira, Yu Yu: *Leakage-Resilient Symmetric Cryptography under Empirically Verifiable Assumptions*. CRYPTO (1) 2013: 335-352. Jake Longo, Daniel P. Martin, Elisabeth Oswald, Daniel Page, Martijn Stam, Michael Tunstall: *Simulatable Leakage: Analysis, Pitfalls, and New Constructions*. ASIACRYPT (1) 2014: 223-242

- Yet, the pragmatic model seems sound
- In practice, how to design F^* is open too

- In theory, the proof challenge remains open

François-Xavier Standaert, Olivier Pereira, Yu Yu: *Leakage-Resilient Symmetric Cryptography under Empirically Verifiable Assumptions*. CRYPTO (1) 2013: 335-352. Jake Longo, Daniel P. Martin, Elisabeth Oswald, Daniel Page, Martijn Stam, Michael Tunstall: *Simulatable Leakage: Analysis, Pitfalls, and New Constructions*. ASIACRYPT (1) 2014: 223-242

- Yet, the pragmatic model seems sound
- In practice, how to design F^* is open too, e.g.,
 - Masking (\Rightarrow bitslice ciphers)

Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, Kerem Varici: *LS-Designs: Bitslice Encryption for Efficient Masked Software Implementations*. FSE 2014: 18-37

- In theory, the proof challenge remains open

François-Xavier Standaert, Olivier Pereira, Yu Yu: *Leakage-Resilient Symmetric Cryptography under Empirically Verifiable Assumptions*. CRYPTO (1) 2013: 335-352. Jake Longo, Daniel P. Martin, Elisabeth Oswald, Daniel Page, Martijn Stam, Michael Tunstall: *Simulatable Leakage: Analysis, Pitfalls, and New Constructions*. ASIACRYPT (1) 2014: 223-242

- Yet, the pragmatic model seems sound
- In practice, how to design F^* is open too, e.g.,
 - Masking (\Rightarrow bitslice ciphers)

Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, Kerem Varici: *LS-Designs: Bitslice Encryption for Efficient Masked Software Implementations*. FSE 2014: 18-37

- **PRFs with non-standard assumptions**

Marcel Medwed, François-Xavier Standaert, Antoine Joux: *Towards Super-Exponential Side-Channel Security with Efficient Leakage-Resilient PRFs*. CHES 2012: 193-212. Marcel Medwed, François-Xavier Standaert, Ventzi Nikov, Martin Feldhofer, *Unknown-Input Attacks in the Parallel Setting: Improving the Security and Performances of the CHES 2012 Leakage-Resilient PRF*, ASIACRYPT 2016 (to appear)

- In theory, the proof challenge remains open

François-Xavier Standaert, Olivier Pereira, Yu Yu: *Leakage-Resilient Symmetric Cryptography under Empirically Verifiable Assumptions*. CRYPTO (1) 2013: 335-352. Jake Longo, Daniel P. Martin, Elisabeth Oswald, Daniel Page, Martijn Stam, Michael Tunstall: *Simulatable Leakage: Analysis, Pitfalls, and New Constructions*. ASIACRYPT (1) 2014: 223-242

- Yet, the pragmatic model seems sound
- In practice, how to design F^* is open too, e.g.,
 - Masking (\Rightarrow bitslice ciphers)

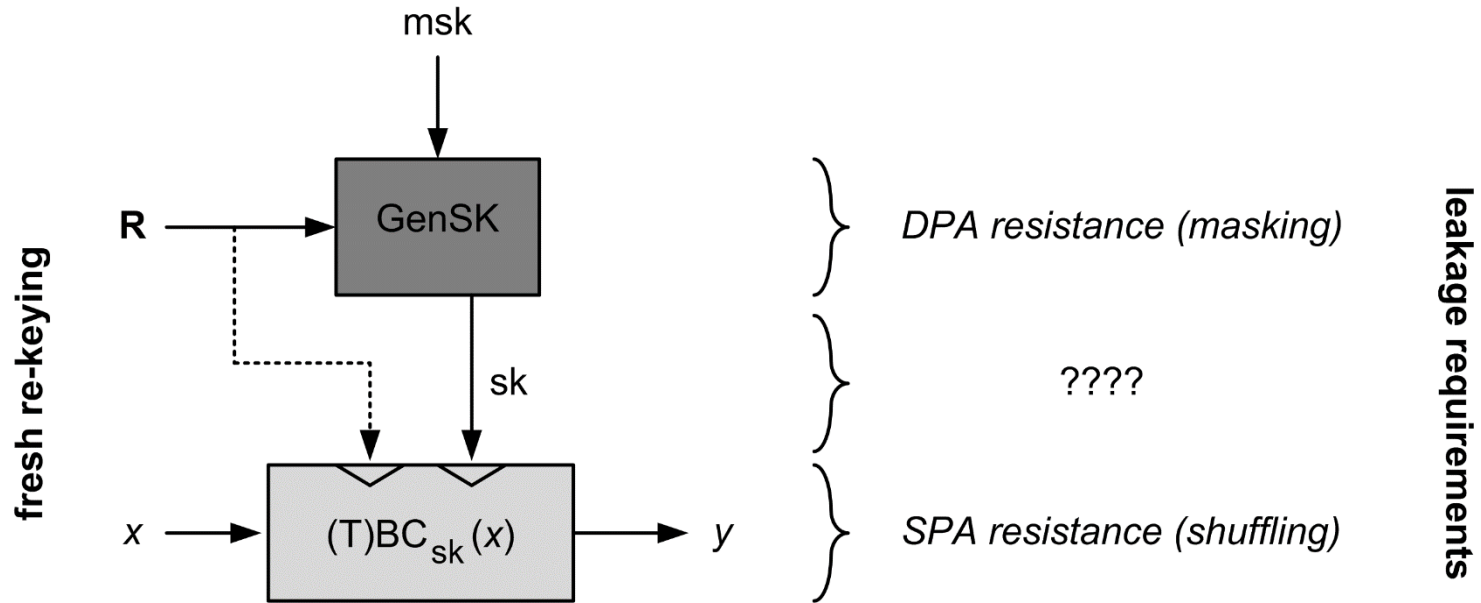
Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, Kerem Varici: *LS-Designs: Bitslice Encryption for Efficient Masked Software Implementations*. FSE 2014: 18-37


- PRFs with non-standard assumptions

Marcel Medwed, François-Xavier Standaert, Antoine Joux: *Towards Super-Exponential Side-Channel Security with Efficient Leakage-Resilient PRFs*. CHES 2012: 193-212. Marcel Medwed, François-Xavier Standaert, Ventzi Nikov, Martin Feldhofer, *Unknown-Input Attacks in the Parallel Setting: Improving the Security and Performances of the CHES 2012 Leakage-Resilient PRF*, ASIACRYPT 2016 (to appear)

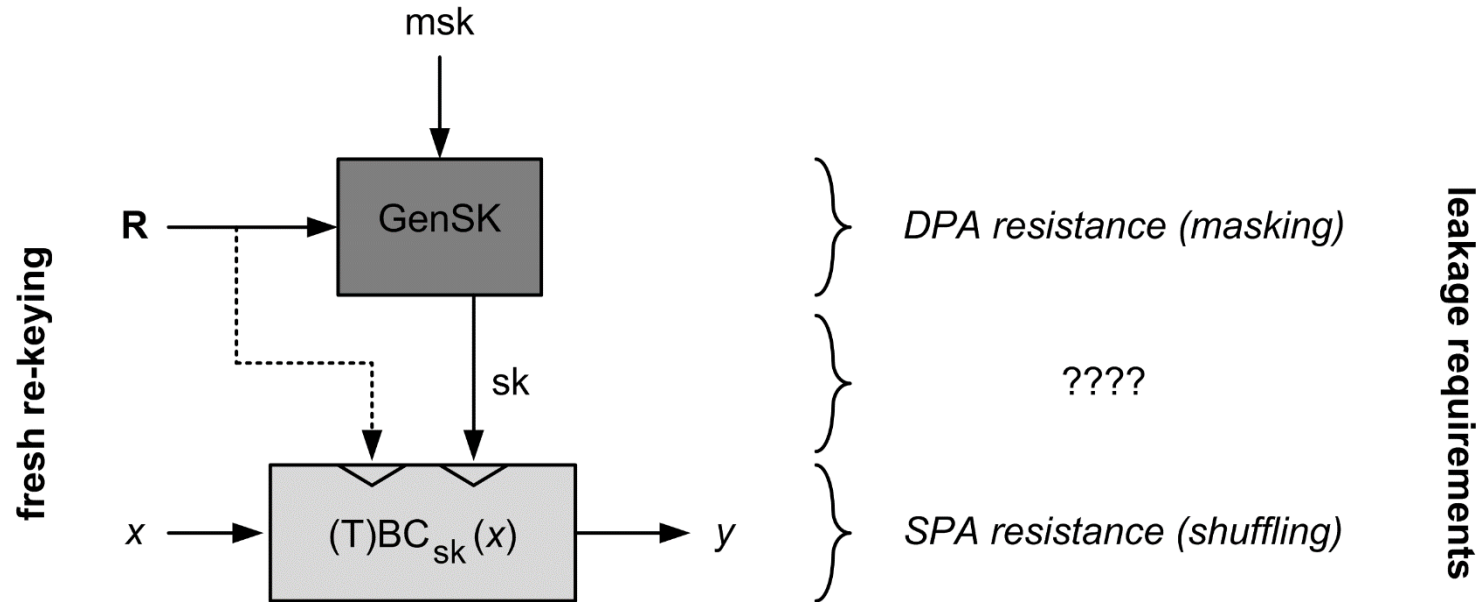
- **Key homomorphism & fresh re-keying**

Christoph Dobraunig, François Koeune, Stefan Mangard, Florian Mendel, François-Xavier Standaert: *Towards Fresh and Hybrid Re-Keying Schemes with Beyond Birthday Security*. CARDIS 2015: 225-241





- Cryptographically strong re-keying function
 - $sk = \langle \mathbf{R}, msk \rangle = \sum (\langle \mathbf{R}, msk_i \rangle)$



- Cryptographically strong re-keying function
 - $sk = \langle \mathbf{R}, msk \rangle = \sum (\langle \mathbf{R}, msk_i \rangle)$
- Security based on hard lattice problems
- Simple & efficient: all computations in \mathbb{Z}_2^m


- Authenticated encryption is also possible
 - But combination with IV misuse is tricky
 - Because controlling the IV transforms ephemeral secrets into long-term ones
 - Same reason makes LR-decryption tricky

- Authenticated encryption is also possible
 - But combination with IV misuse is tricky
 - Because controlling the IV transforms ephemeral secrets into long-term ones
 - Same reason makes LR-decryption tricky
 - Full misuse resistance does not seem possible
 - (In the symmetric crypto setting)
- ⇒ Current answer: ciphertext integrity with misuse (possible because unpredictability-based)

Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- **Mitigating hardware defaults (is hard)**
- Transparency is needed (open source)
- Summary and conclusions


- Can break leakage independence requirements



- e.g., recombine the shares of a masking scheme
⇒ Makes secure masked implementation hard to obtain

Stefan Mangard, Thomas Popp, Berndt M. Gammel: *Side-Channel Leakage of Masked CMOS Gates*. CT-RSA 2005: 351-365. Josep Balasch, Benedikt Gierlichs, Vincent Grosso, Oscar Reparaz, François-Xavier Standaert: *On the Cost of Lazy Engineering for Masked Software Implementations*. CARDIS 2014: 64-81. Svetla Nikova, Vincent Rijmen, Martin Schl affer: *Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches*. J. Cryptology 24(2): 292-321 (2011)

- Can break leakage independence requirements



- e.g., recombine the shares of a masking scheme
⇒ Makes secure masked implementation hard to obtain
- Default-tolerant protections would be (very) handy

Stefan Mangard, Thomas Popp, Berndt M. Gammel: *Side-Channel Leakage of Masked CMOS Gates*. CT-RSA 2005: 351-365. Josep Balasch, Benedikt Gierlichs, Vincent Grosso, Oscar Reparaz, François-Xavier Standaert: *On the Cost of Lazy Engineering for Masked Software Implementations*. CARDIS 2014: 64-81. Svetla Nikova, Vincent Rijmen, Martin Schl affer: *Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches*. J. Cryptology 24(2): 292-321 (2011)


Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- Mitigating hardware defaults (is hard)
- **Transparency is needed (open source)**
- Summary and conclusions


standard practice




attack-based evaluations



standard practice



attack-based evaluations



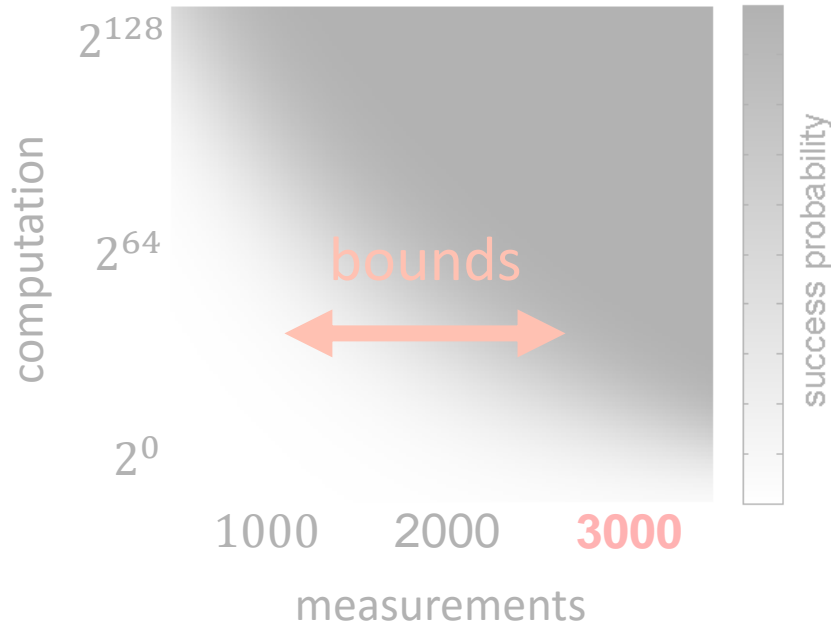
standard practice




transparency
helps evaluations



attack-based evaluations



proof-based evaluations



- As masking order increases, the # of d -tuples of informative samples increases (say by d)
⇒ the gap between “simple” attacks targeting one d -tuple and d ones increase by a factor d

- As masking order increases, the # of d -tuples of informative samples increases (say by d)
⇒ the gap between “simple” attacks targeting one d -tuple and d ones increase by a factor d
- If shares are re-used (allowing averaging before combination) this factor becomes d^d

- As masking order increases, the # of d -tuples of informative samples increases (say by d)

⇒ the gap between “simple” attacks targeting one d -tuple and d ones increase by a factor d

- If shares are re-used (allowing averaging before combination) this factor becomes d^d

⇒ It means security depends on efficiency (in cycles), e.g., parallelism reduces # of leaking tuples

Outline

- Preliminary questions / definitions
- Side-channel basics (attack steps)
- Noise (aka hardware) is not enough
- Noise amplification (aka masking)
- Reductions help (aka leakage resilience)
- Transparency is needed (open source)
- Mitigating hardware defaults (is hard)
- **Summary and conclusions**

- Effective countermeasures against side-channel attacks always combine sound hardware assumptions & mathematical amplification

- Effective countermeasures against side-channel attacks always combine sound hardware assumptions & mathematical amplification
- High physical security is not mission impossible but has a cost! (e.g., time $\times > 10$, area $\times > 2$)
 - Yet, good designs can mitigate this cost

- Effective countermeasures against side-channel attacks always combine sound hardware assumptions & mathematical amplification
- High physical security is not mission impossible but has a cost! (e.g., time $\times > 10$, area $\times > 2$)
 - Yet, good designs can mitigate this cost
- Metric I: # of operations per sensitive variable
 - Physical security \propto efficiency
- Metric II: non-linearity (because hard to mask)

- Systematic ways to deal with hardware defaults also has a price (e.g., doubling the # of shares)
 - But is probably worth it (to reduce risk)

- Systematic ways to deal with hardware defaults also has a price (e.g., doubling the # of shares)
 - But is probably worth it (to reduce risk)
- Security should not depend on adversaries
 - Beware of too specific evaluations (T-tests)
 - Especially for protected implementations

- Systematic ways to deal with hardware defaults also has a price (e.g., doubling the # of shares)
 - But is probably worth it (to reduce risk)
- Security should not depend on adversaries
 - Beware of too specific evaluations (T-tests)
 - Especially for protected implementations
- Long-term: open source codes/chips that can be used by any (non SCA expert) engineer

THANKS

<http://perso.uclouvain.be/fstandae/>