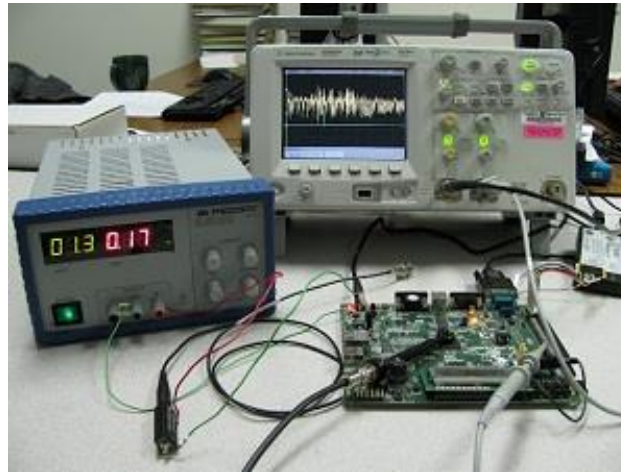


Towards Fair and Efficient Evaluations of Leaking Cryptographic Devices (Overview of the ERC Project CRASH, part I)



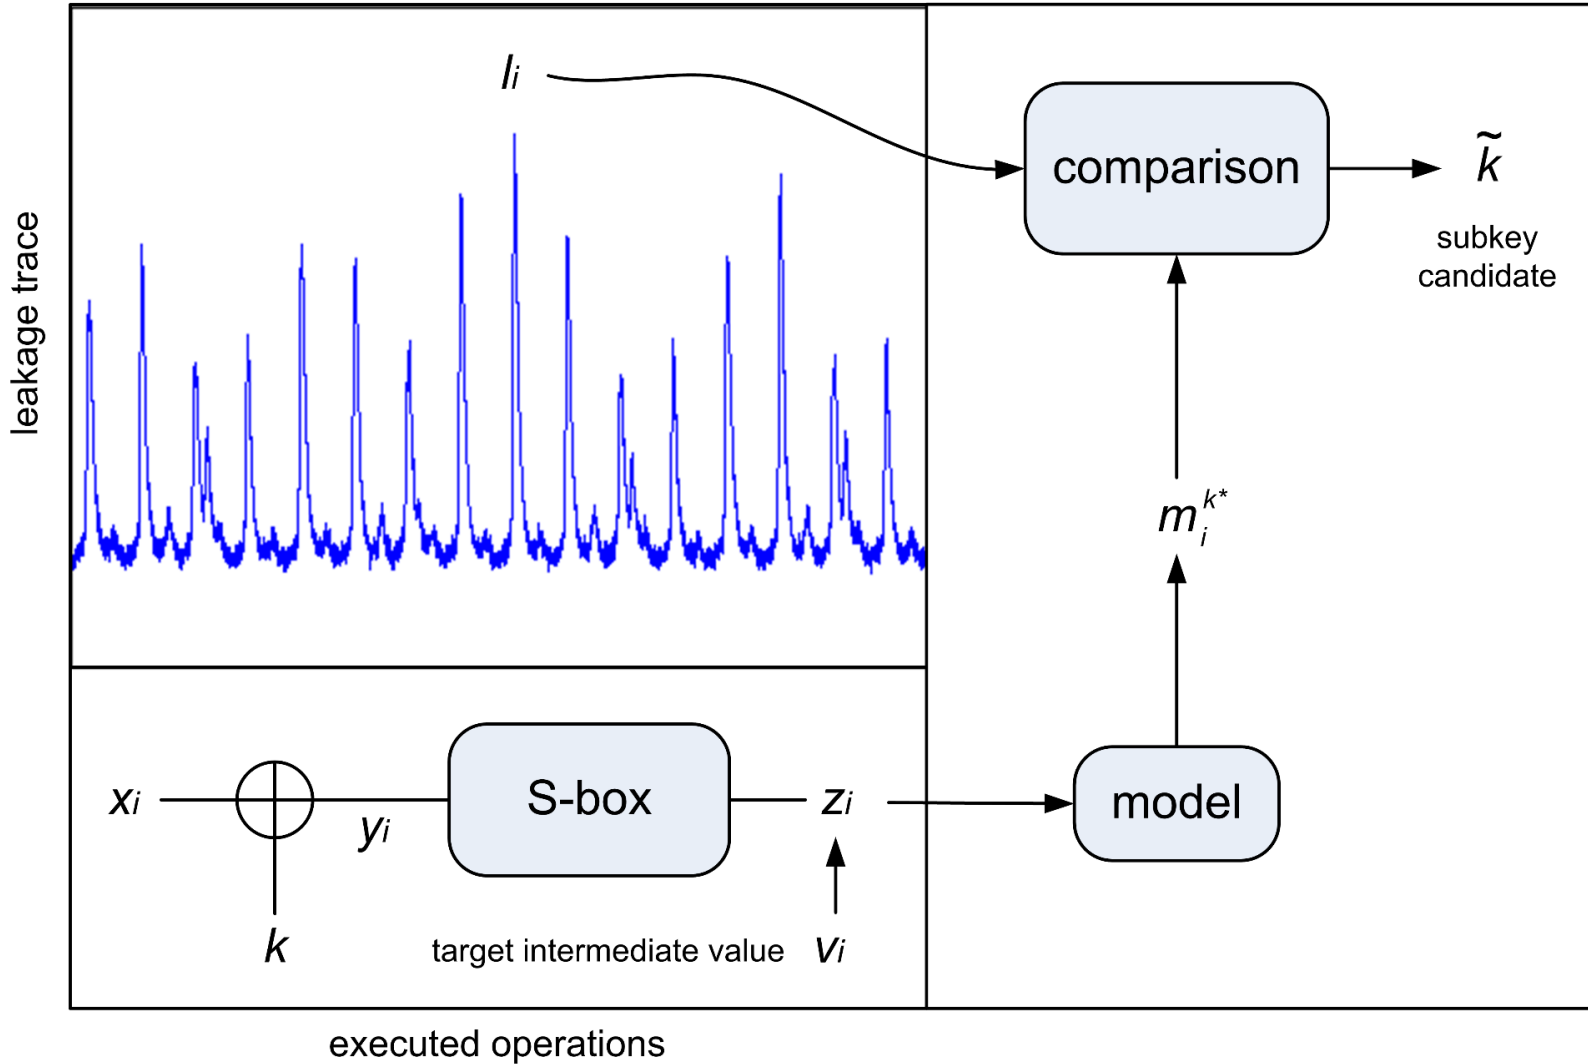
François-Xavier Standaert
UCL Crypto Group, Belgium
SPACE, December 2016

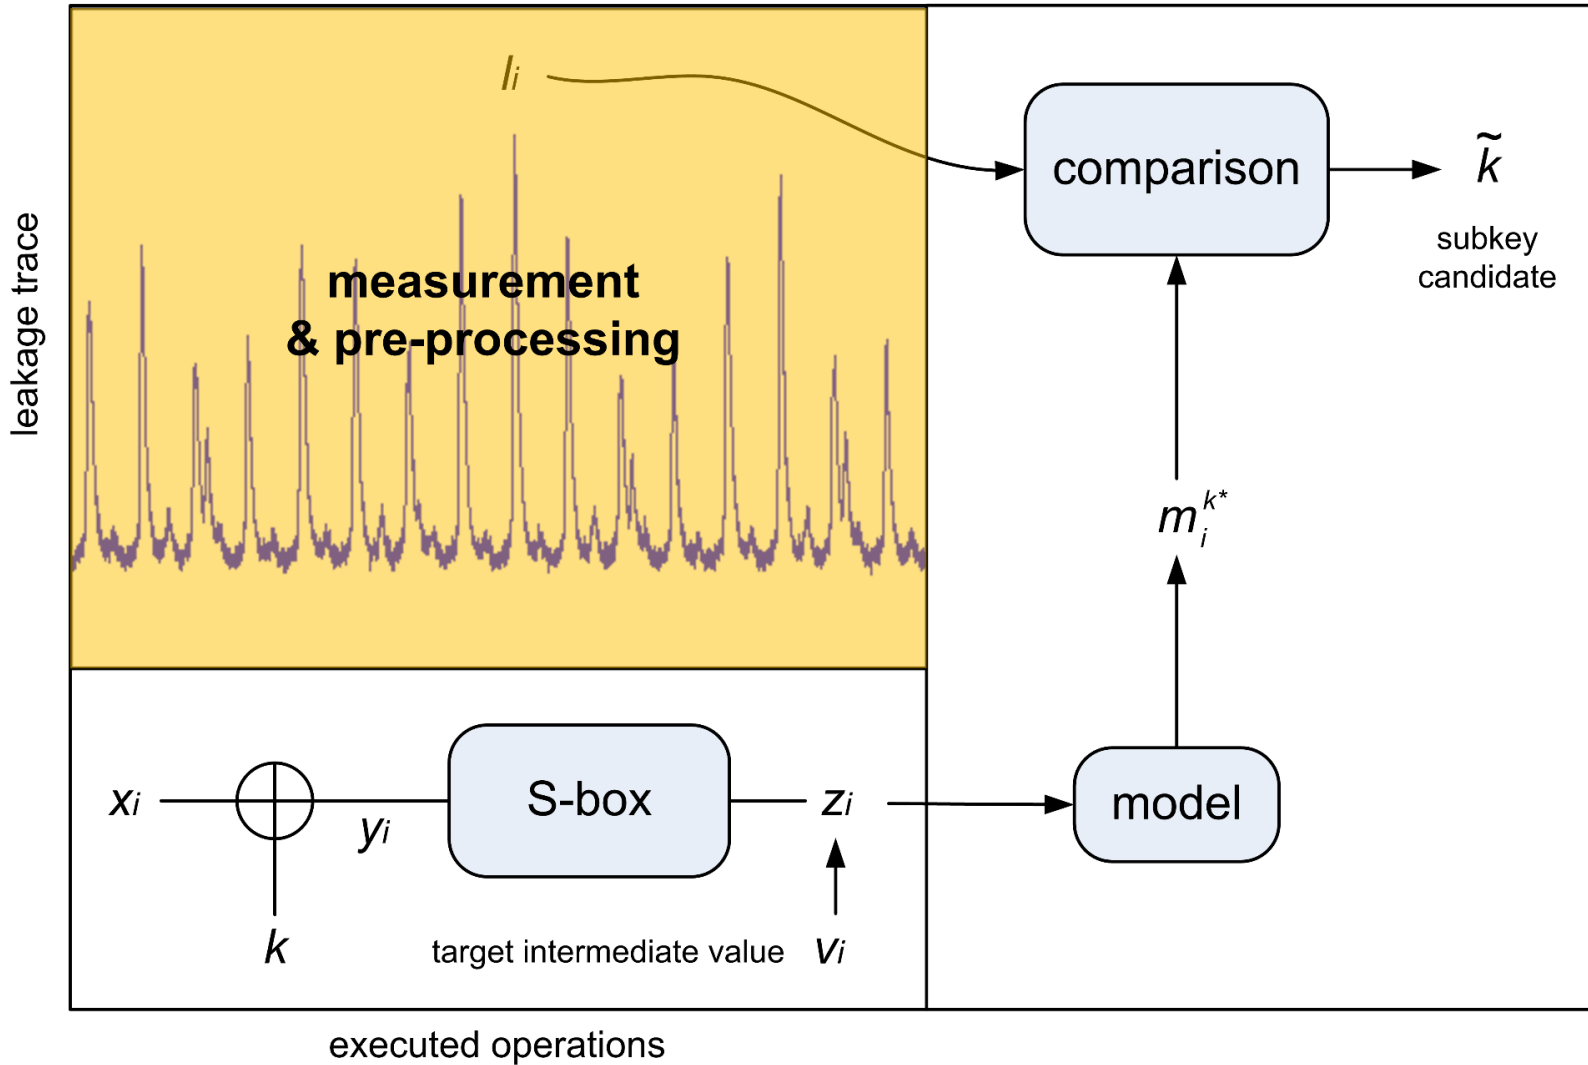
Outline

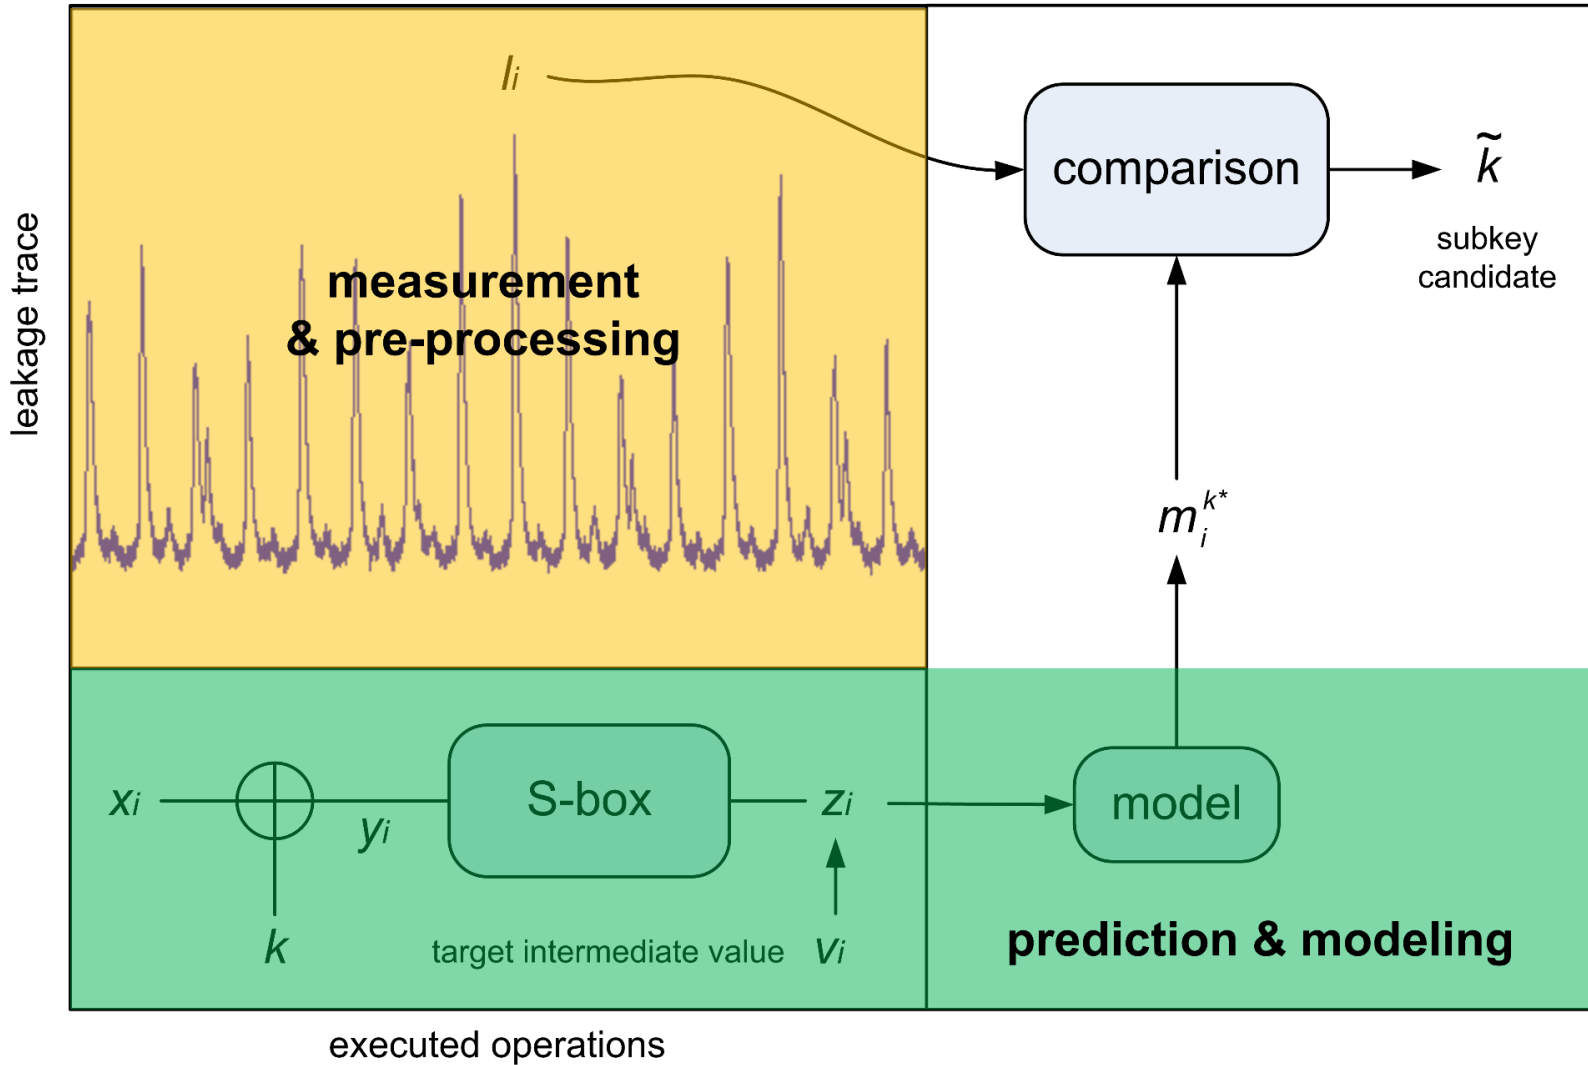
- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs

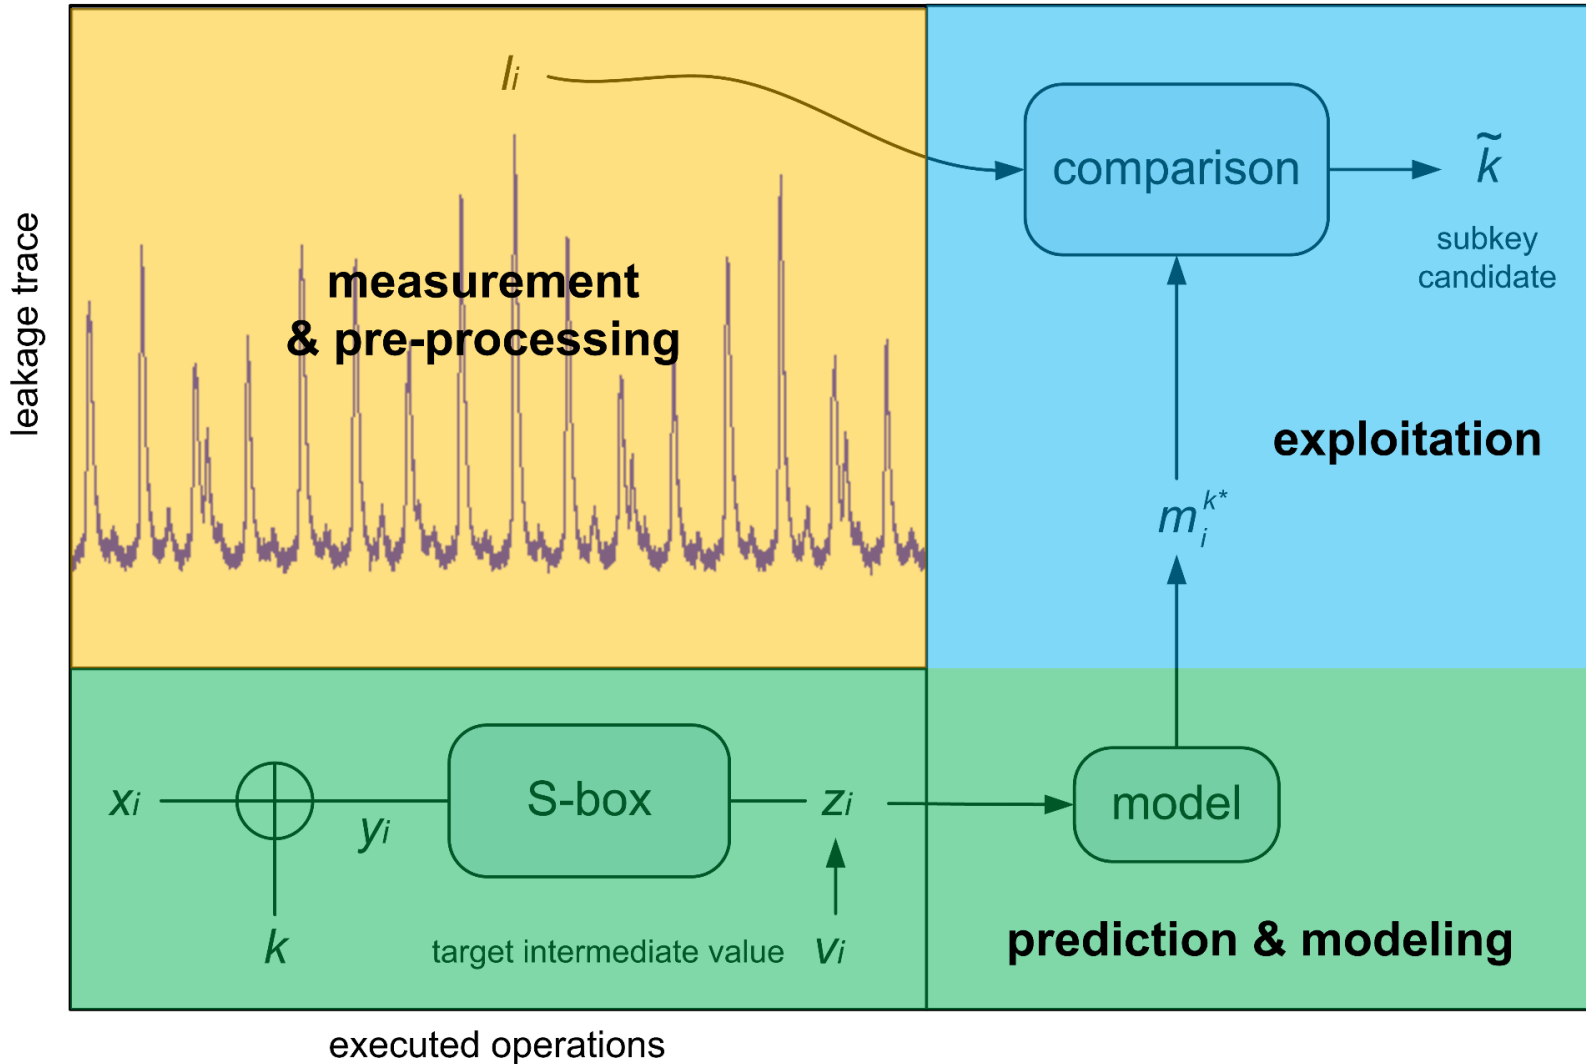
Outline

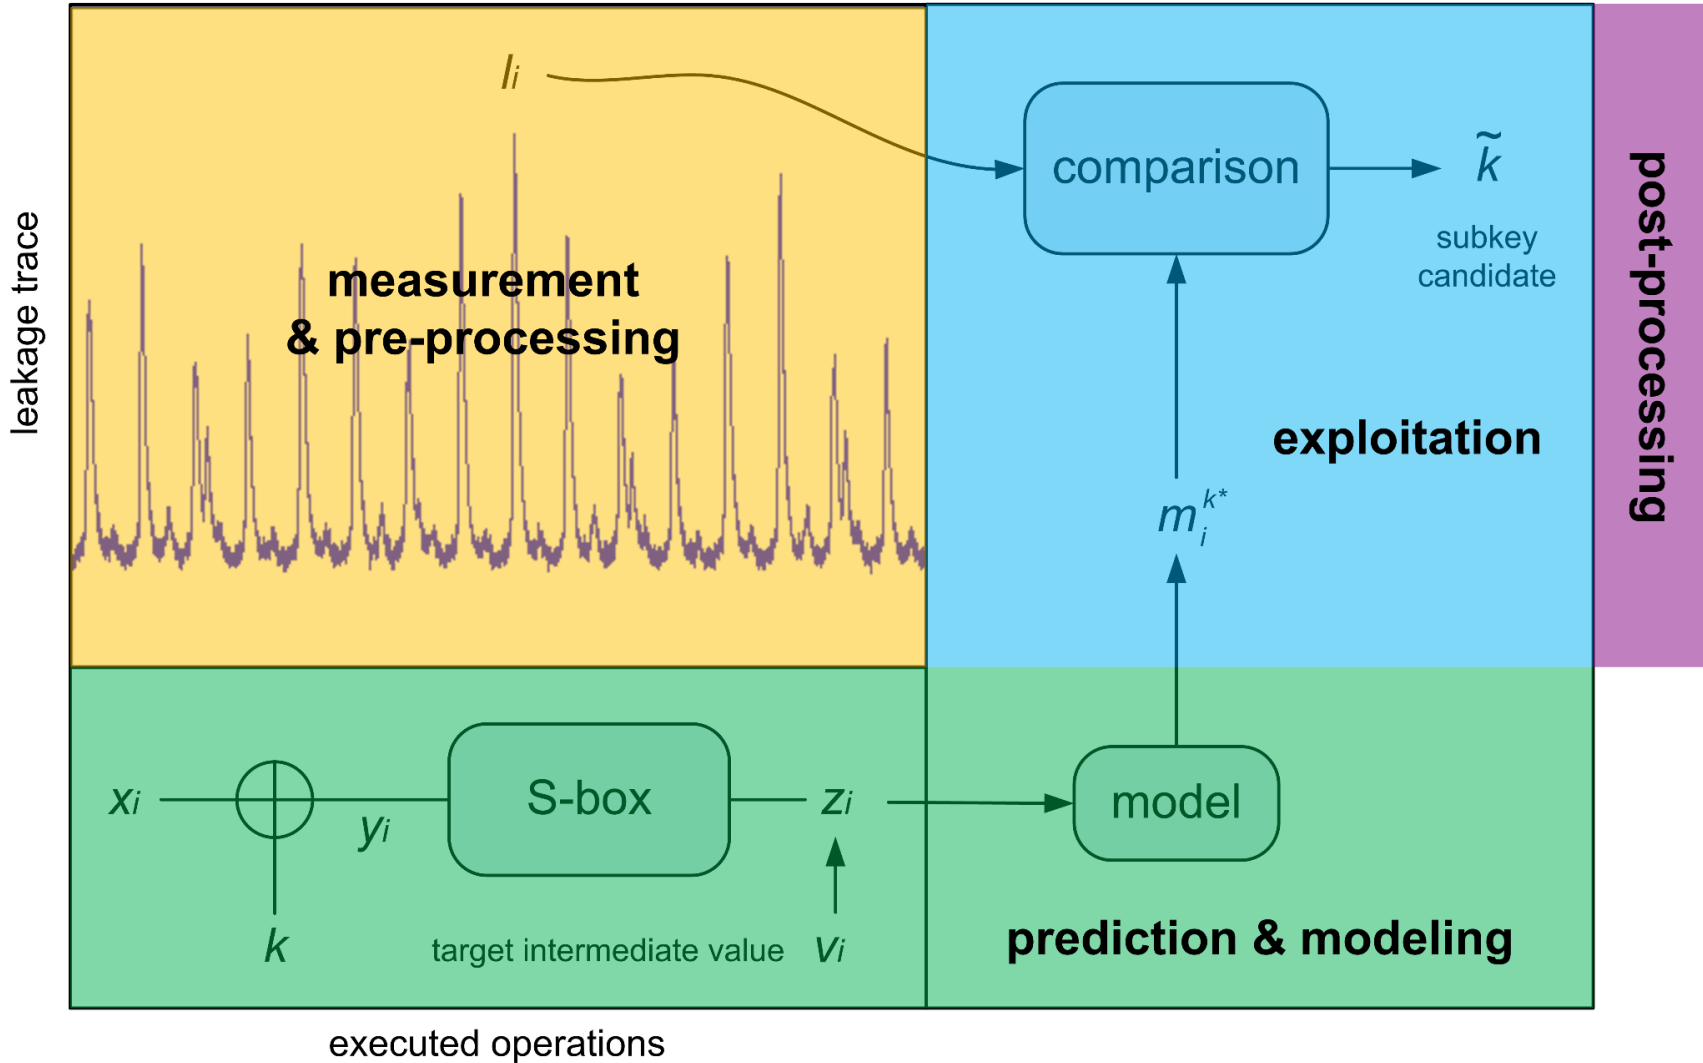
- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs

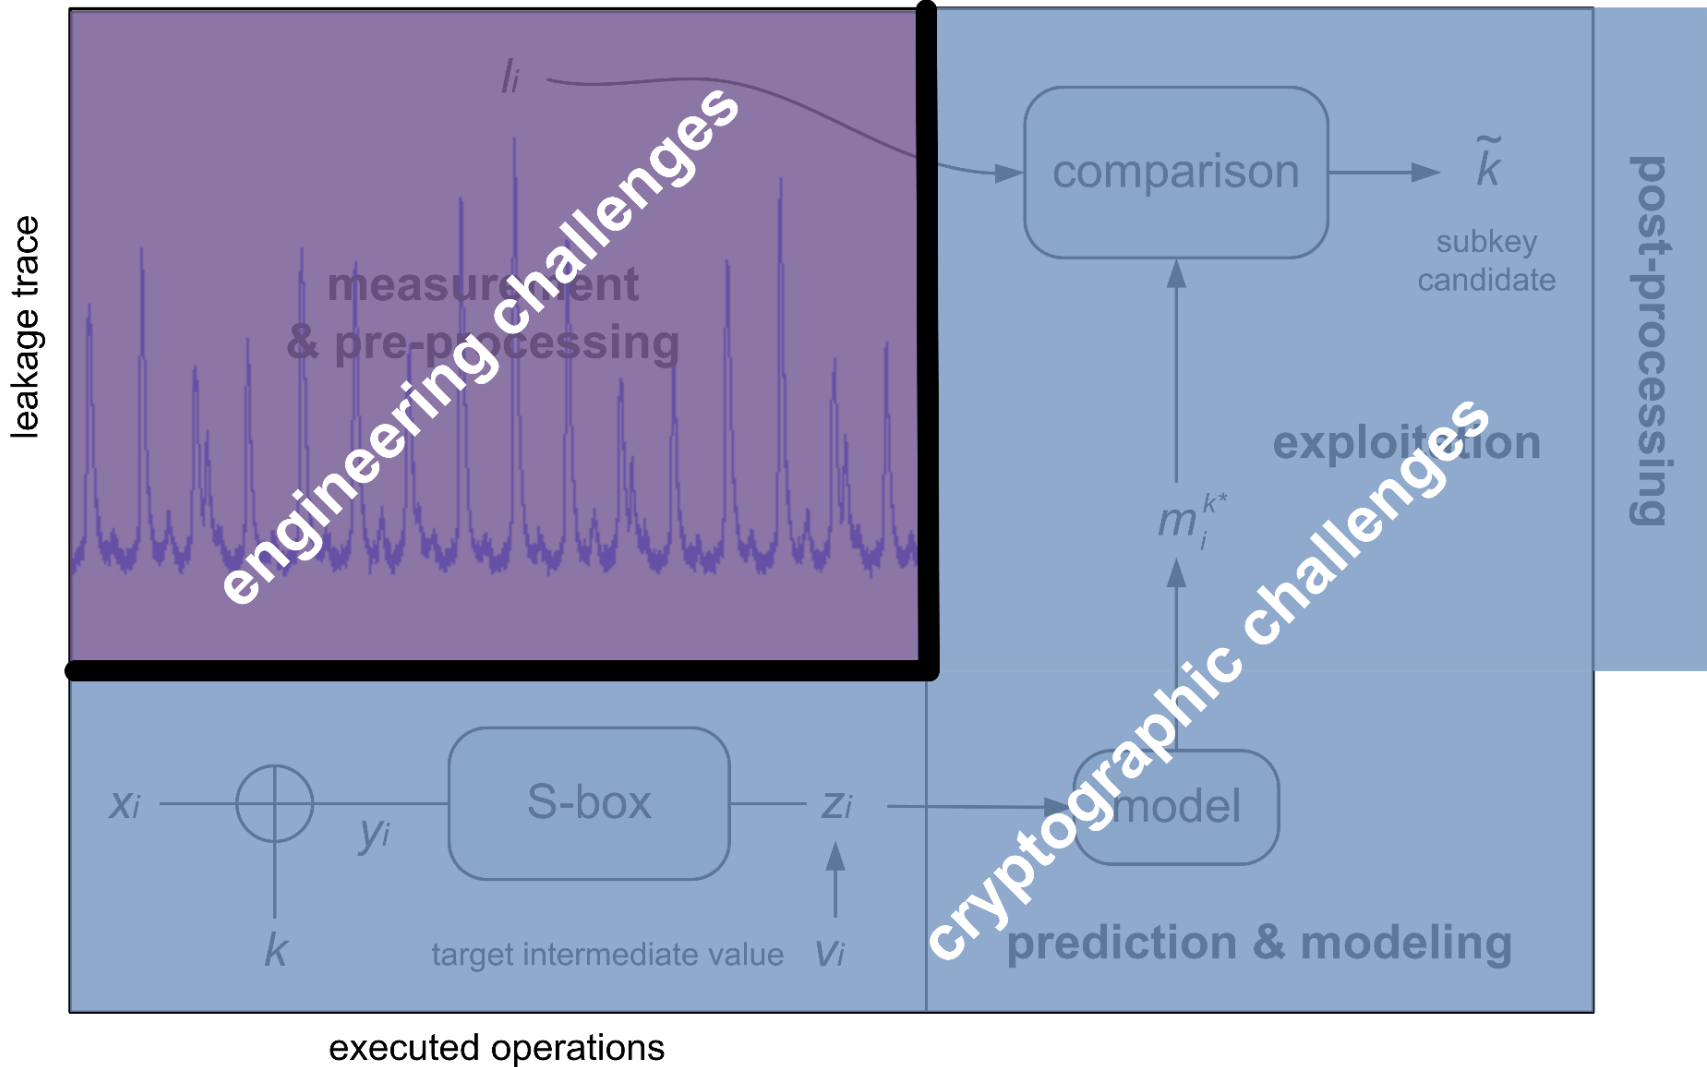






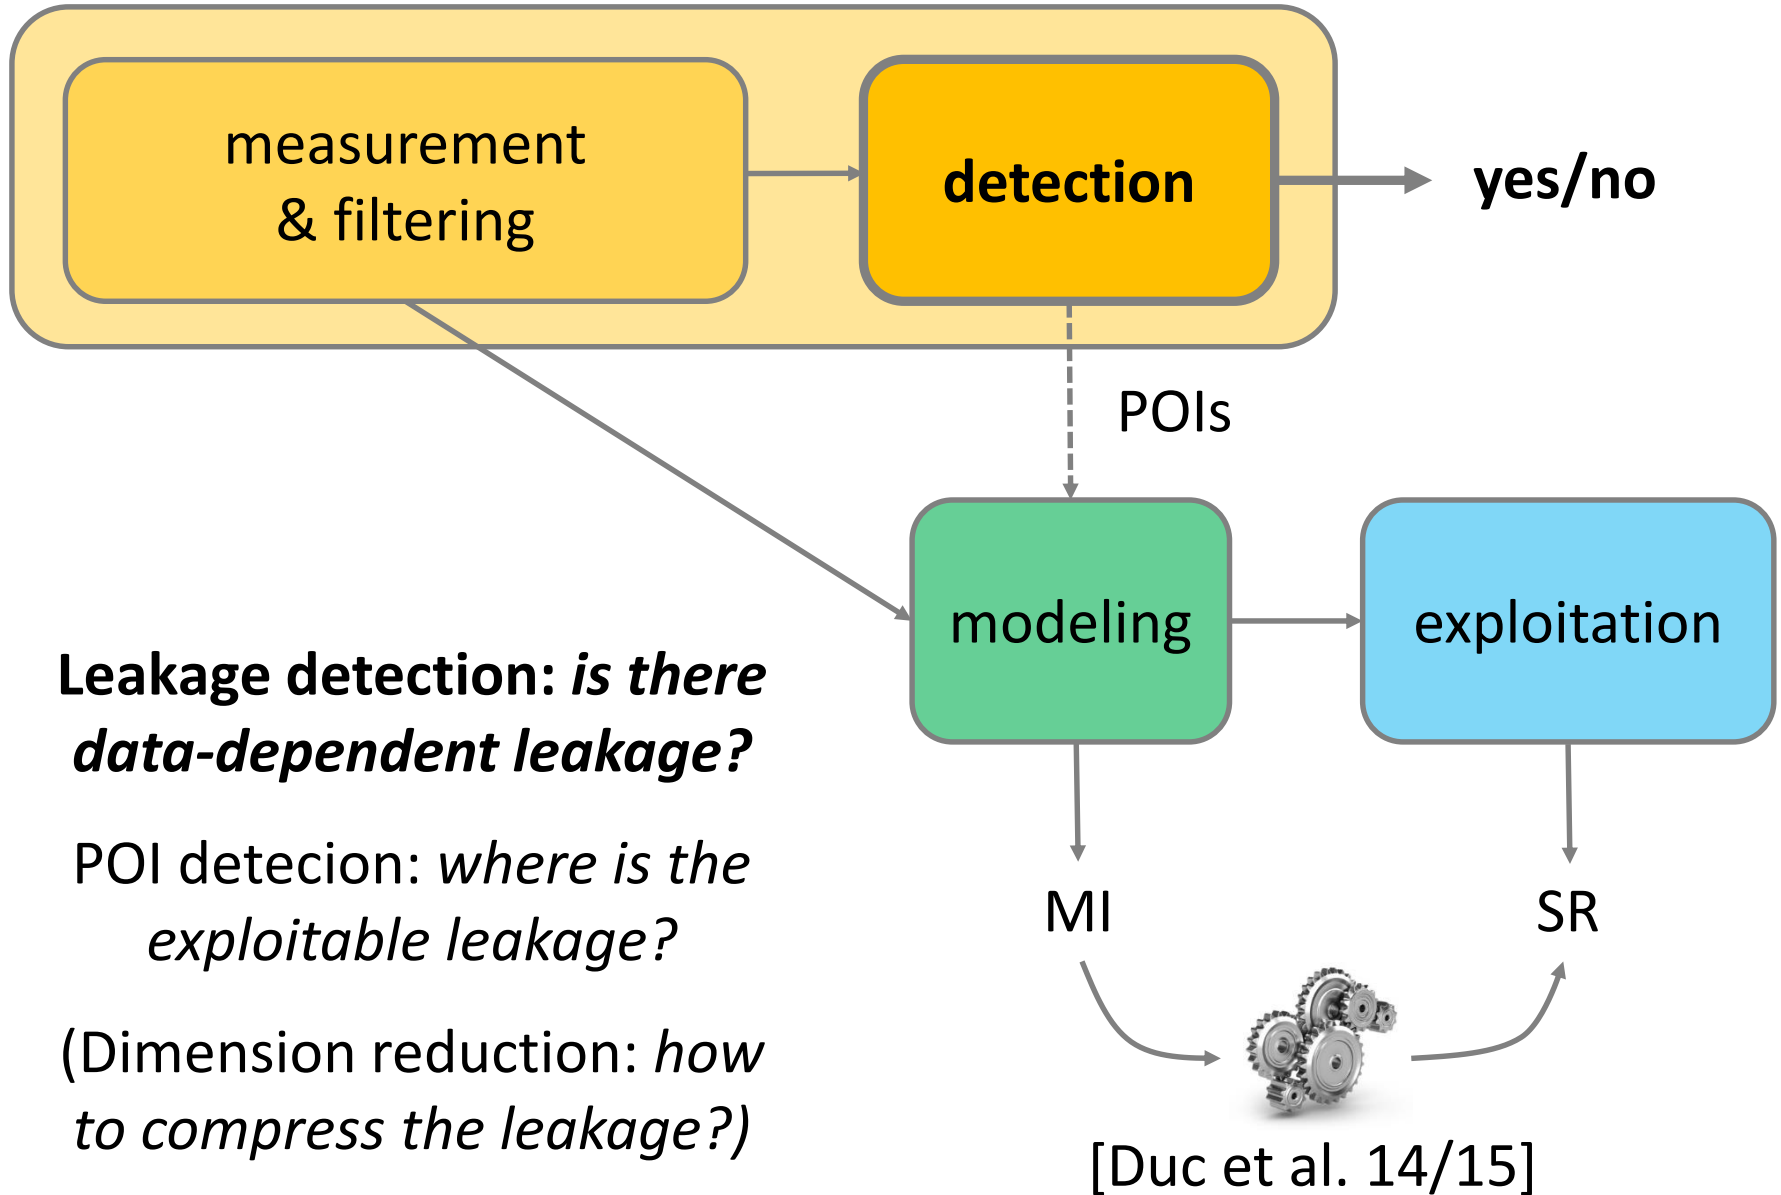






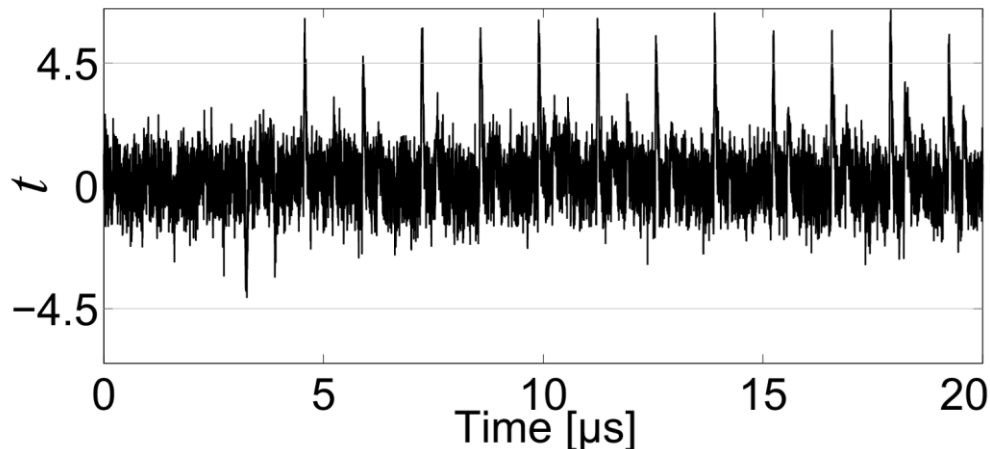
Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- **Measurement & preprocessing**
 - Filtering, **leakage/POI detection**, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs



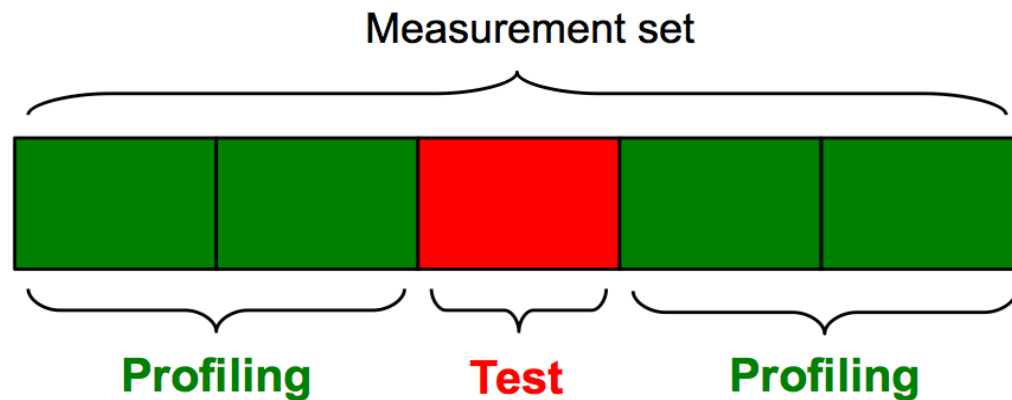
- AES Rijndael example
 - 128-bit key fixed
 - N_r traces with random plaintexts
 - N_f traces with a fixed plaintexts
 - Apply Student's t-test to the f&r classes:
 - $$\Delta(t) = [\hat{\mu}_f(t) - \hat{\mu}_r(t)] / [(\hat{\sigma}_f^2(t)/N_f) + (\hat{\sigma}_r^2(t)/N_r)]$$

- AES Rijndael example
 - 128-bit key fixed
 - N_r traces with random plaintexts
 - N_f traces with a fixed plaintexts
 - Apply Student's t-test to the f&r classes:
 - $$\Delta(t) = [\hat{\mu}_f(t) - \hat{\mu}_r(t)] / [(\hat{\sigma}_f^2(t)/N_f) + (\hat{\sigma}_r^2(t)/N_r)]$$
- Large t statistic: « some data dependency detected »

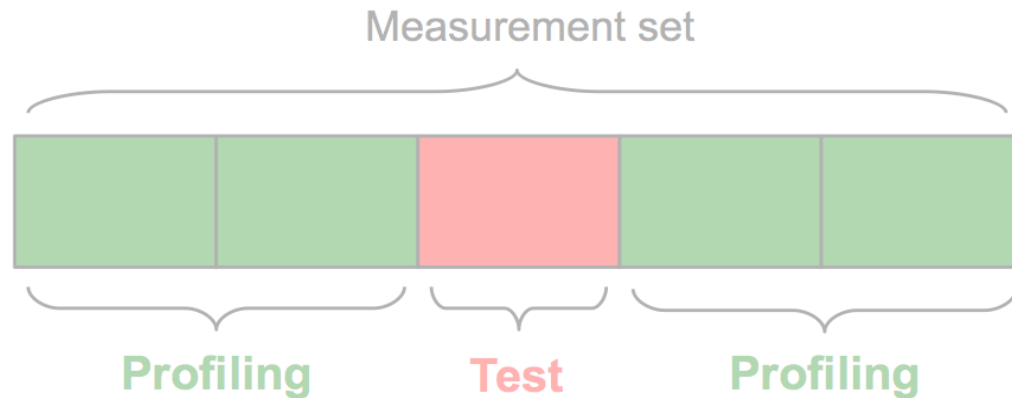


- ρ -test
 - 128-bit key fixed
 - N traces with random plaintexts
 - Targets an enumerable intermediate value X
 - Estimate Pearson's coefficient: $\hat{r}(t) = \hat{\rho}(L_X(t), \hat{model}_t(X))$

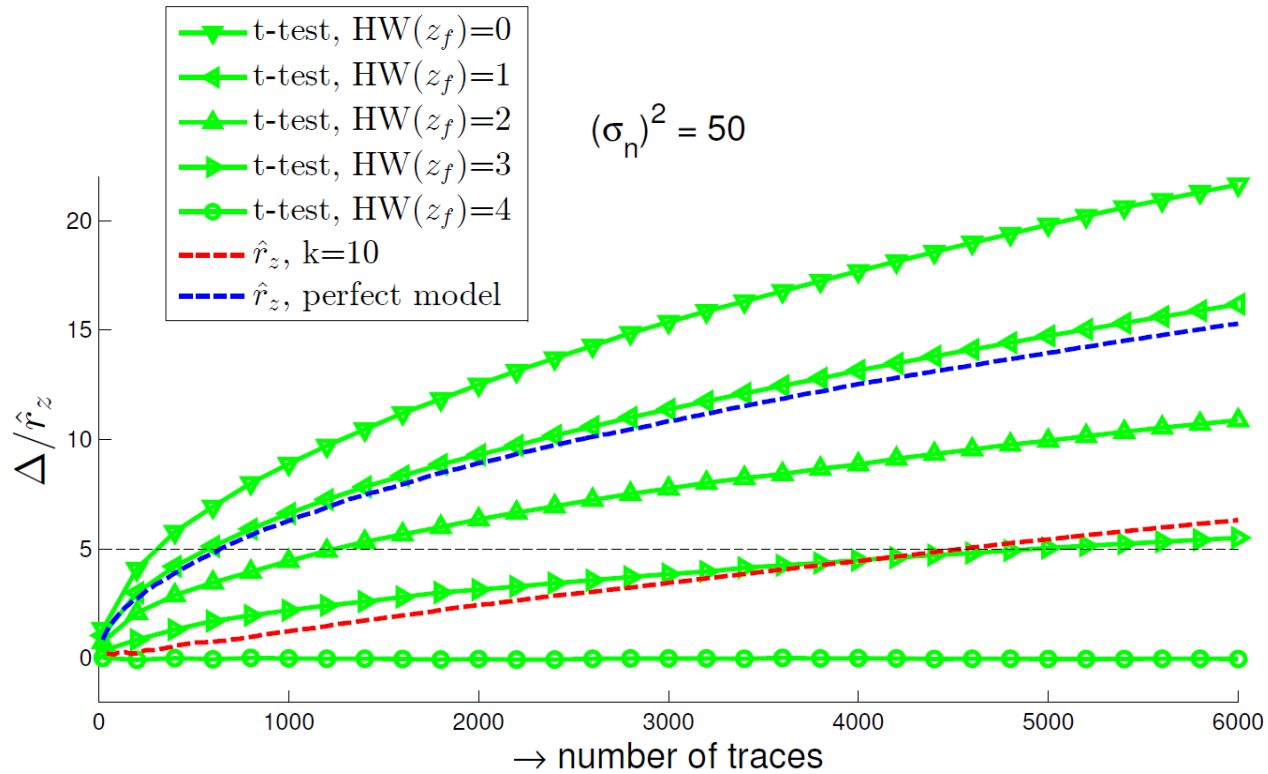
- ρ -test
 - 128-bit key fixed
 - N traces with random plaintexts
 - Targets an enumerable intermediate value X
 - Estimate Pearson's coefficient: $\hat{r}(t) = \hat{\rho}(L_X(t), \text{model}_t(X))$
- Model estimated with cross-validation

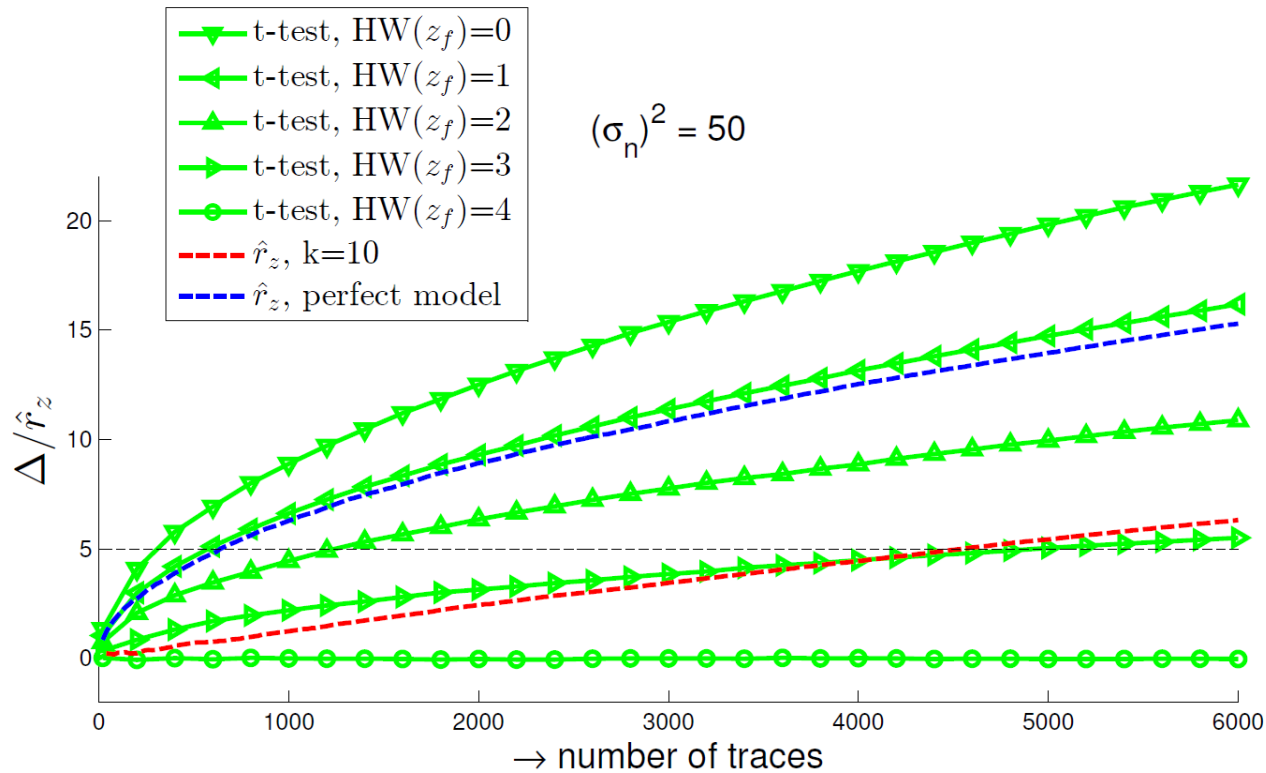


- ρ -test
 - 128-bit key fixed
 - N traces with random plaintexts
 - Targets an enumerable intermediate value X
 - Estimate Pearson's coefficient: $\hat{r}(t) = \hat{\rho}(L_X(t), \text{model}_t(X))$
- Model estimated with cross-validation

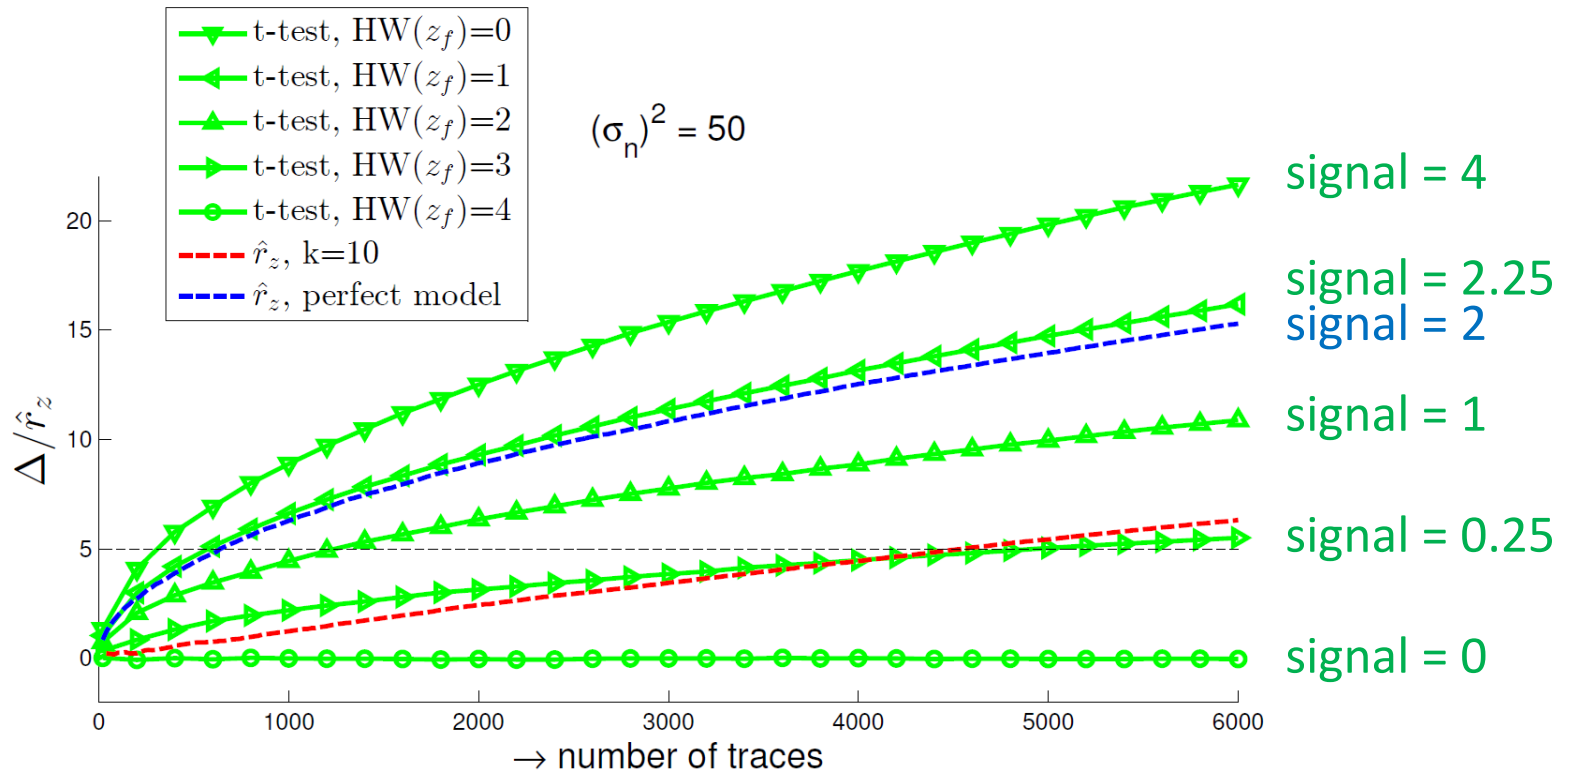


- Hypothesis test (Z transform): $\frac{1}{2} \times \ln \left(\frac{1+\hat{r}(t)}{1-\hat{r}(t)} \right) \sim N \left(0, \frac{1}{\sqrt{N-3}} \right)$

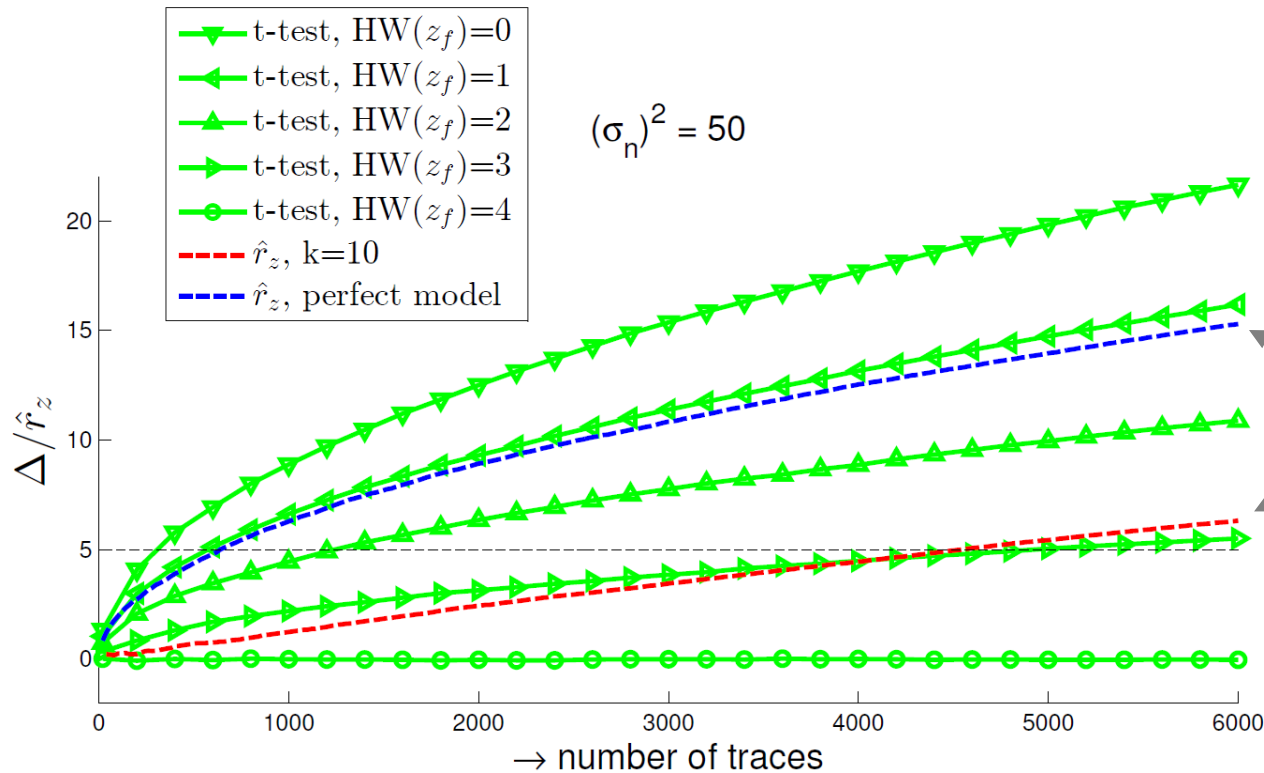




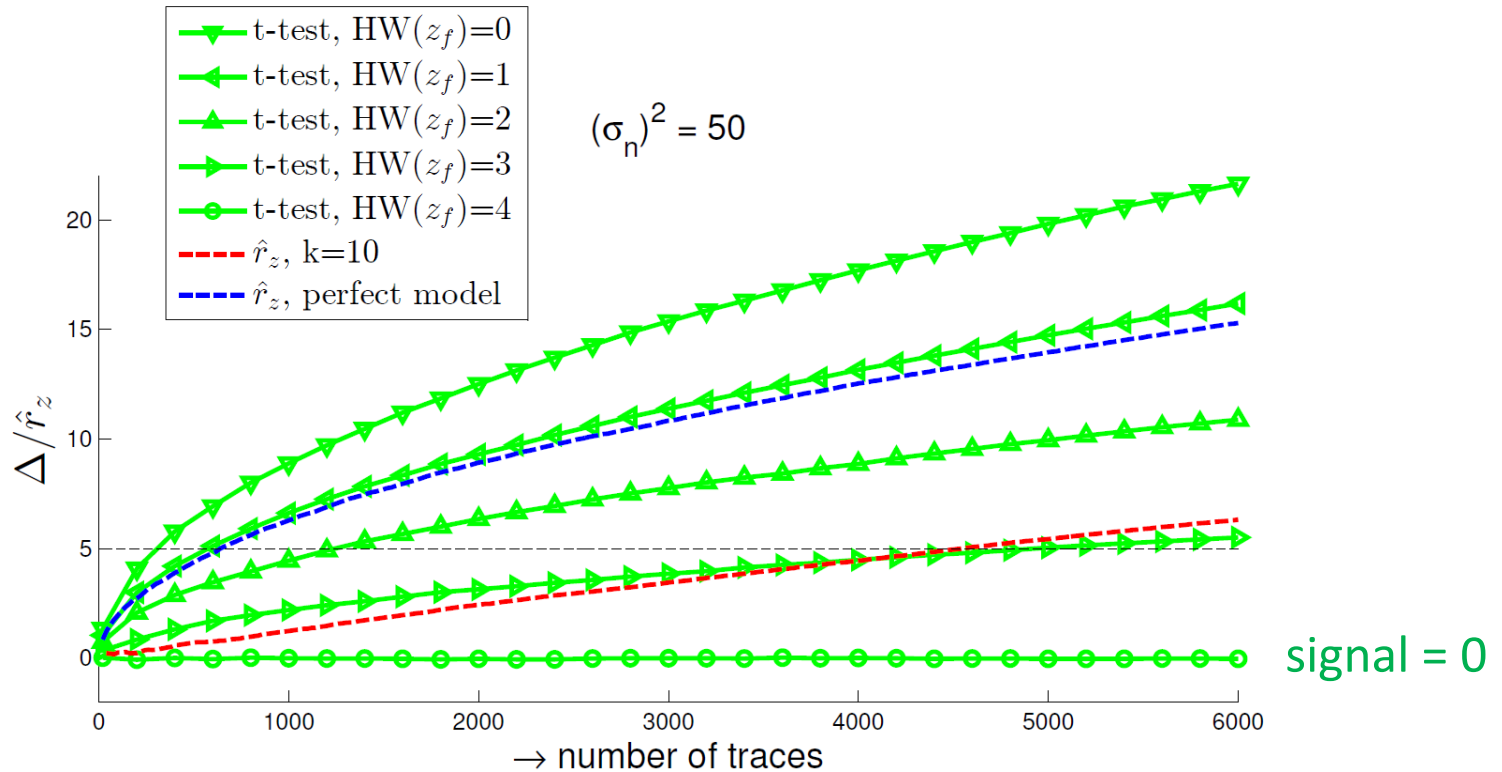
- CRI's t-test pro: sampling complexity!



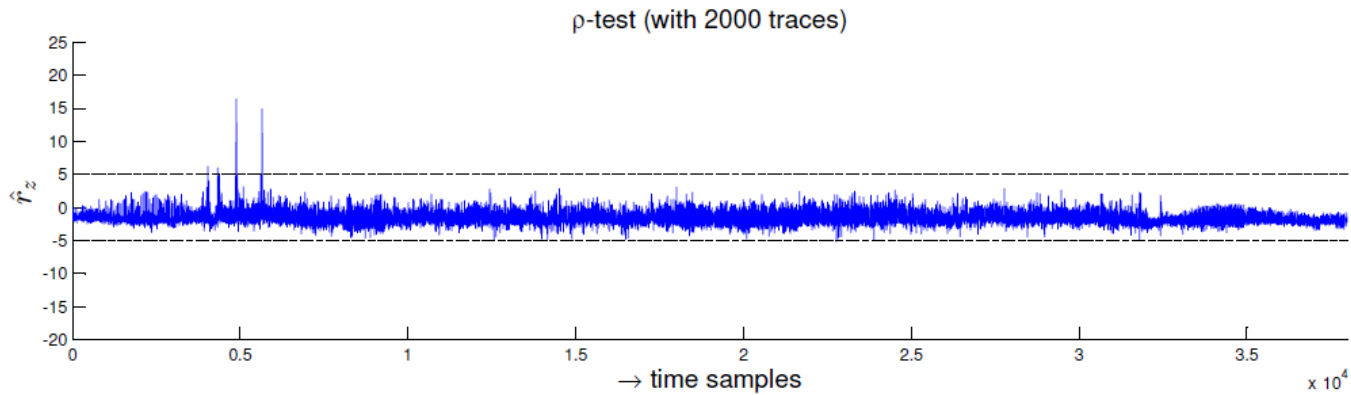
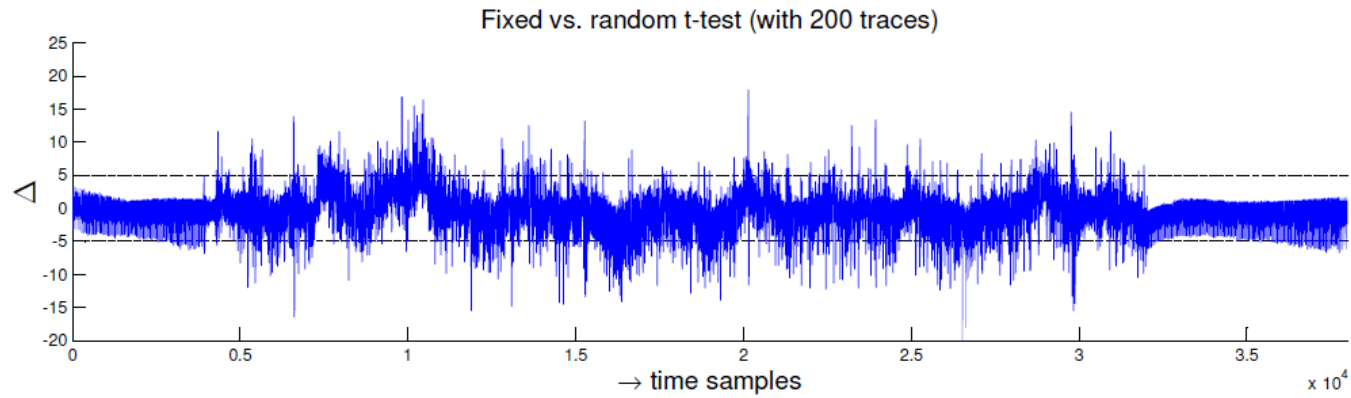
- CRI's t-test pro: sampling complexity!
 - Better **signal** for well-chosen fixed classes

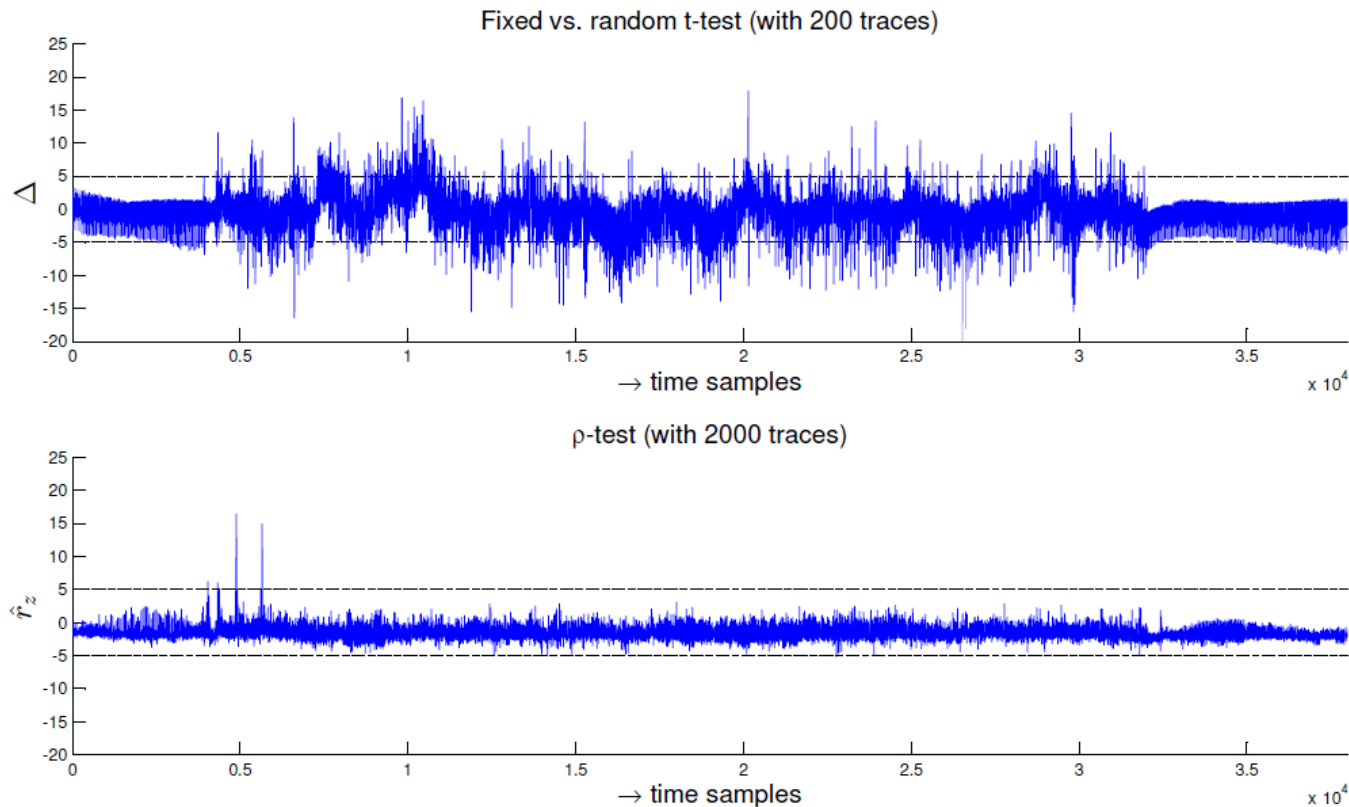


- CRI's t-test pro: sampling complexity!
 - Better **signal** for well-chosen fixed classes
 - Easier estimation (**2** classes vs. **256** classes)

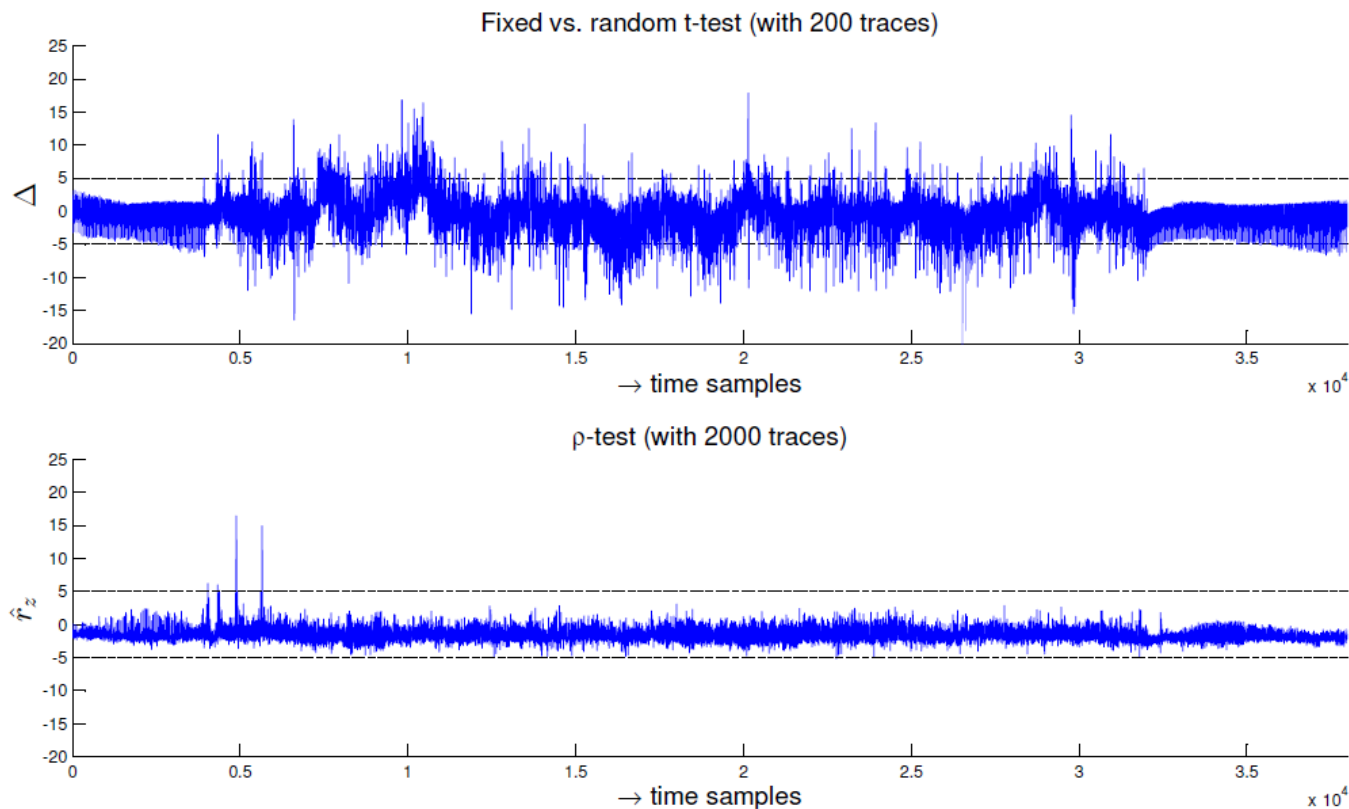


- CRI's t-test con: possible false negative!
 - Possibly **no signal** for badly-chosen fixed classes



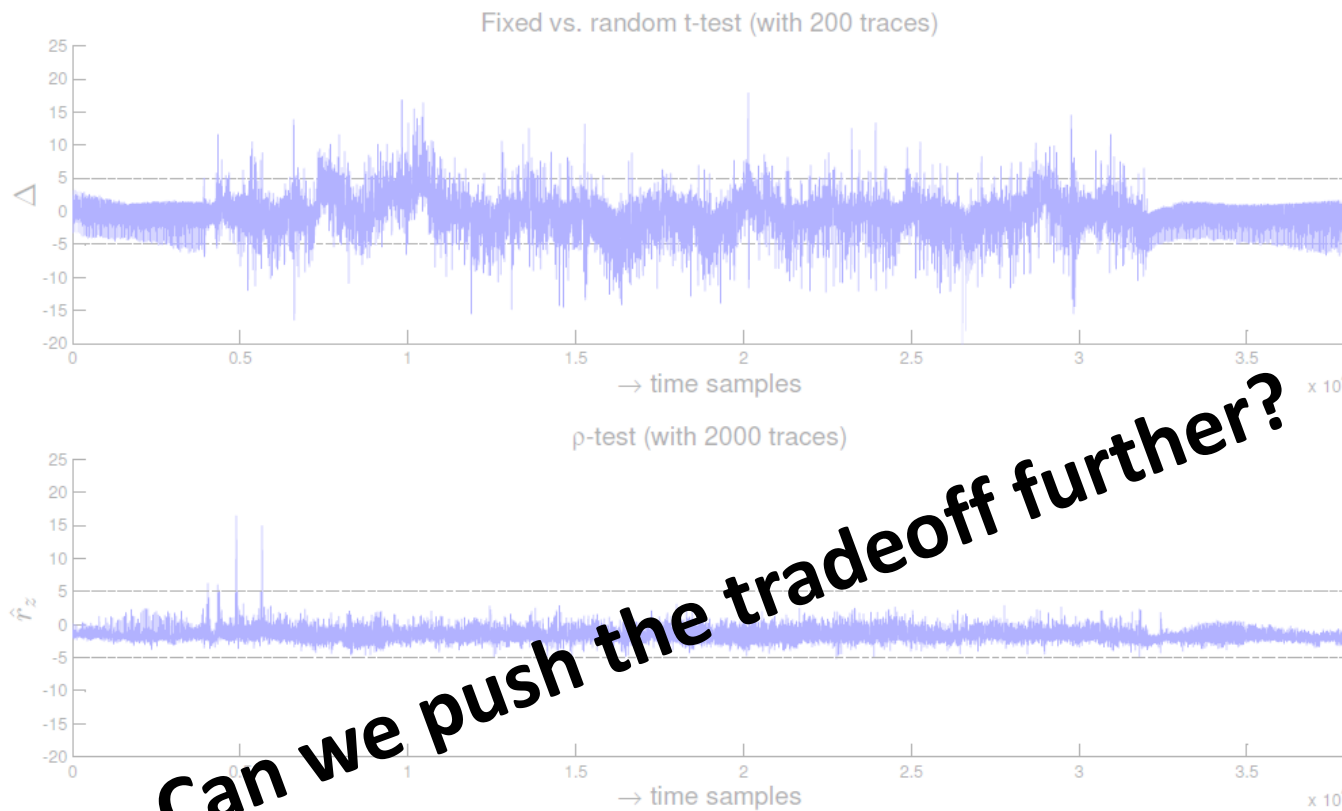


- t-test false negatives do not hurt (integrated over time)
- But there are many false positives (w.r.t. DPA)
- Only the value of ρ is connected with key recovery SR



sampling complexity vs.
informativeness tradeoff

- t-test false negatives do not hurt (integrated over time)
- But there are many false positives (w.r.t. DPA)
- Only the value of ρ is connected with key recovery SR

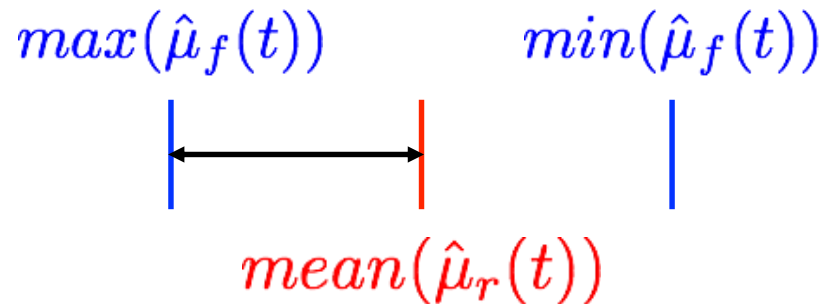


Can we push the tradeoff further?

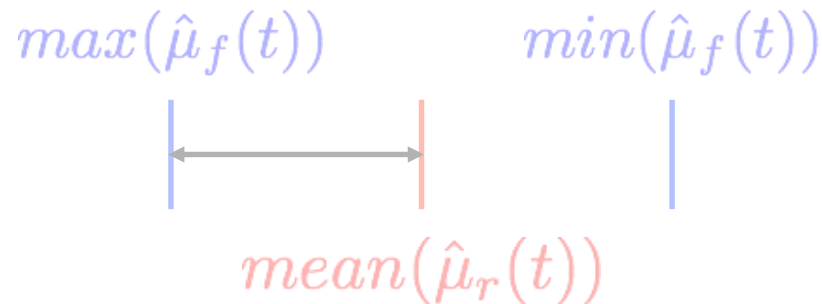
sampling complexity vs.
informativeness tradeoff

- t-test false negatives do not hurt (integrated over time)
- But there are many false positives (w.r.t. DPA)
- Only the value of ρ is connected with key recovery SR

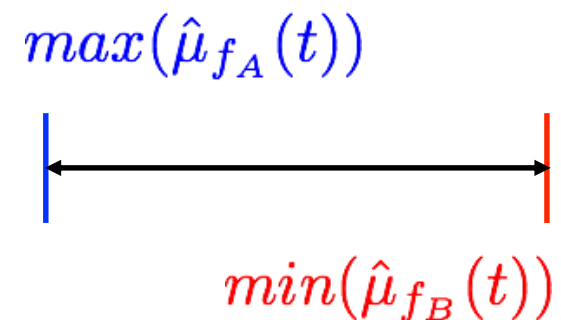
- CRI's fixed vs. random test (e.g. HW leakages)
 - Maximum HW difference observed = 4
 - “Algorithmic noise” due to the random class

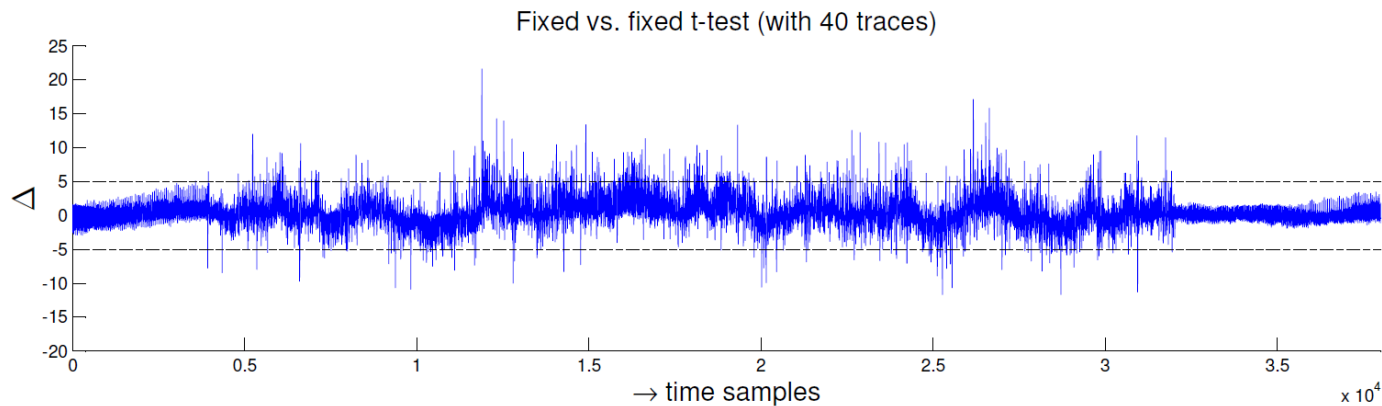
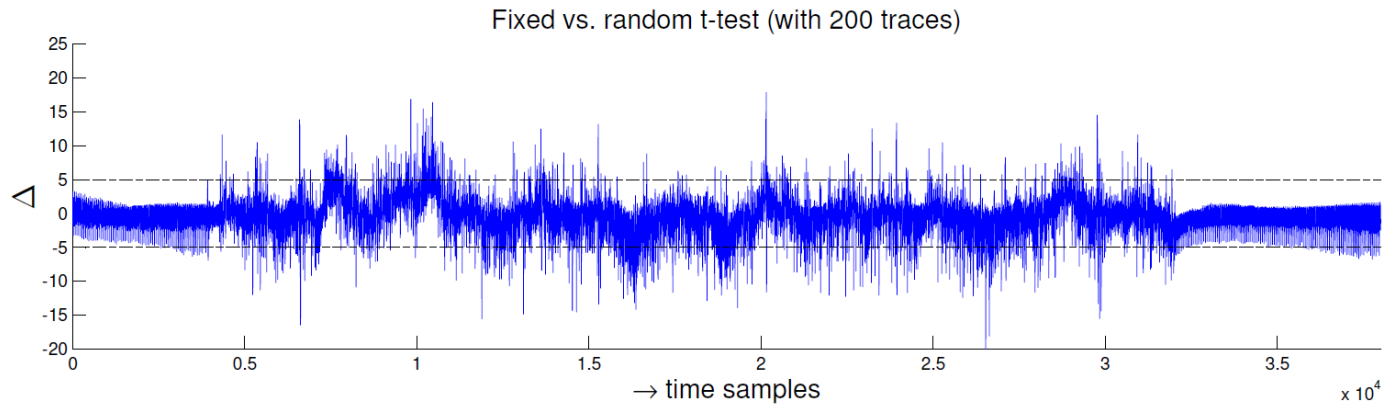


- CRI's fixed vs. random test (e.g. HW leakages)
 - Maximum HW difference observed = 4
 - “Algorithmic noise” due to the random class

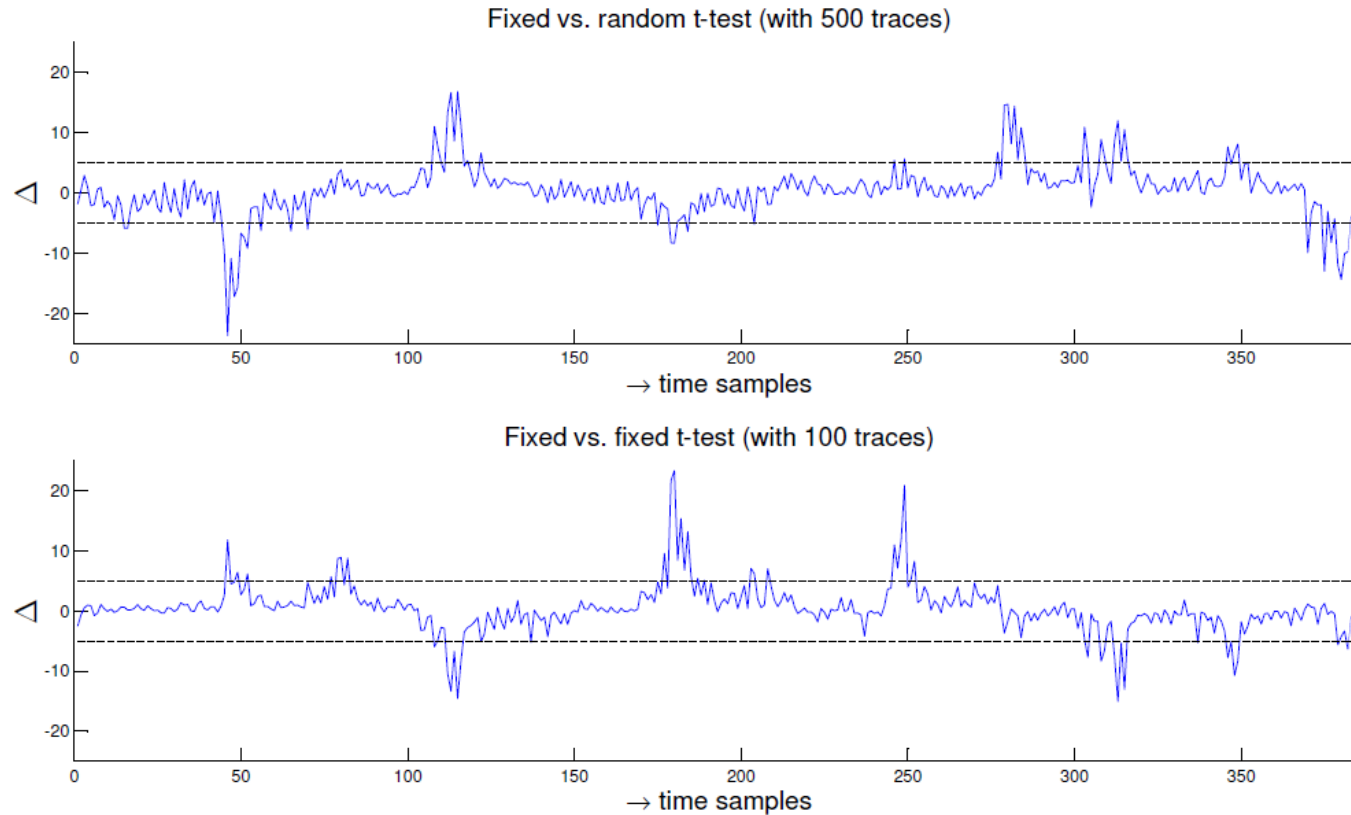


- Natural extension: *fixed vs. fixed test*
 - Maximum HW difference = 8
 - Average signal multiplied 2
 - No algorithmic noise!

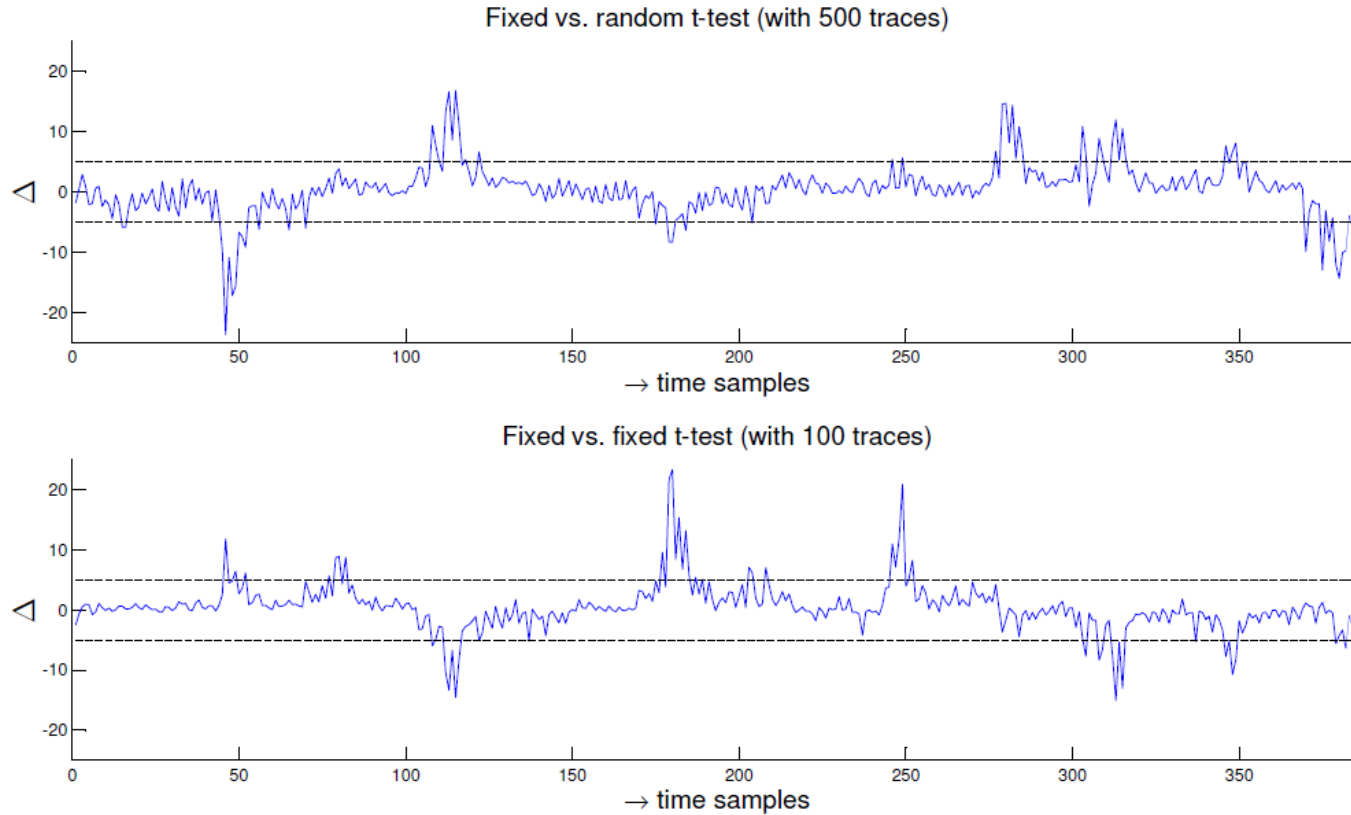




- Gain dominated by the increased signal
- Reduction of the sampling complexity by a factor ≈ 5



- Gain dominated by the reduced noise
- Reduction of the sampling complexity by a factor ≈ 5



F vs. F better deal on average; similar risks of false positives & negatives

- Gain dominated by the reduced noise
- Reduction of the sampling complexity by a factor ≈ 5

Outline

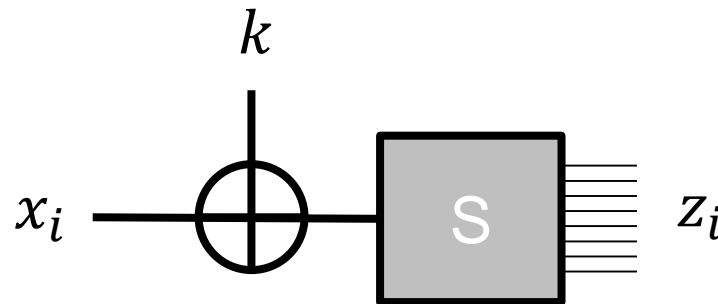
- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- **Predictions & modeling**
 - **Profiled vs. non-profiled separation**, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs

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

- General case: profiled DPA
 - Build “*templates*”, i.e. $\hat{f}(l_i|k, x_i)$
 - e.g., Gaussian, regression-based
 - Which directly leads to $\widehat{\text{Pr}}[k|l_i, x_i]$
- “Simplified” case: non-profiled DPA
 - Just assumes some model
 - e.g., CPA with $m_i^{k^*} = \text{HW}(z_i)$
 - e.g., DPA with $m_i^{k^*} = z_i[1]$
 - ...

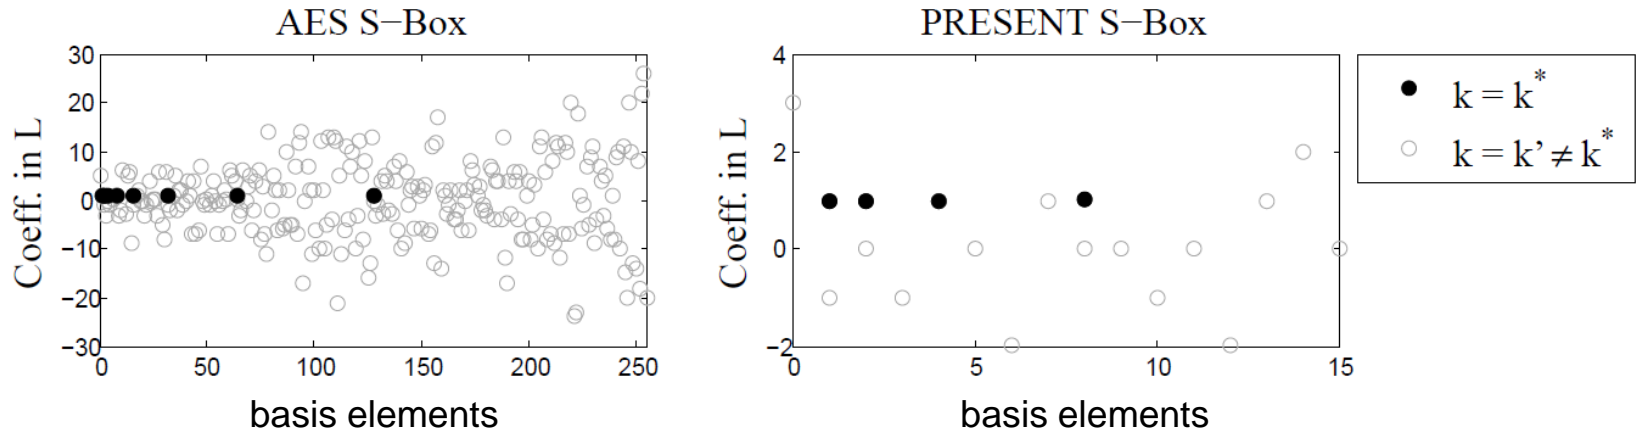
- Only profiled DPA is *guaranteed* to succeed!

- Only profiled DPA is *guaranteed* to succeed!



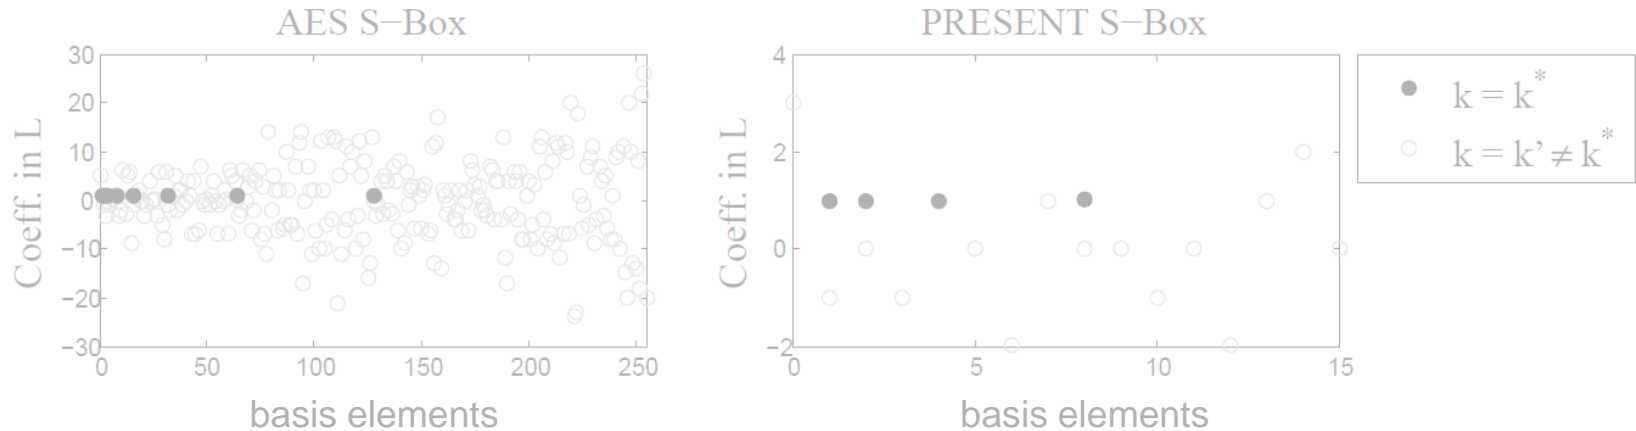
- Regression: $L(z_i) \approx F(z_i) + N_i$ with:
 - Linear basis (i.e., 8 bits)
 - Quadratic basis (i.e., add 28 products)
 - ...
 - Full basis (i.e., 256 elements)

- e.g., if $L(z_i) = HW(z_i)$



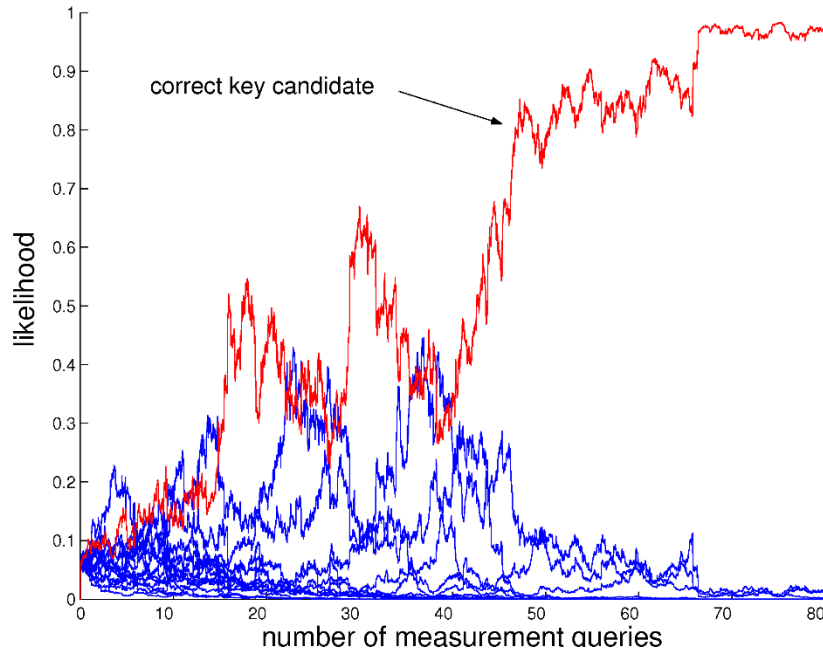
- Full basis perfectly explains any L by overfitting
 - Even for incorrect key candidates!

- e.g., if $L(z_i) = HW(z_i)$



- Full basis perfectly explains any L by overfitting
 - Even for incorrect key candidates!
- ⇒ Non-profiled DPA needs a good assumption
- e.g., the model is linear, simple, ...
 - This, *in general*, is only provided by profiling

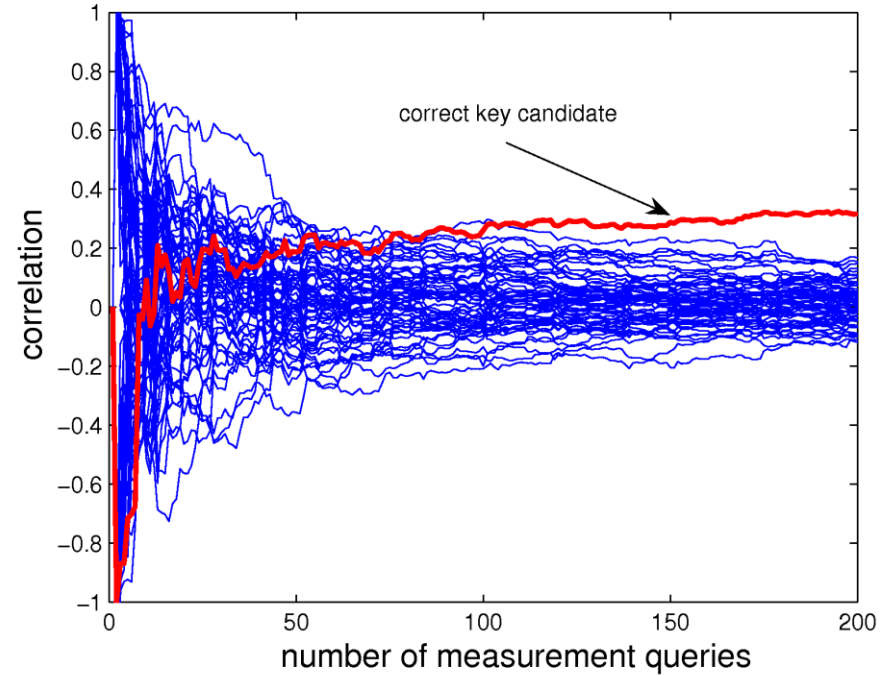
Gaussian templates



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

- More efficient (better model)
- Outputs probabilities

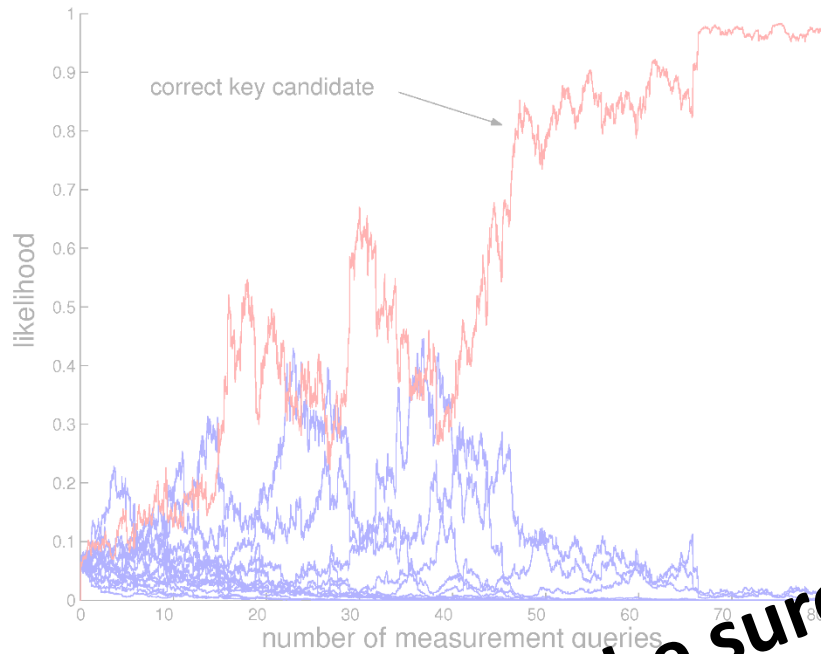
CPA



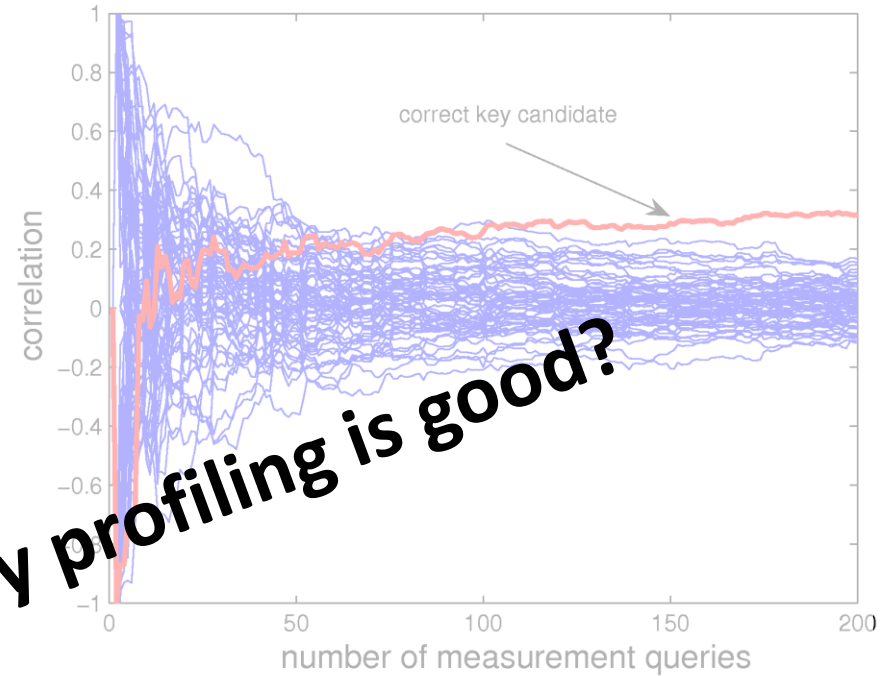
$$\tilde{k} = \operatorname{argmax}_{k^*} \frac{E(L \cdot M^{k^*}) - E(L) \cdot E(M^{k^*})}{\sigma(L) \cdot \sigma(M^{k^*})}$$

- Less efficient (worse model)
- Outputs scores

Gaussian templates



CPA



How can I be sure my profiling is good?

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

- More efficient (better model)
- Outputs probabilities

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

- Less efficient (worse model)
- Outputs scores

Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- **Predictions & modeling**
 - Profiled vs. non-profiled separation, **leakage certification**
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs

- A model is optimal if $\hat{\text{Pr}}_{model} [l|k] = \text{Pr}_{chip} [l|k]$

- A model is optimal if $\hat{\Pr}_{model} [l|k] = \Pr_{chip} [l|k]$

⇒ Theory would say it is ε -close to optimal if

$$SD(\hat{\Pr}_{model} [l|k], \Pr_{chip} [l|k]) < \varepsilon$$

- (with SD a statistical distance)

- A model is optimal if $\hat{\Pr}_{model} [l|k] = \Pr_{chip} [l|k]$

⇒ Theory would say it is ε -close to optimal if

$$SD(\hat{\Pr}_{model} [l|k], \Pr_{chip} [l|k]) < \varepsilon$$

- (with SD a statistical distance)
- **Convenient since ε would quantify the loss**
 - That could be reported in SR bounds [DFS15]

- A model is optimal if $\hat{\Pr}_{model} [l|k] = \Pr_{chip} [l|k]$

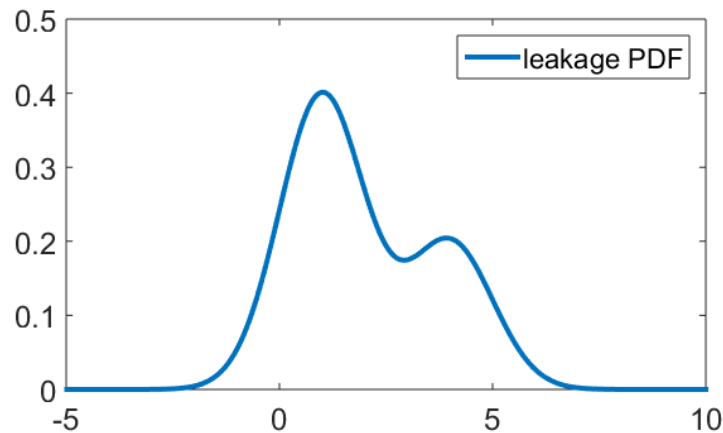
⇒ Theory would say it is ε -close to optimal if

$$SD(\hat{\Pr}_{model} [l|k], \Pr_{chip} [l|k]) < \varepsilon$$

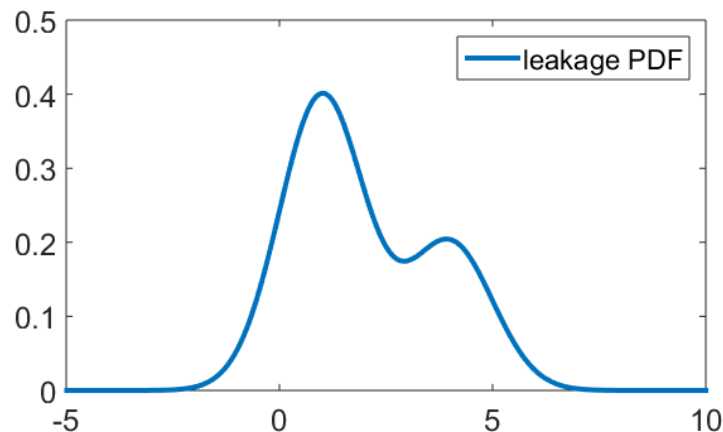
- (with SD a statistical distance)
- Convenient since ε would quantify the loss
 - That could be reported in SR bounds [DFS15]
- **Problem:** $\Pr_{chip} [l|k]$ is unknown

- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.

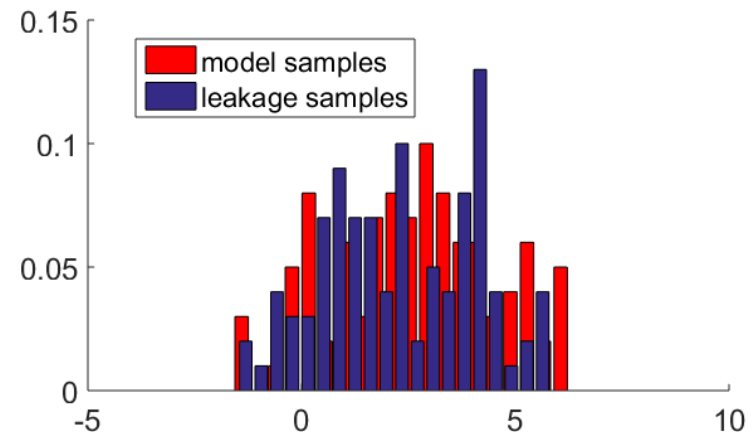
- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.
- **Example:**



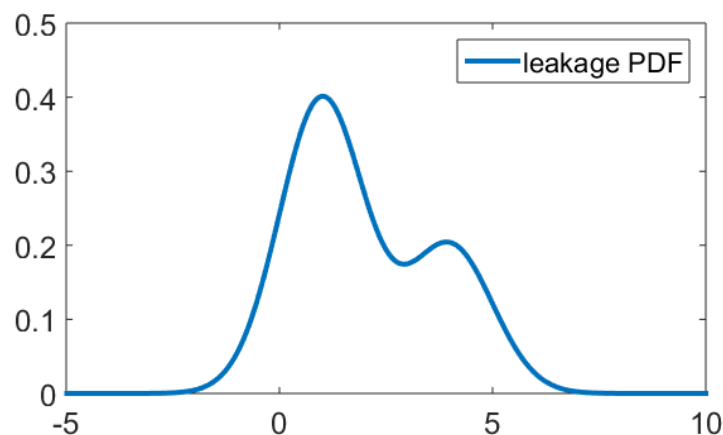
- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.
- Example:



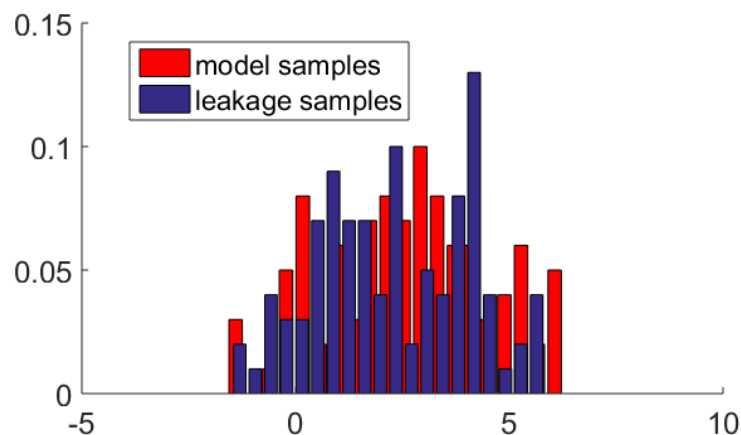
No samples



- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.
- Example:



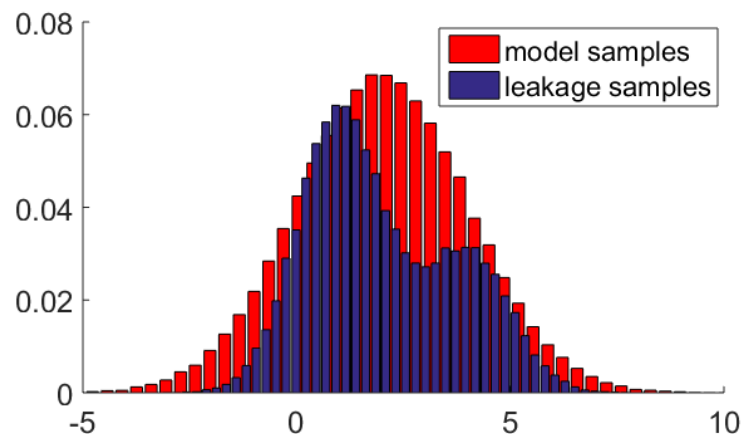
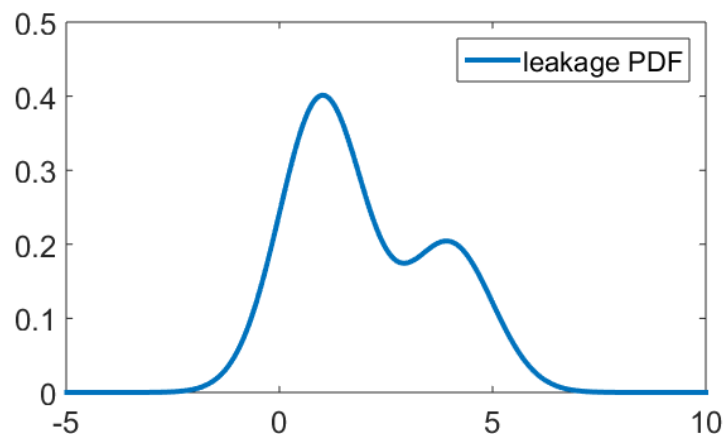
estimation errors dominate



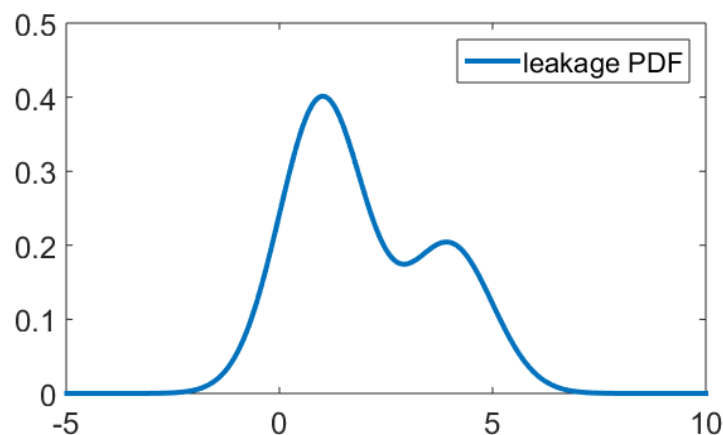
⇒ need to measure more

- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.
- Example:

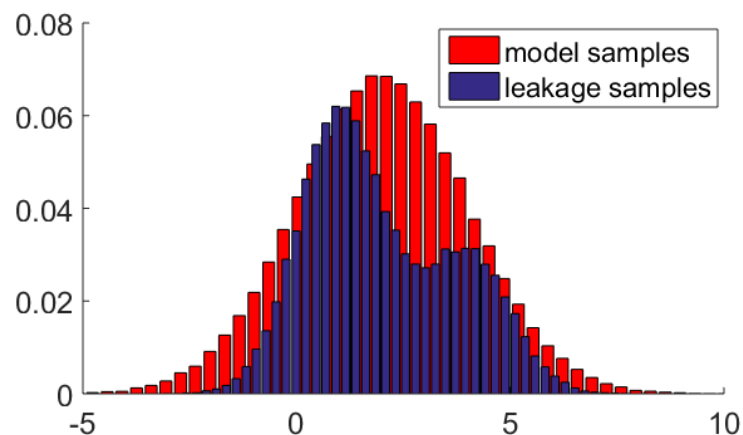
$N_1 > N_0$ samples



- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.
- Example:

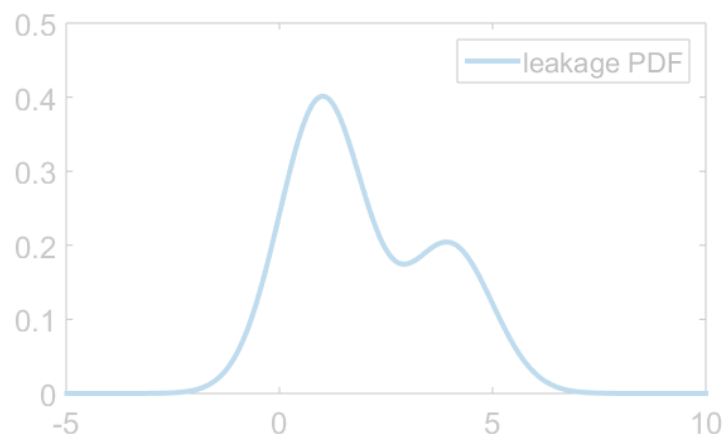


assumption errors dominate

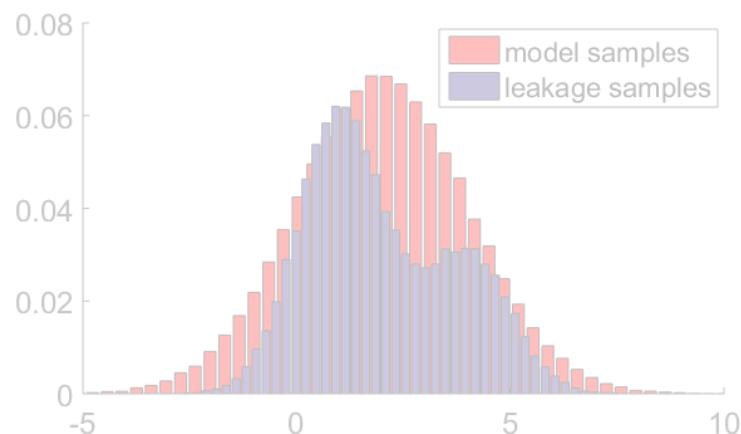


⇒ need another model

- Distinguish estimation & assumption errors
 - Recall estimation errors decrease with # meas.
- Example:



assumption errors dominate



⇒ need another model

⇒ good enough model: *ass. err* \ll *est. err.* given N

- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Output a p-value $p(N)$
 - Small p 's indicate hyp. is likely incorrect

- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Output a p-value $p(N)$ **Eval. lab. limit**
 - Small p's indicate hyp. is likely incorrect

- Test the hypothesis that

$$\hat{\Pr}_{model} [l|k] \stackrel{N}{=} \Pr_{chip} [l|k]$$

- Taking advantage of cross-validation



modeling samples



test samples



- Output a p-value $p(N)$ Eval. lab. limit
 - Small p's indicate hyp. is likely incorrect

- Main drawback: cost (of sampling distributions)

- Compare moments (rather than distributions)

1. $\widehat{M}_d \stackrel{N}{\leftarrow} \widehat{\text{Pr}}_{model} [l|k]$

2. $\widetilde{M}_d \stackrel{N}{\leftarrow} \text{Pr}_{chip} [l|k]$

3. Test equality

$$\widehat{M}_d = \widetilde{M}_d$$

- Compare moments (rather than distributions)

1. $\hat{M}_d \stackrel{N}{\leftarrow} \hat{\Pr}_{model} [l|k]$

2. $\tilde{M}_d \stackrel{N}{\leftarrow} \Pr_{chip} [l|k]$

3. Test equality

$$\hat{M}_d = \tilde{M}_d$$

+ Can be done with simple univariate tests

- e.g., T-test (assuming \hat{M}_d, \tilde{M}_d are Gaussian)

- Compare moments (rather than distributions)

1. $\hat{M}_d \stackrel{N}{\leftarrow} \hat{\Pr}_{model} [l|k]$
2. $\tilde{M}_d \stackrel{N}{\leftarrow} \Pr_{chip} [l|k]$
3. Test equality
 $\hat{M}_d = \tilde{M}_d$

+ Can be done with simple univariate tests

- e.g., T-test (assuming \hat{M}_d, \tilde{M}_d are Gaussian)

– Is it theoretically sound? No!

- Compare moments (rather than distributions)

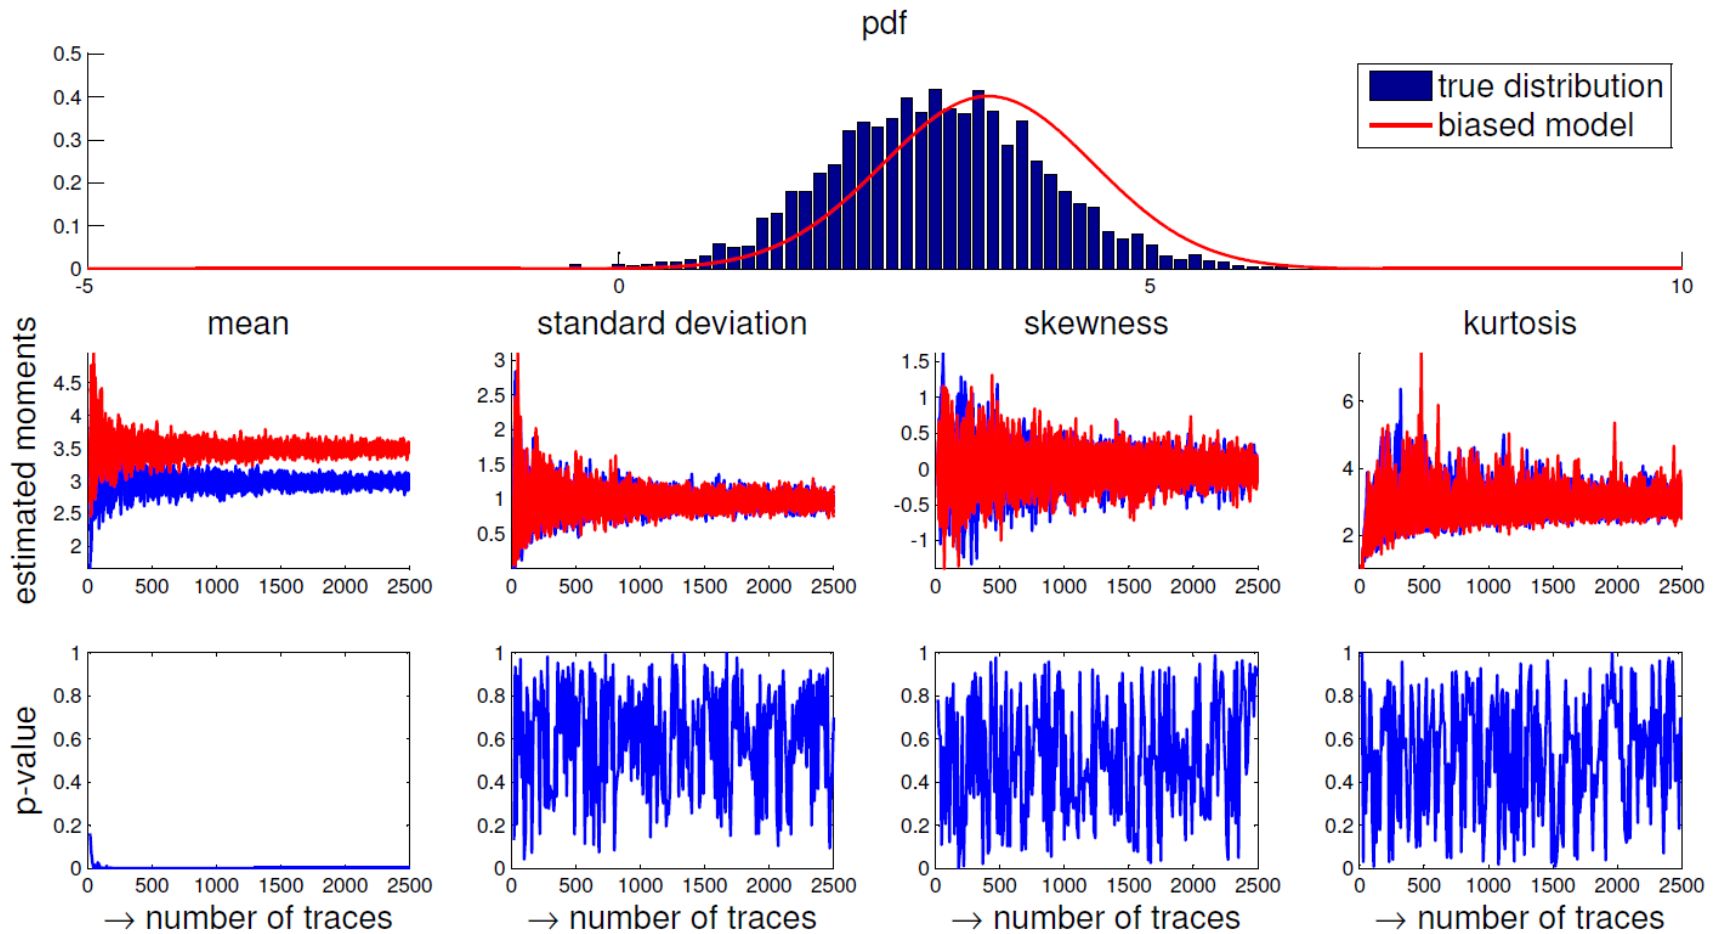
1. $\hat{M}_d \stackrel{N}{\leftarrow} \hat{\Pr}_{model} [l|k]$
2. $\tilde{M}_d \stackrel{N}{\leftarrow} \Pr_{chip} [l|k]$
3. Test equality
 $\hat{M}_d = \tilde{M}_d$

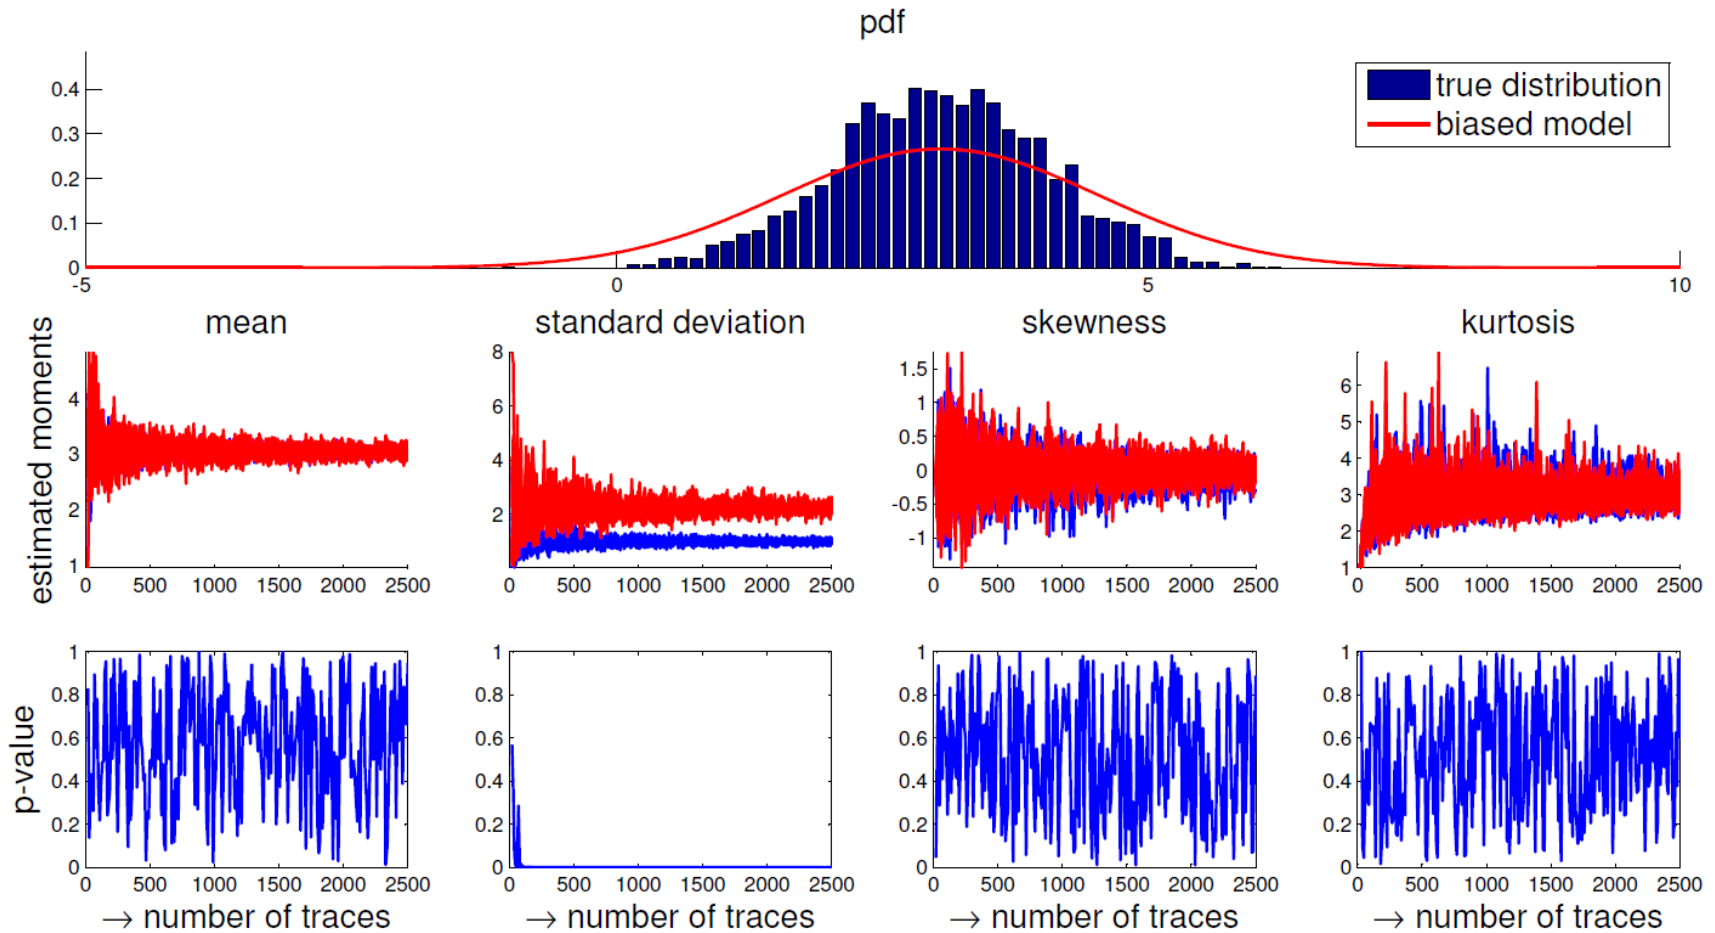
+ Can be done with simple univariate tests

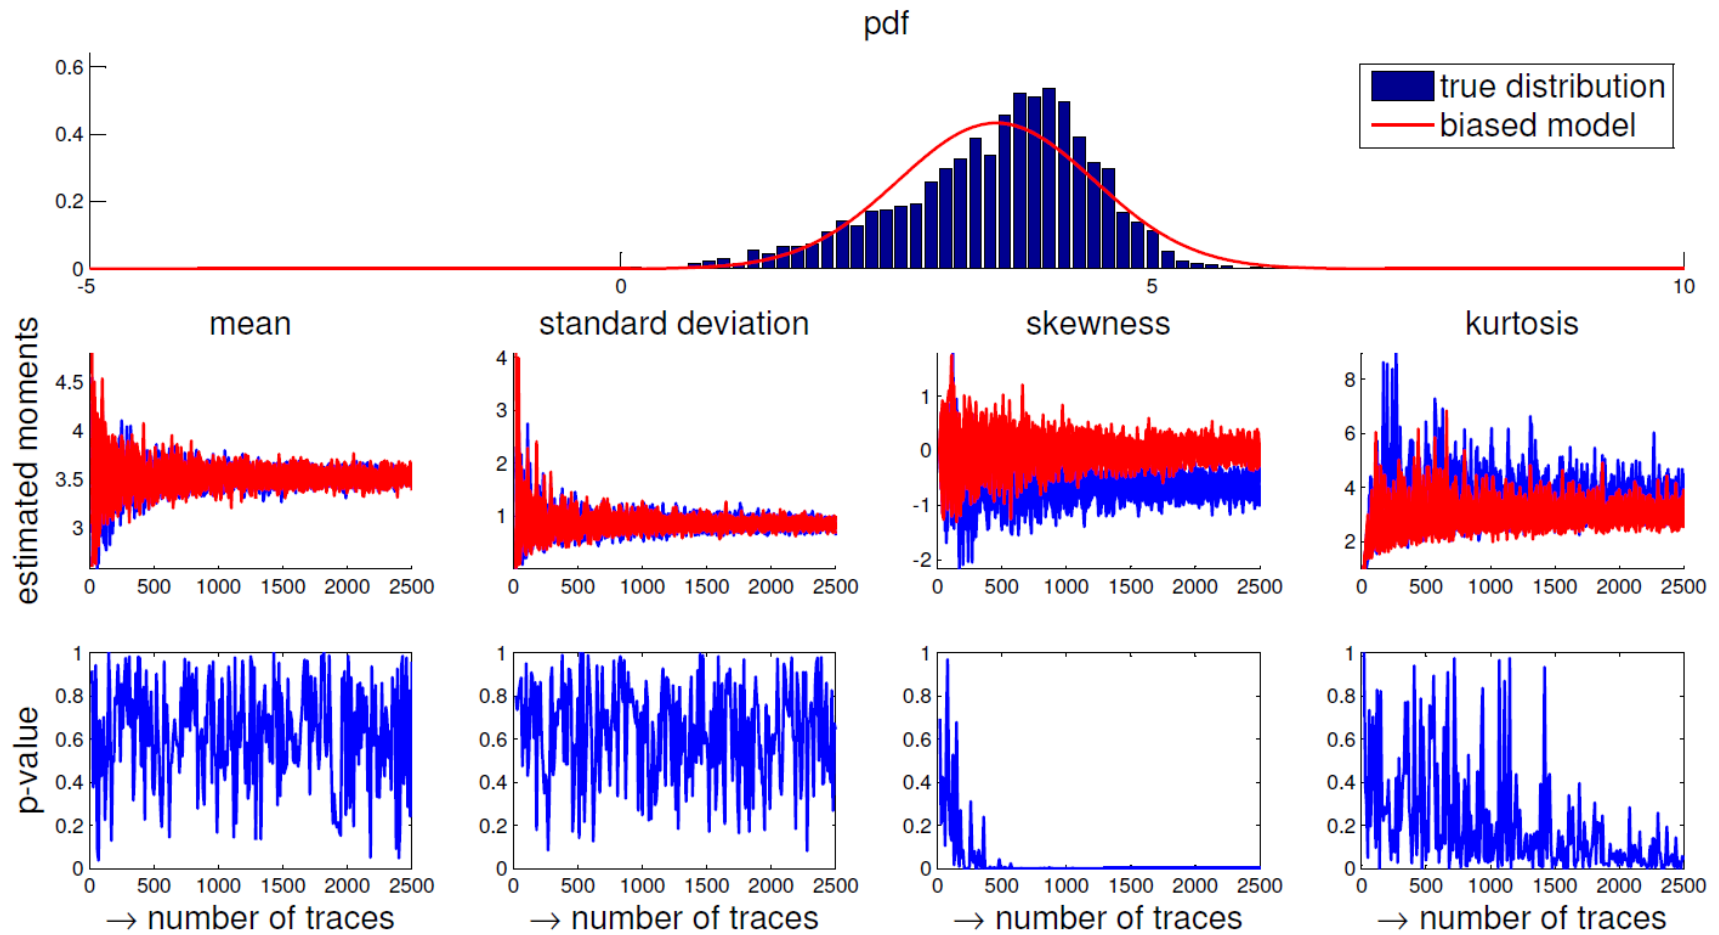
- e.g., T-test (assuming \hat{M}_d, \tilde{M}_d are Gaussian)

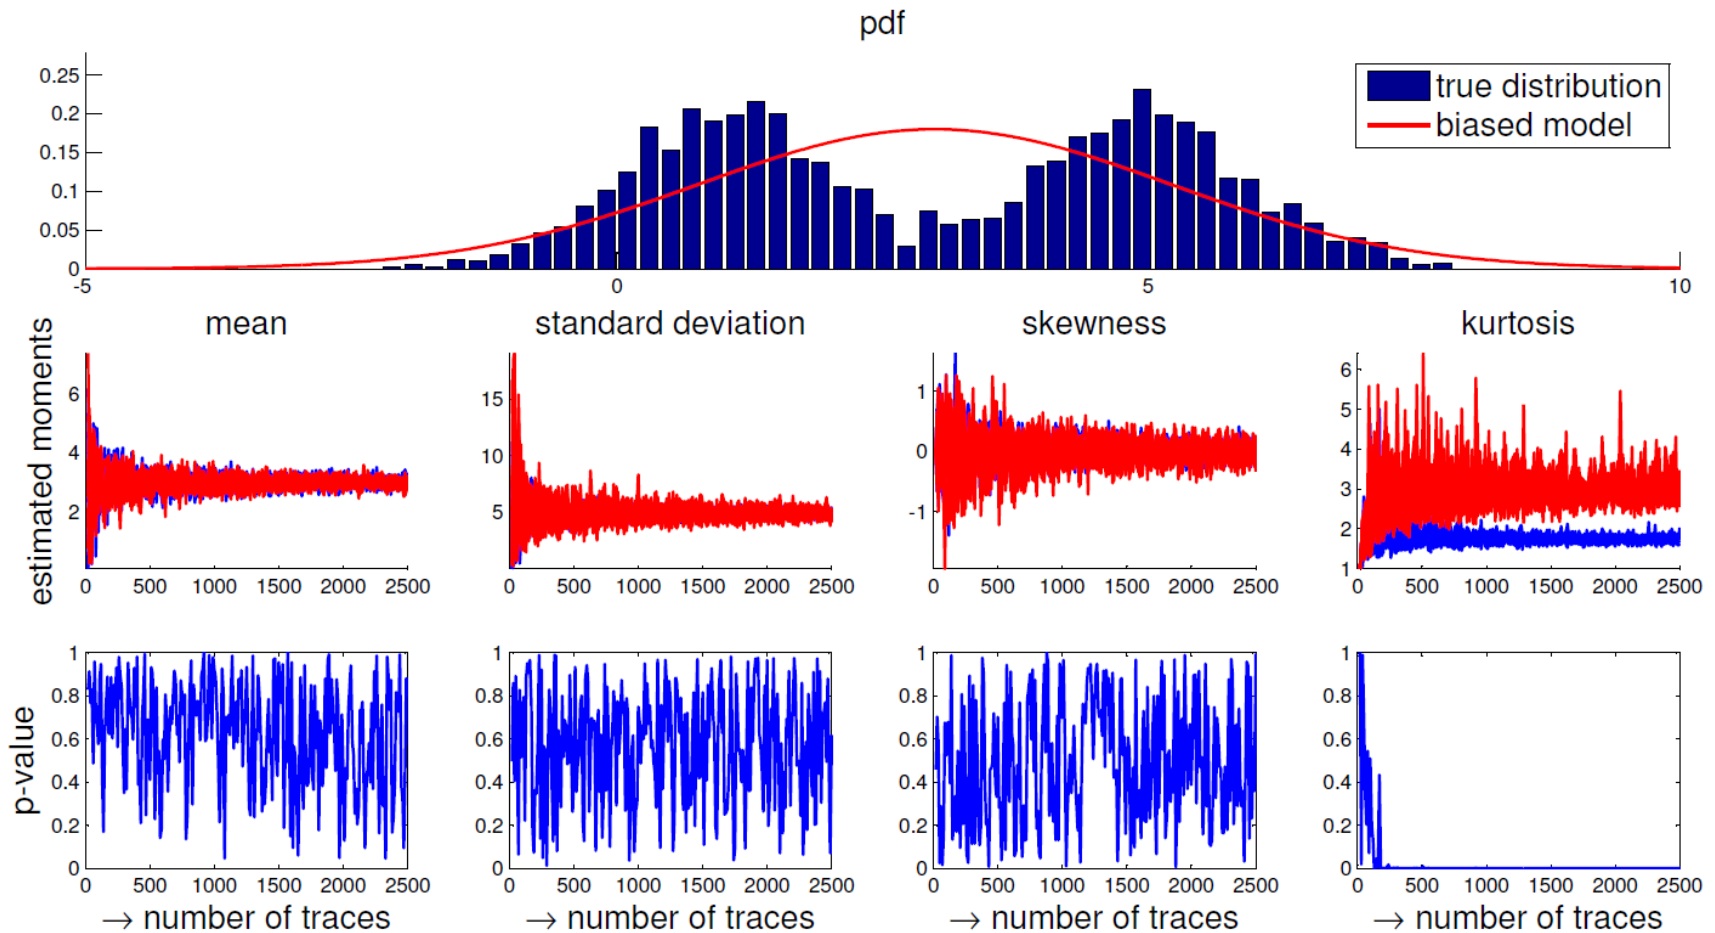
– Is it theoretically sound? No!

- But counterexamples are involved
- & SCA literature frequently does it
 - Leakage detection, HO attacks, ...





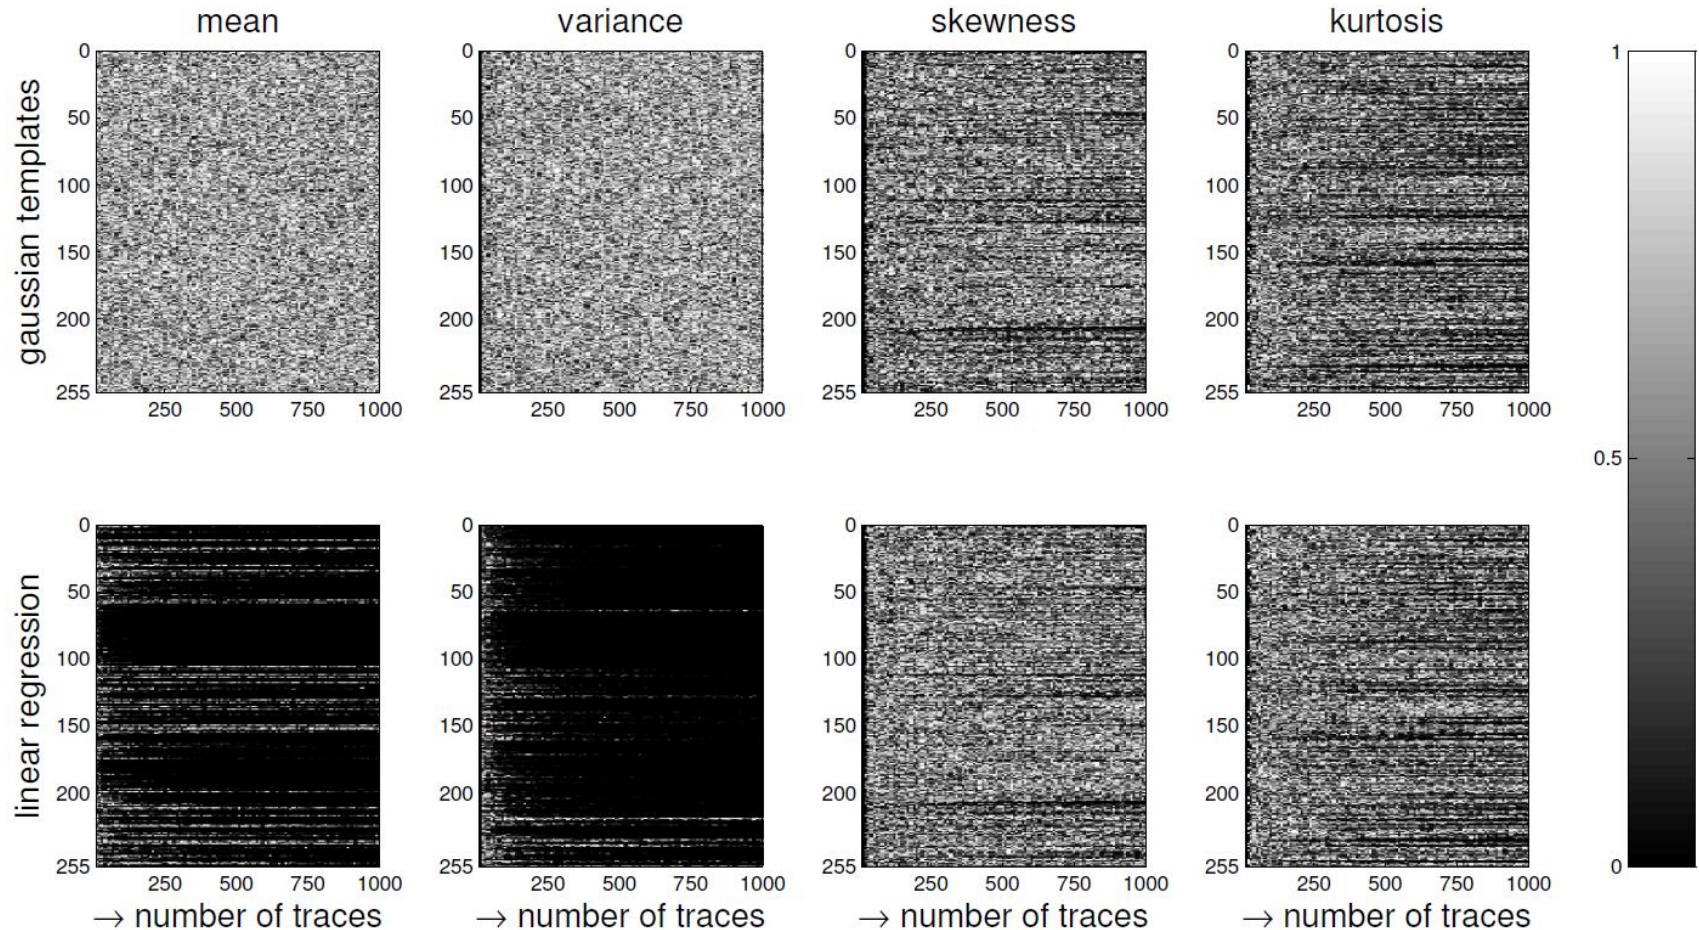




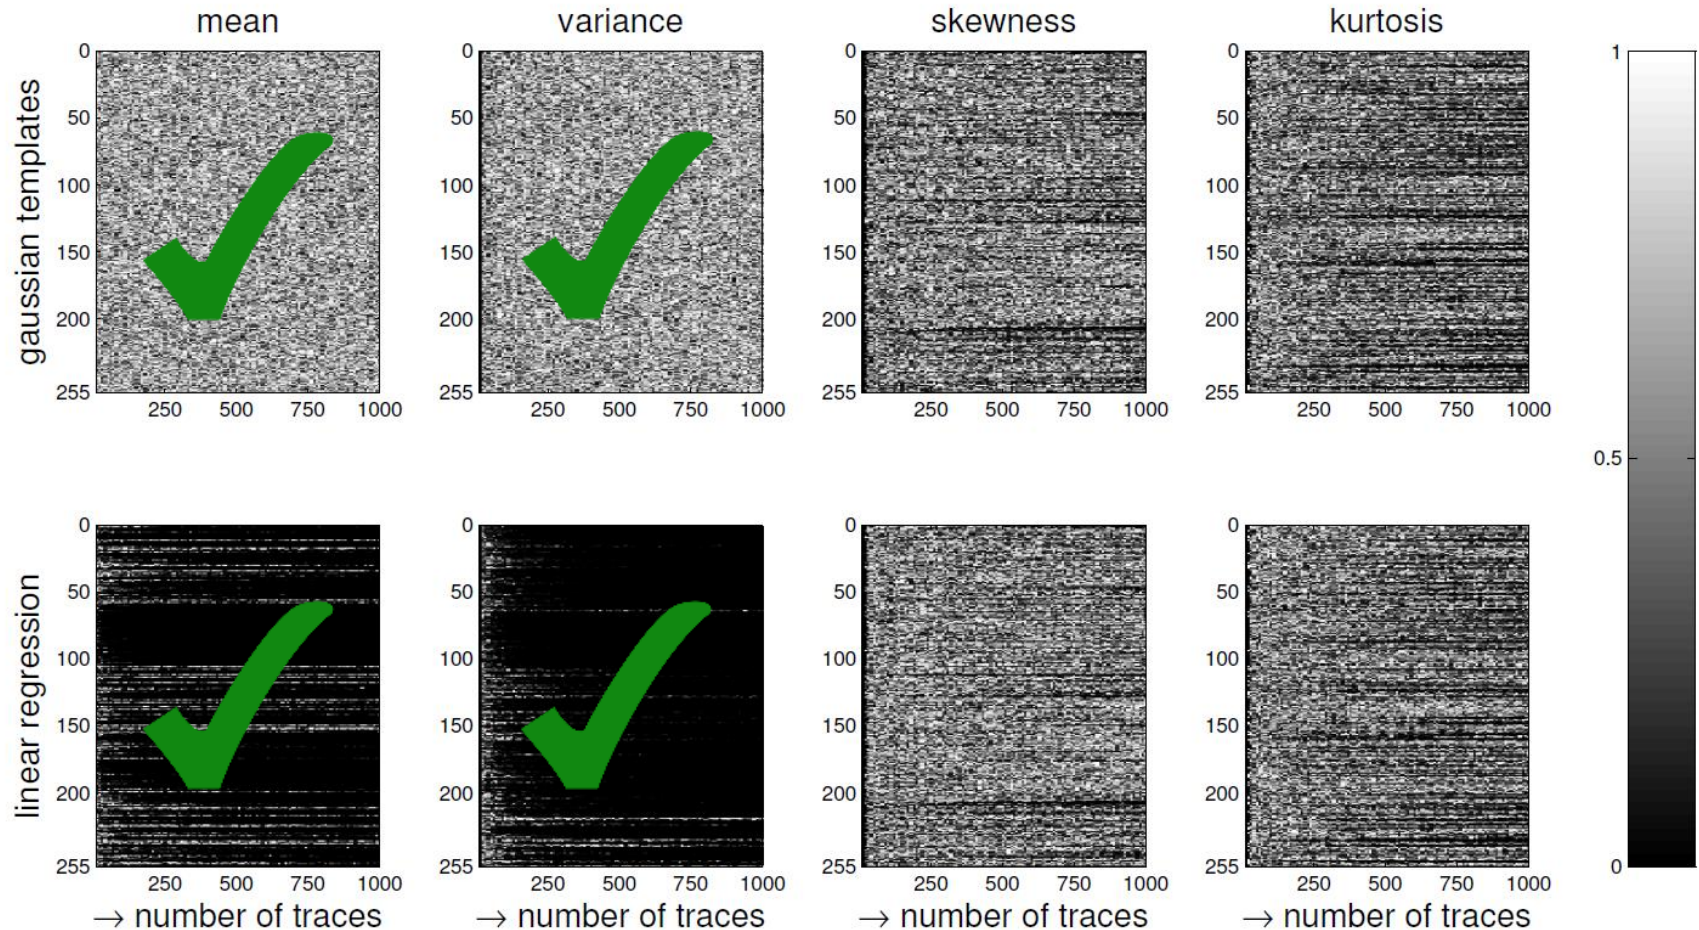
- Repeating the Eurocrypt 2014 case study

- Unprotected AES implementation, Atmel AVR

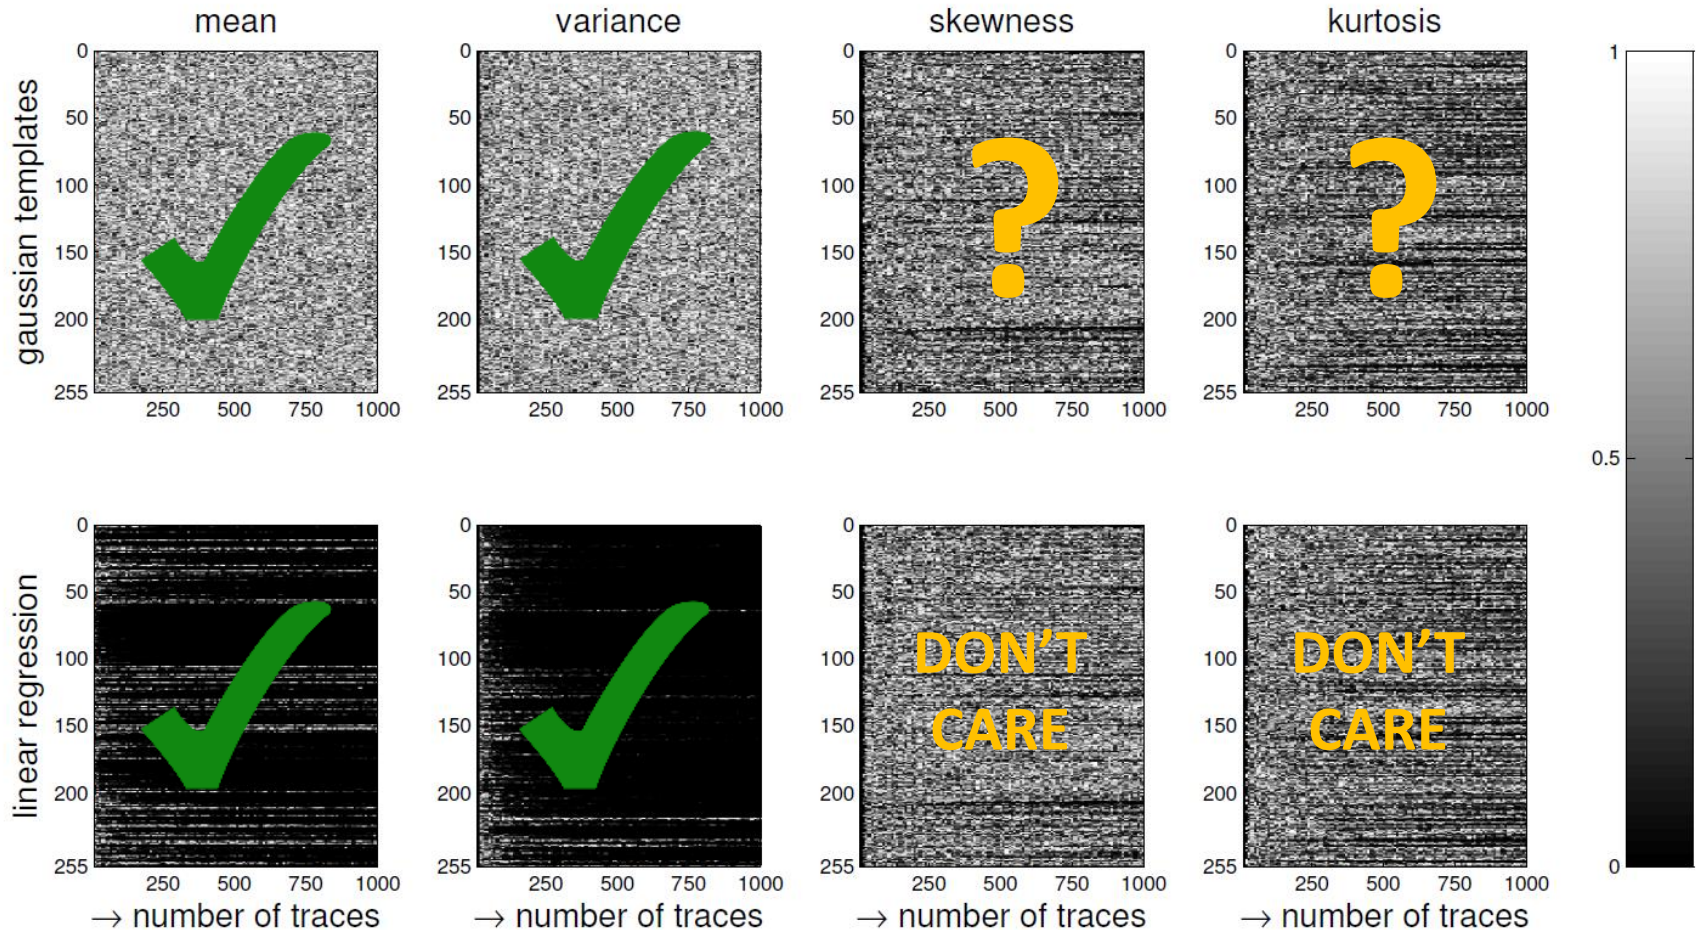
- Unprotected AES implementation, Atmel AVR



- Unprotected AES implementation, Atmel AVR



- Unprotected AES implementation, Atmel AVR



- Eurocrypt 2014: no errors detected with up to 256x1000 measurements & Gaussian template
- CHES 2016: small errors in \tilde{M}_3 and \tilde{M}_4

⇒ *Is there an inconsistency in our results?*

⇒ Do these errors lead to significant information loss

- Eurocrypt 2014: no errors detected with up to 256x1000 measurements & Gaussian template
- CHES 2016: small errors in \tilde{M}_3 and \tilde{M}_4

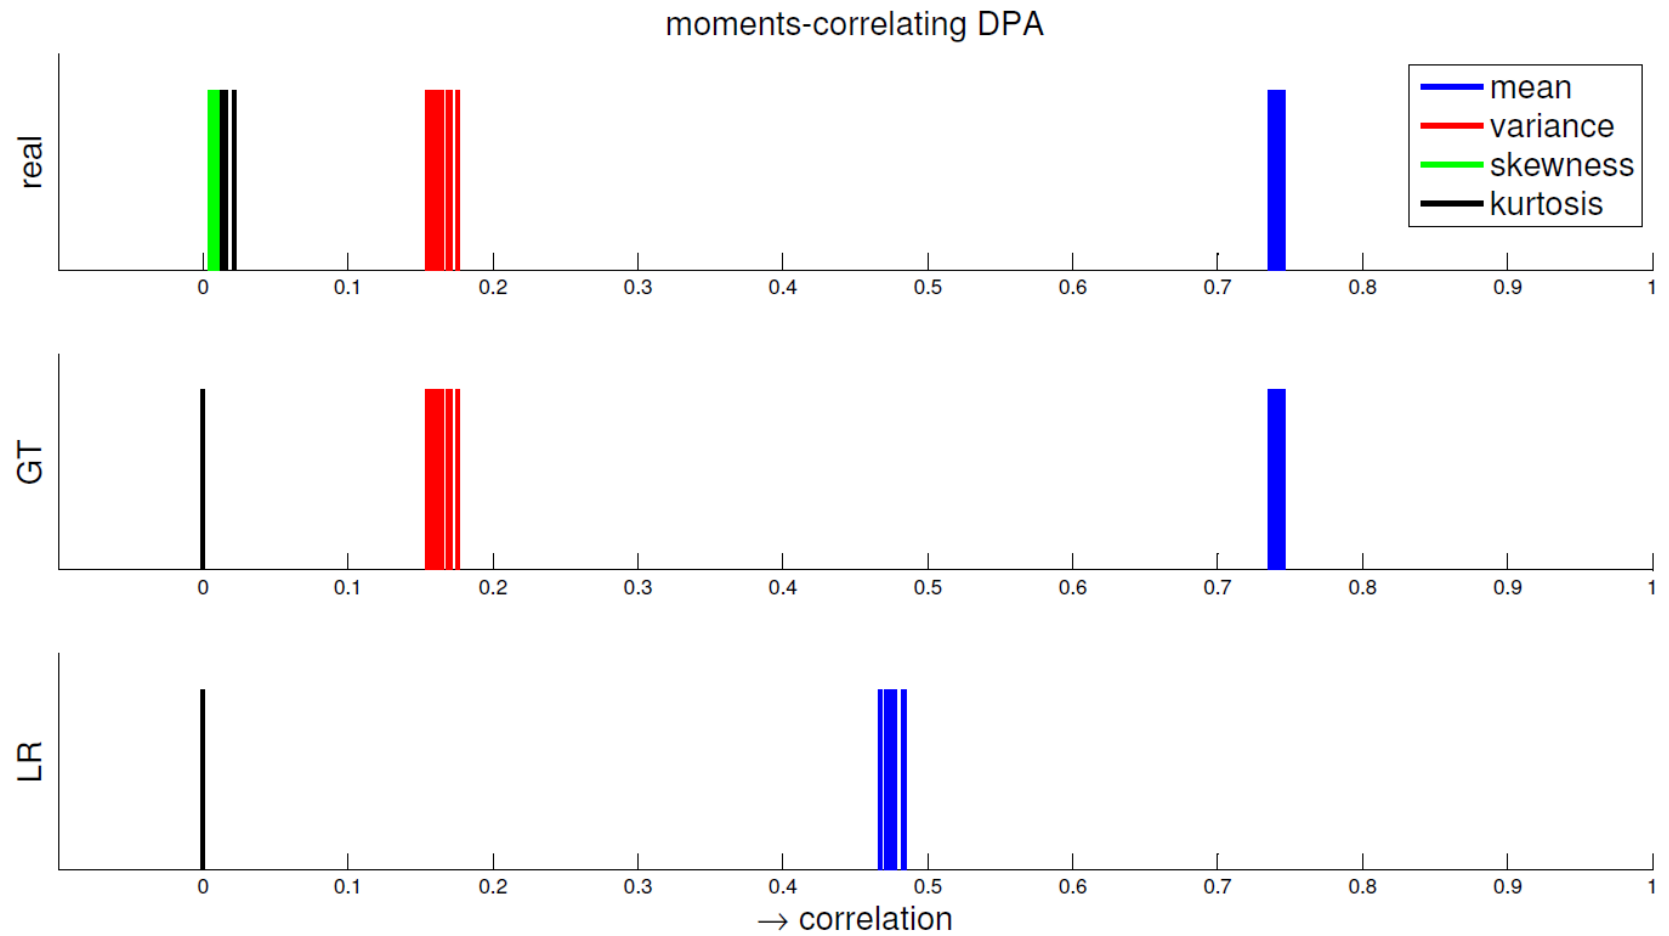
⇒ *Is there an inconsistency in our results?*

⇒ Do these errors lead to significant information loss

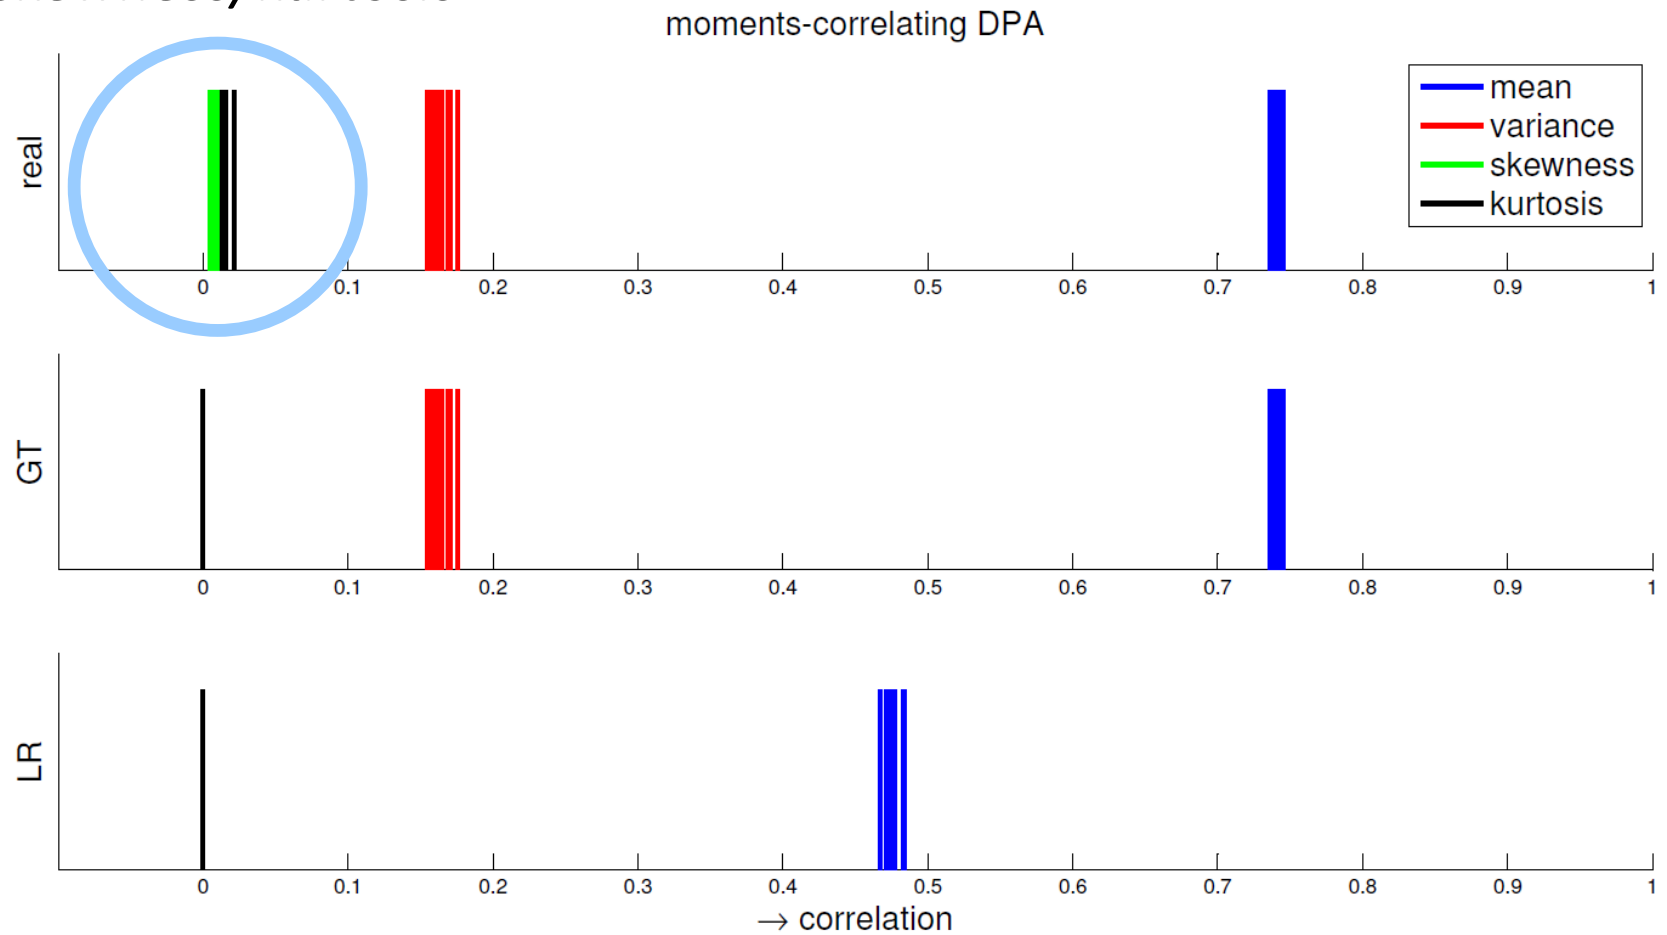
- Additional test: Moments-Correlating DPA [MS14]

$$\text{MPC-DPA}(d) = \hat{\rho}(\hat{M}_d, l^d)$$

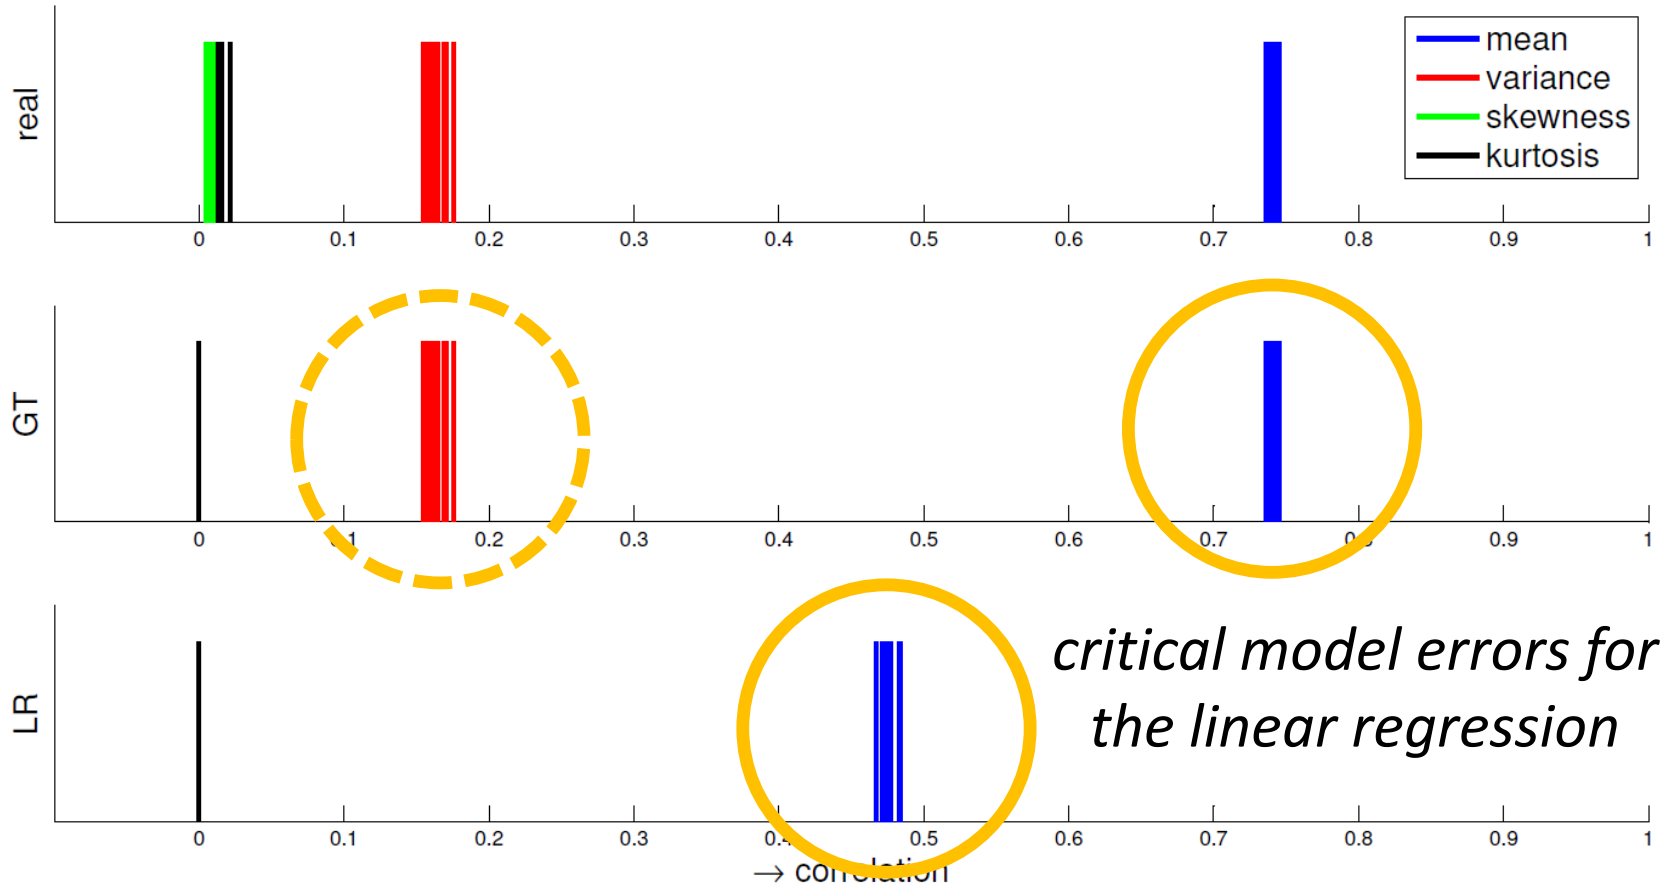
- Metric intuition: $N_s = \frac{c}{\hat{\rho}(\hat{M}_d, l^d)^2}$



*little information in
skewness/kurtosis*

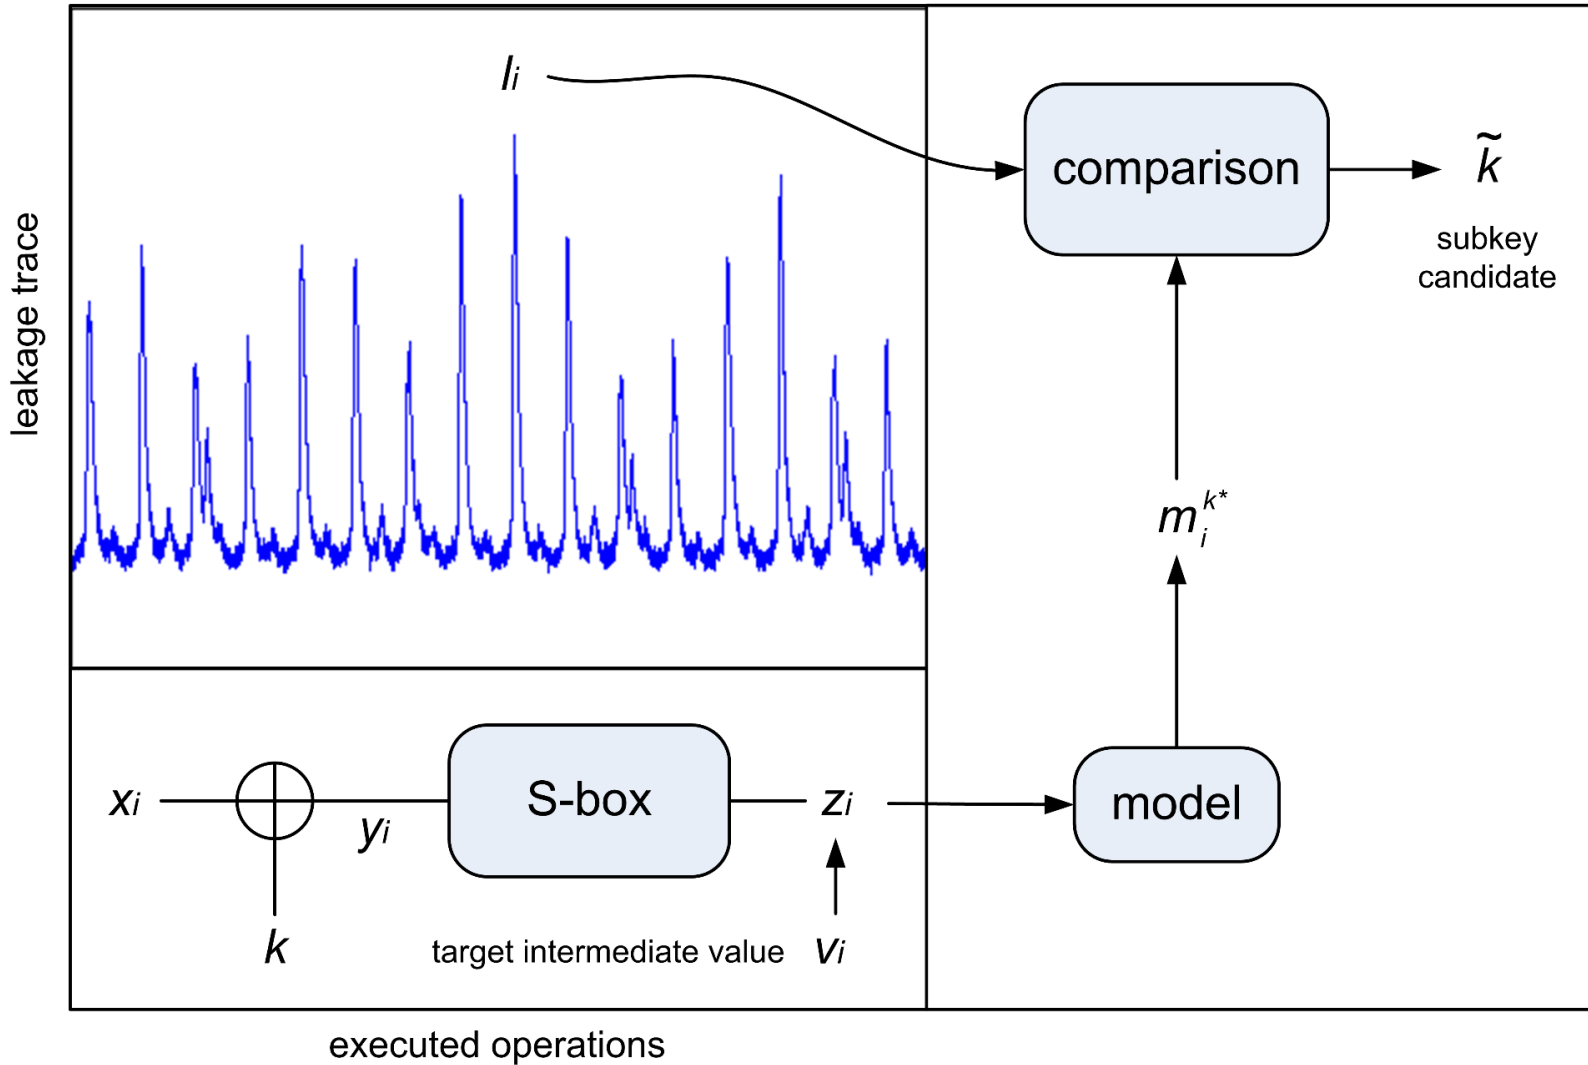


moments-correlating DPA



Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- **Exploitation**
 - **Soft Analytical Side-Channel Attacks**
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs



- More samples per intermediate value
 - Multivariate DPA
 - e.g., dimensionality reduction
 - Principal Component Analysis (PCA)
 - Linear Discriminant Analysis (LDA)
 - ...

- More samples per intermediate value
 - Multivariate DPA
 - e.g., dimensionality reduction
 - Principal Component Analysis (PCA)
 - Linear Discriminant Analysis (LDA)
 - ...
- **More intermediate values**

ARK	SB	SR	MC	ARK	ROUND 2	ROUND 3	ROUND 4	OTHER ROUNDS	FINAL ROUND
-----	----	----	----	-----	---------	---------	---------	--------------	-------------

- More samples per intermediate value
 - Multivariate DPA
 - e.g., dimensionality reduction
 - Principal Component Analysis (PCA)
 - Linear Discriminant Analysis (LDA)
 - ...
- More intermediate values (CRYPTO 1998)



STANDARD
DPA

- More samples per intermediate value
 - Multivariate DPA
 - e.g., dimensionality reduction
 - Principal Component Analysis (PCA)
 - Linear Discriminant Analysis (LDA)
 - ...
- More intermediate values (ASIACRYPT 2014)



MULTI-TARGET
DPA

- More samples per intermediate value
 - Multivariate DPA
 - e.g., dimensionality reduction
 - Principal Component Analysis (PCA)
 - Linear Discriminant Analysis (LDA)
 - ...
- More intermediate values (FSE 2003)



COLLISION ATTACKS

- More samples per intermediate value
 - Multivariate DPA
 - e.g., dimensionality reduction
 - Principal Component Analysis (PCA)
 - Linear Discriminant Analysis (LDA)
 - ...
- More intermediate values (CHES 2009)



ALGEBRAIC SIDE-CHANNEL ATTACKS (ASCA)

- Cannot deal with measurement noise
 - Despite progresses, e.g., Tolerant ASCA
- Large (time and) memory complexities
 - Limited to a single plaintext, typically
 - Sometimes even less

- Cannot deal with measurement noise
 - Despite progresses, e.g., Tolerant ASCA
- Large (time and) memory complexities
 - Limited to a single plaintext, typically
 - Sometimes even less

⇒ Emerging intuition: ASCA require “hard” information and only standard DPA can efficiently exploit “soft” (probabilistic) information obtained from the measurements of multiple plaintexts

- Cannot deal with measurement noise
 - Despite progresses, e.g., Tolerant ASCA
 - Large (time and) memory complexities
 - Limited to a single plaintext, typically
 - Sometimes even less
- ⇒ Emerging intuition: ASCA require “hard” information and only standard DPA can efficiently exploit “soft” (probabilistic) information obtained from the measurements of multiple plaintexts
- ***Our contribution:*** show this intuition is incorrect!

- Representation of the algorithm/implementation

- Representation of the algorithm/implementation
- Based on two types of nodes

- Representation of the algorithm/implementation
- Based on two types of nodes
 - Variable nodes x_i (i.e., intermediate values)

- Representation of the algorithm/implementation
- Based on two types of nodes
 - Variable nodes x_i (i.e., intermediate values)
 - Factor nodes of two types
 - A priori knowledge $f_i(x_i)=\Pr[x_i|L]$
 - ***Exactly the output of standard DPA!***
 - Operations: $f(x_1, x_2, x_3) = \begin{matrix} 1 & \text{if } OP(x_1, x_2)=x_3 \\ 0 & \text{otherwise} \end{matrix}$

- Representation of the algorithm/implementation
- Based on two types of nodes
 - Variable nodes x_i (i.e., intermediate values)
 - Factor nodes of two types
 - A priori knowledge $f_i(x_i) = \Pr[x_i|L]$
 - ***Exactly the output of standard DPA!***
 - Operations: $f(x_1, x_2, x_3) = \begin{cases} 1 & \text{if } OP(x_1, x_2) = x_3 \\ 0 & \text{otherwise} \end{cases}$
- Edges carry two types of messages
 - Type q messages: from variables to factors
 - Type r messages: from factors to variables

- Propagates the information (probabilities) through the factor graph via message passing

- Propagates the information (probabilities) through the factor graph via message passing

- From variables to factors

$$q_{v_n \rightarrow f_m}(x_n) = \prod_{m' \in M \setminus m} r_{f_{m'} \rightarrow v_n}(x_n)$$

- \approx product over all incoming messages excluding the one of the target factor node

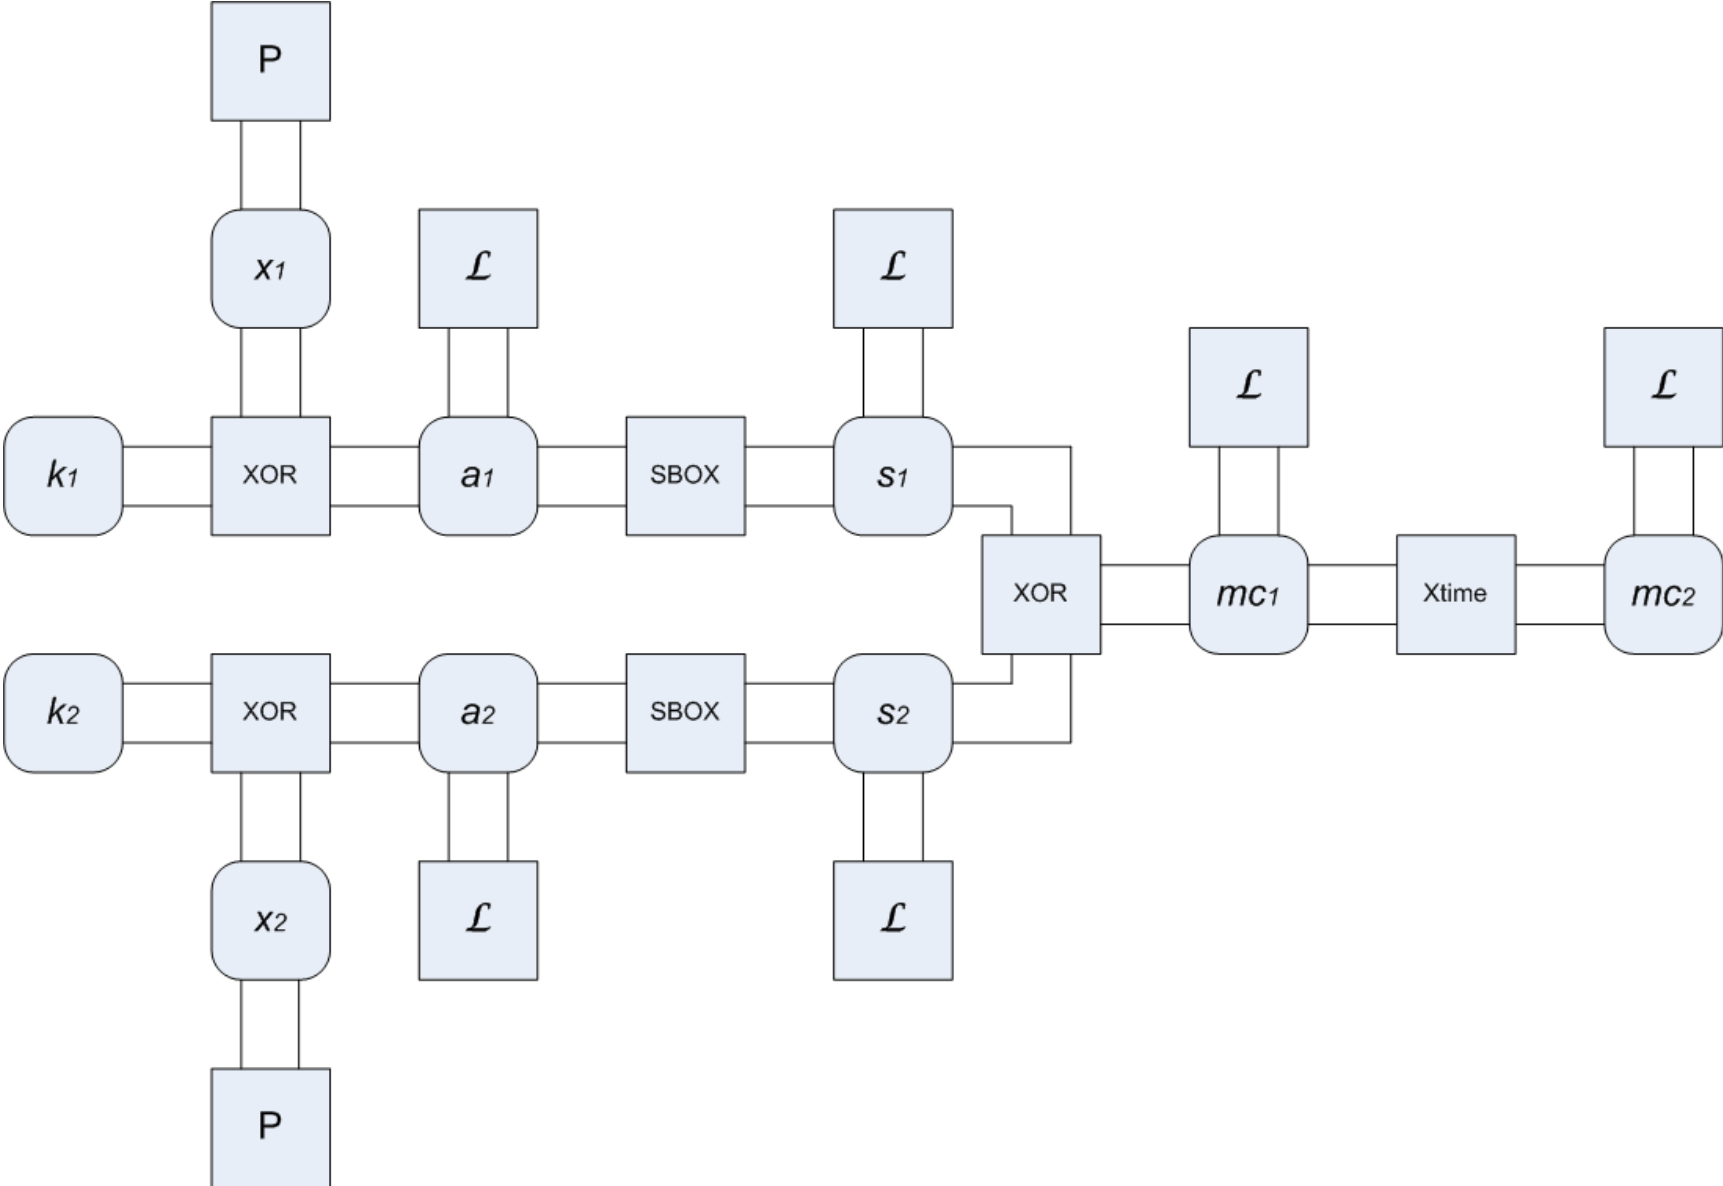
- Propagates the information (probabilities) through the factor graph via message passing

- From factors to variables

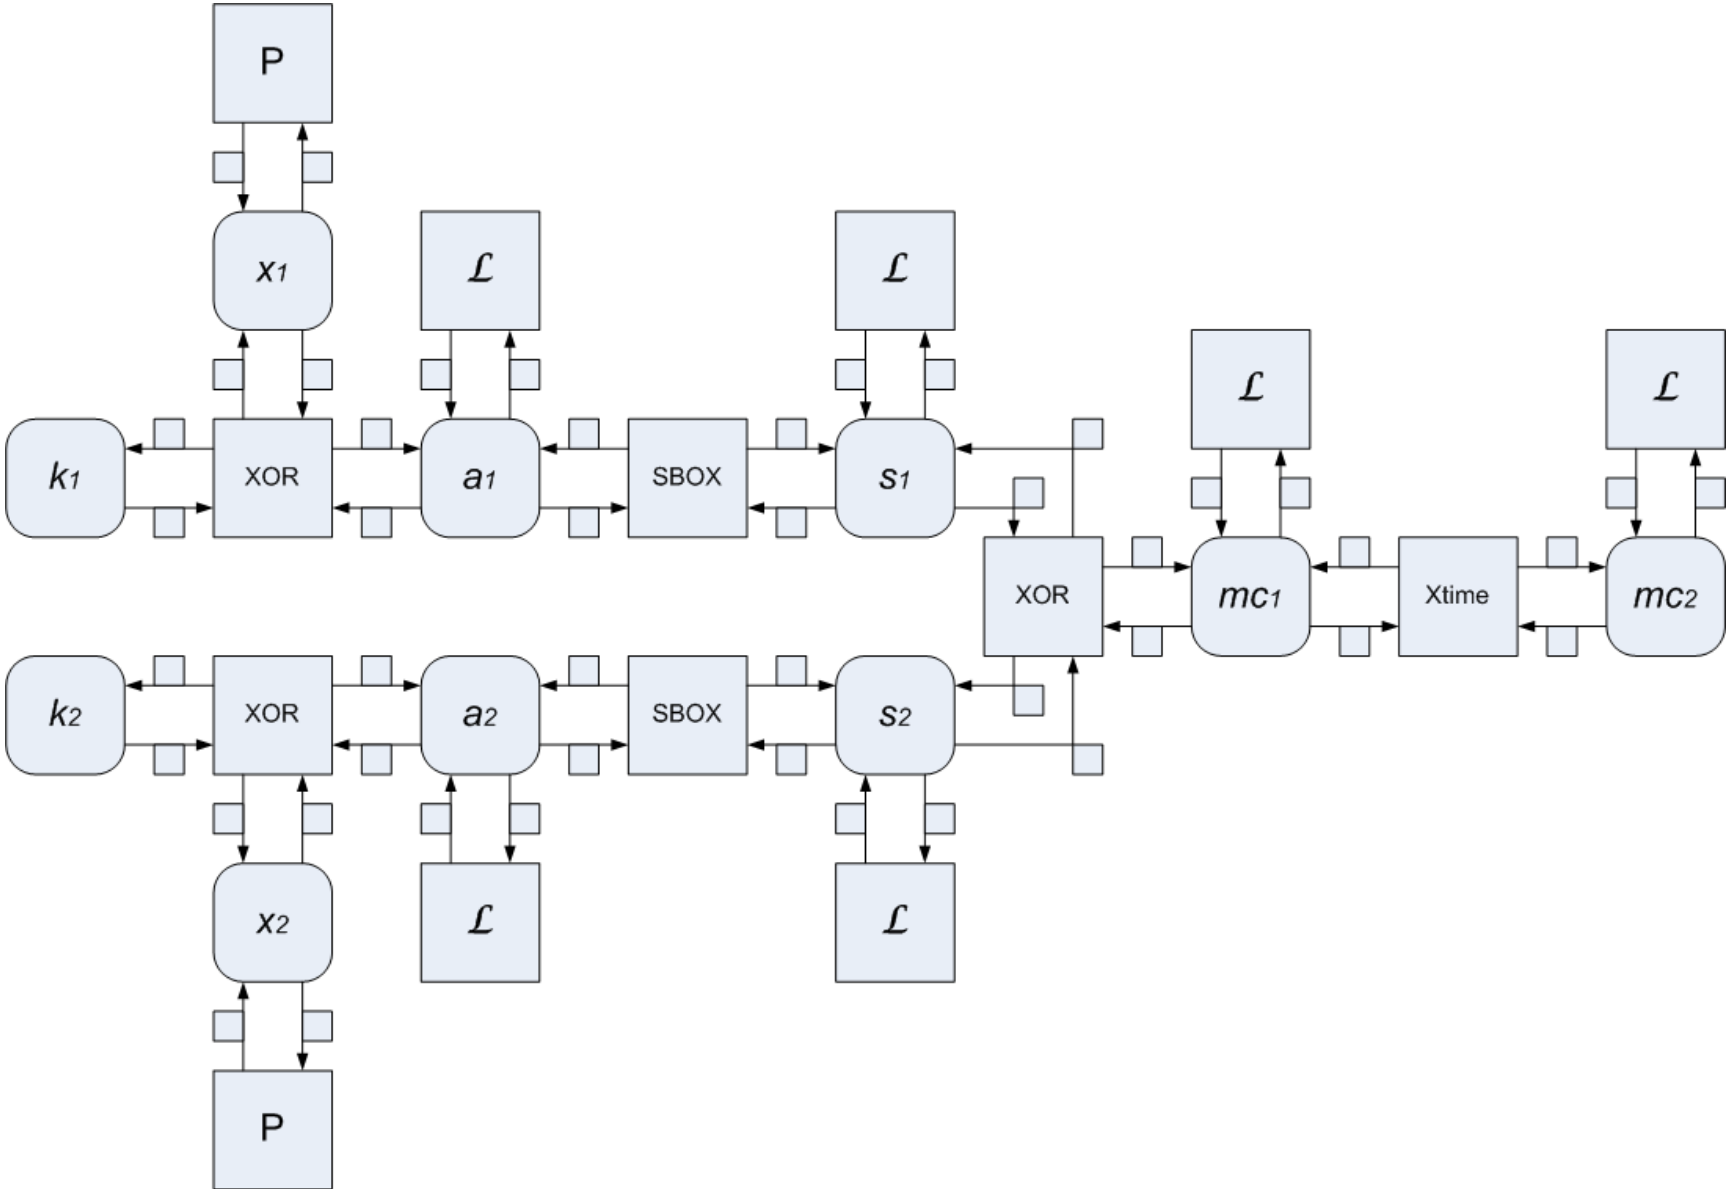
$$r_{f_m \rightarrow v_n}(x_n) = \sum_{x_m \neq x_n} f_m(x_m) \cdot \prod_{n' \in N \setminus n} q_{v_{n'} \rightarrow f_m}(x_{n'})$$

- \approx weighted sum of products over all incoming messages excluding the target variable node

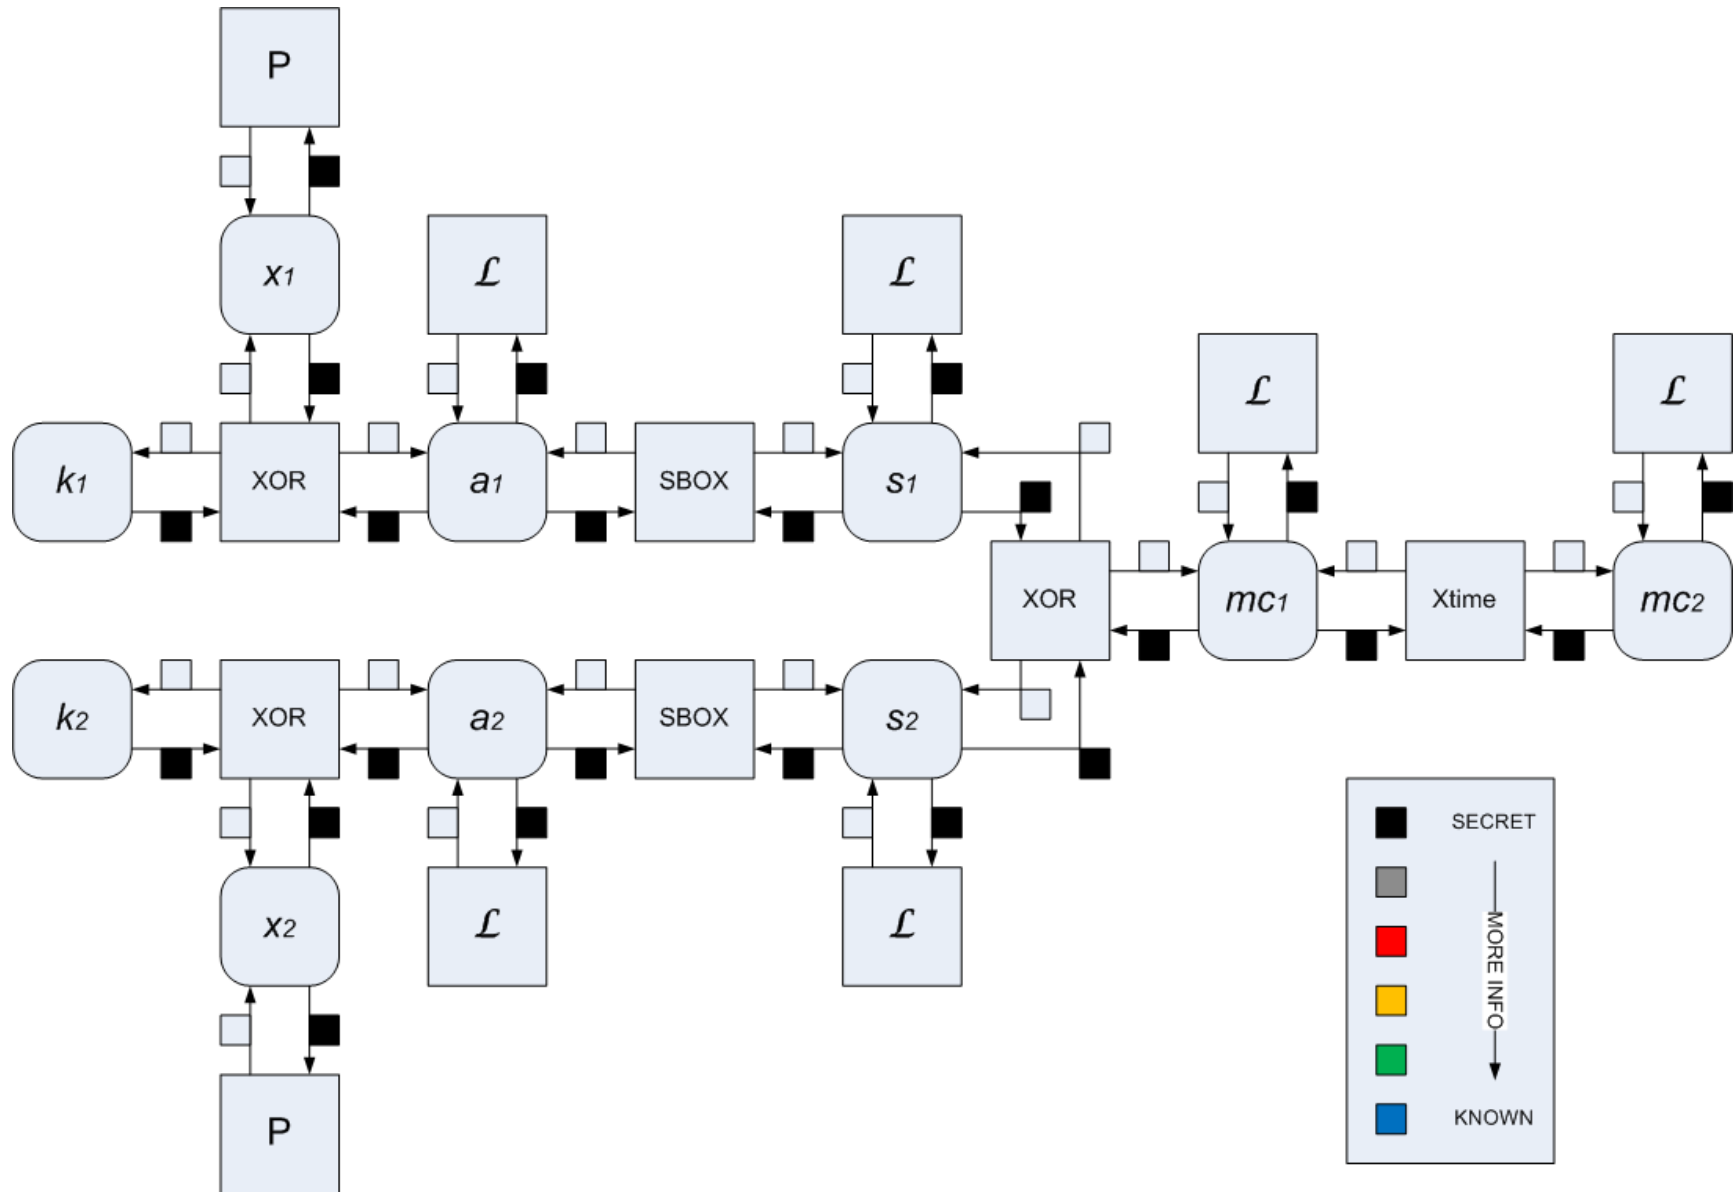
Example (I): factor graph



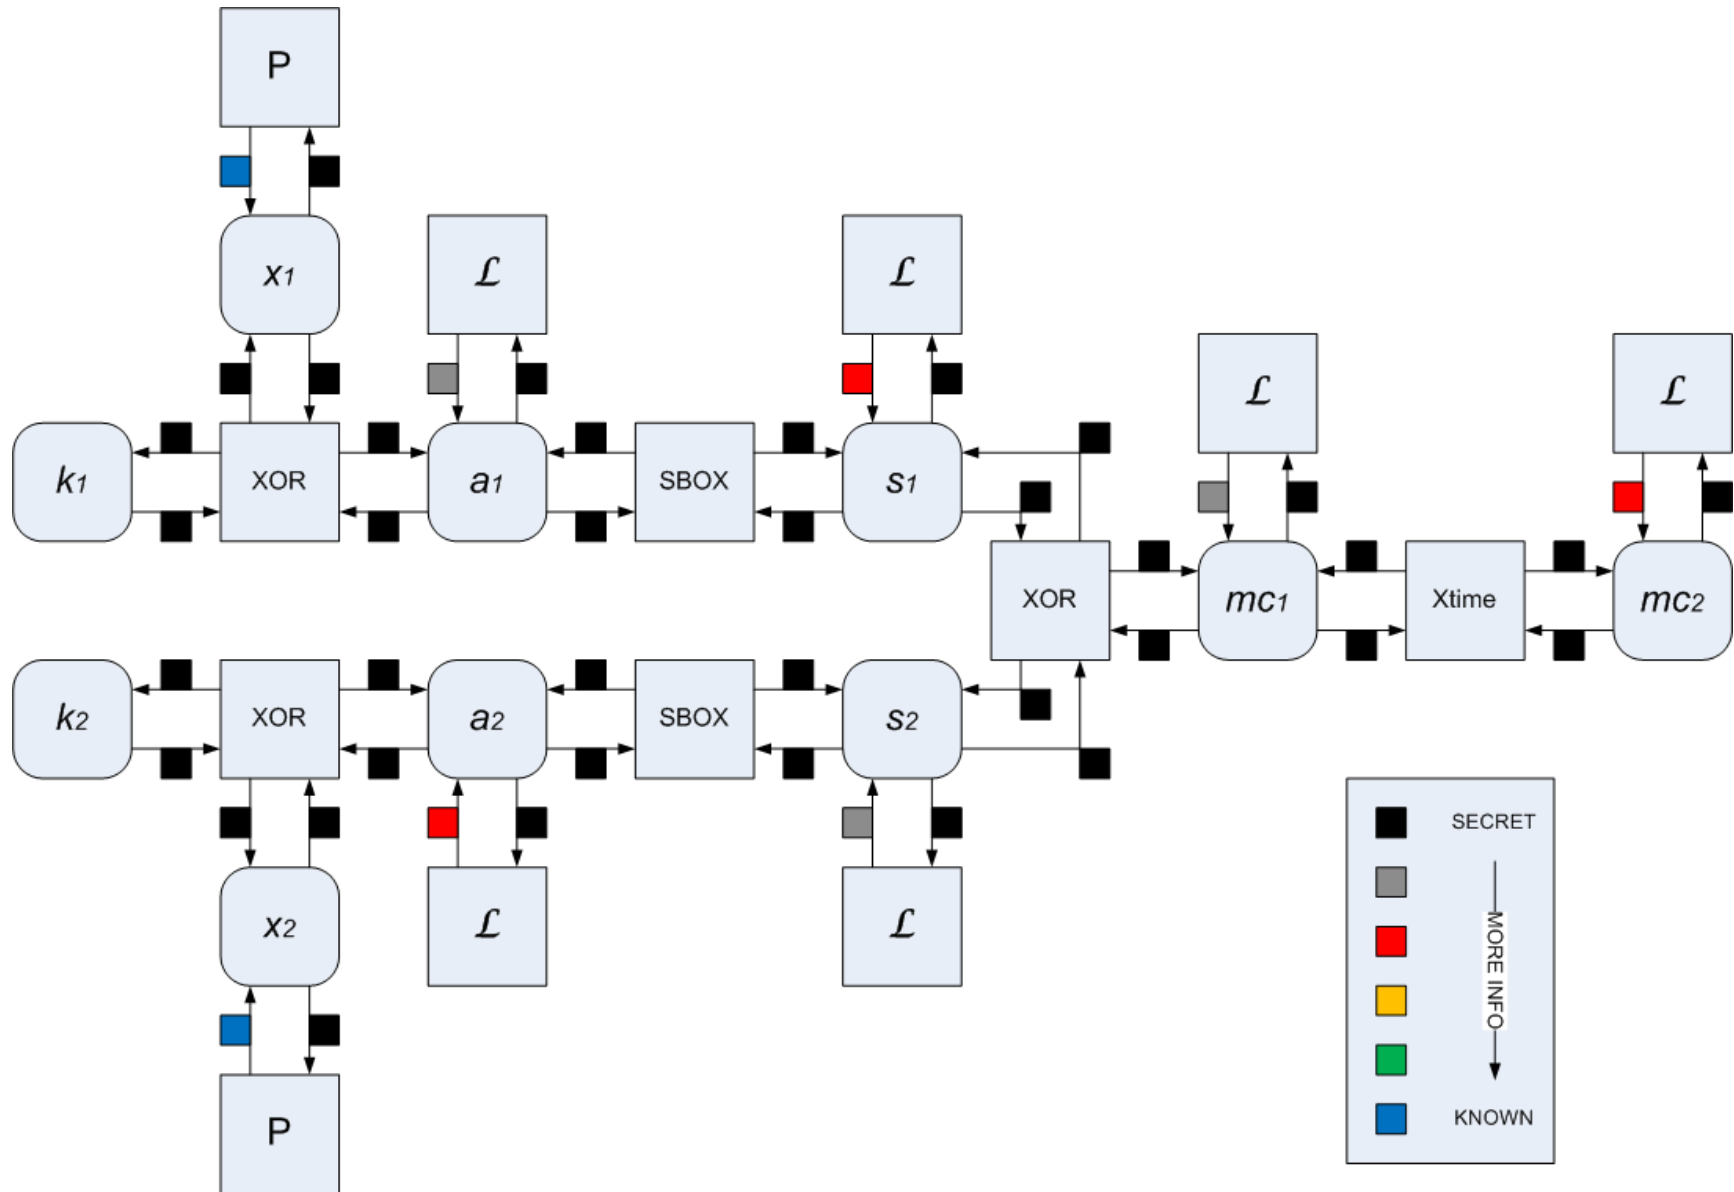
Example (II): adding the messages



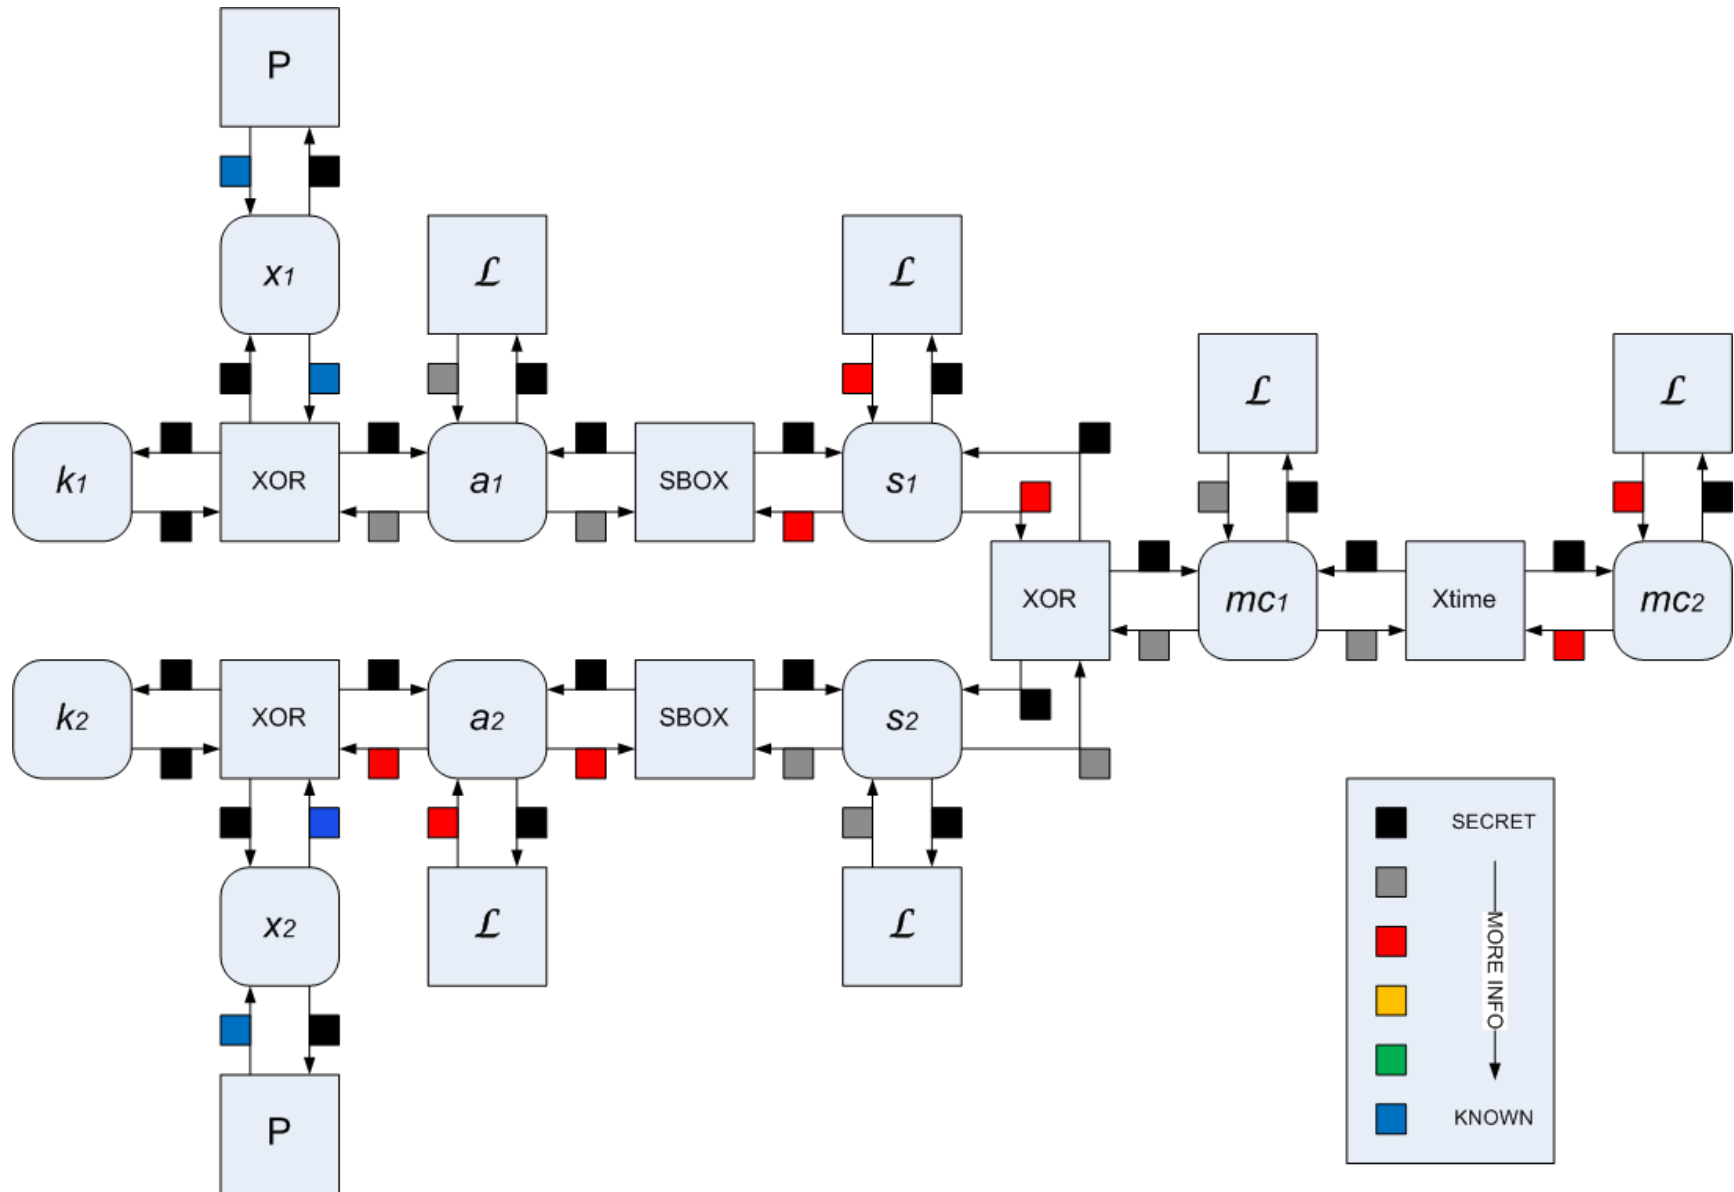
Example (III): initialize the q 's (v to f)



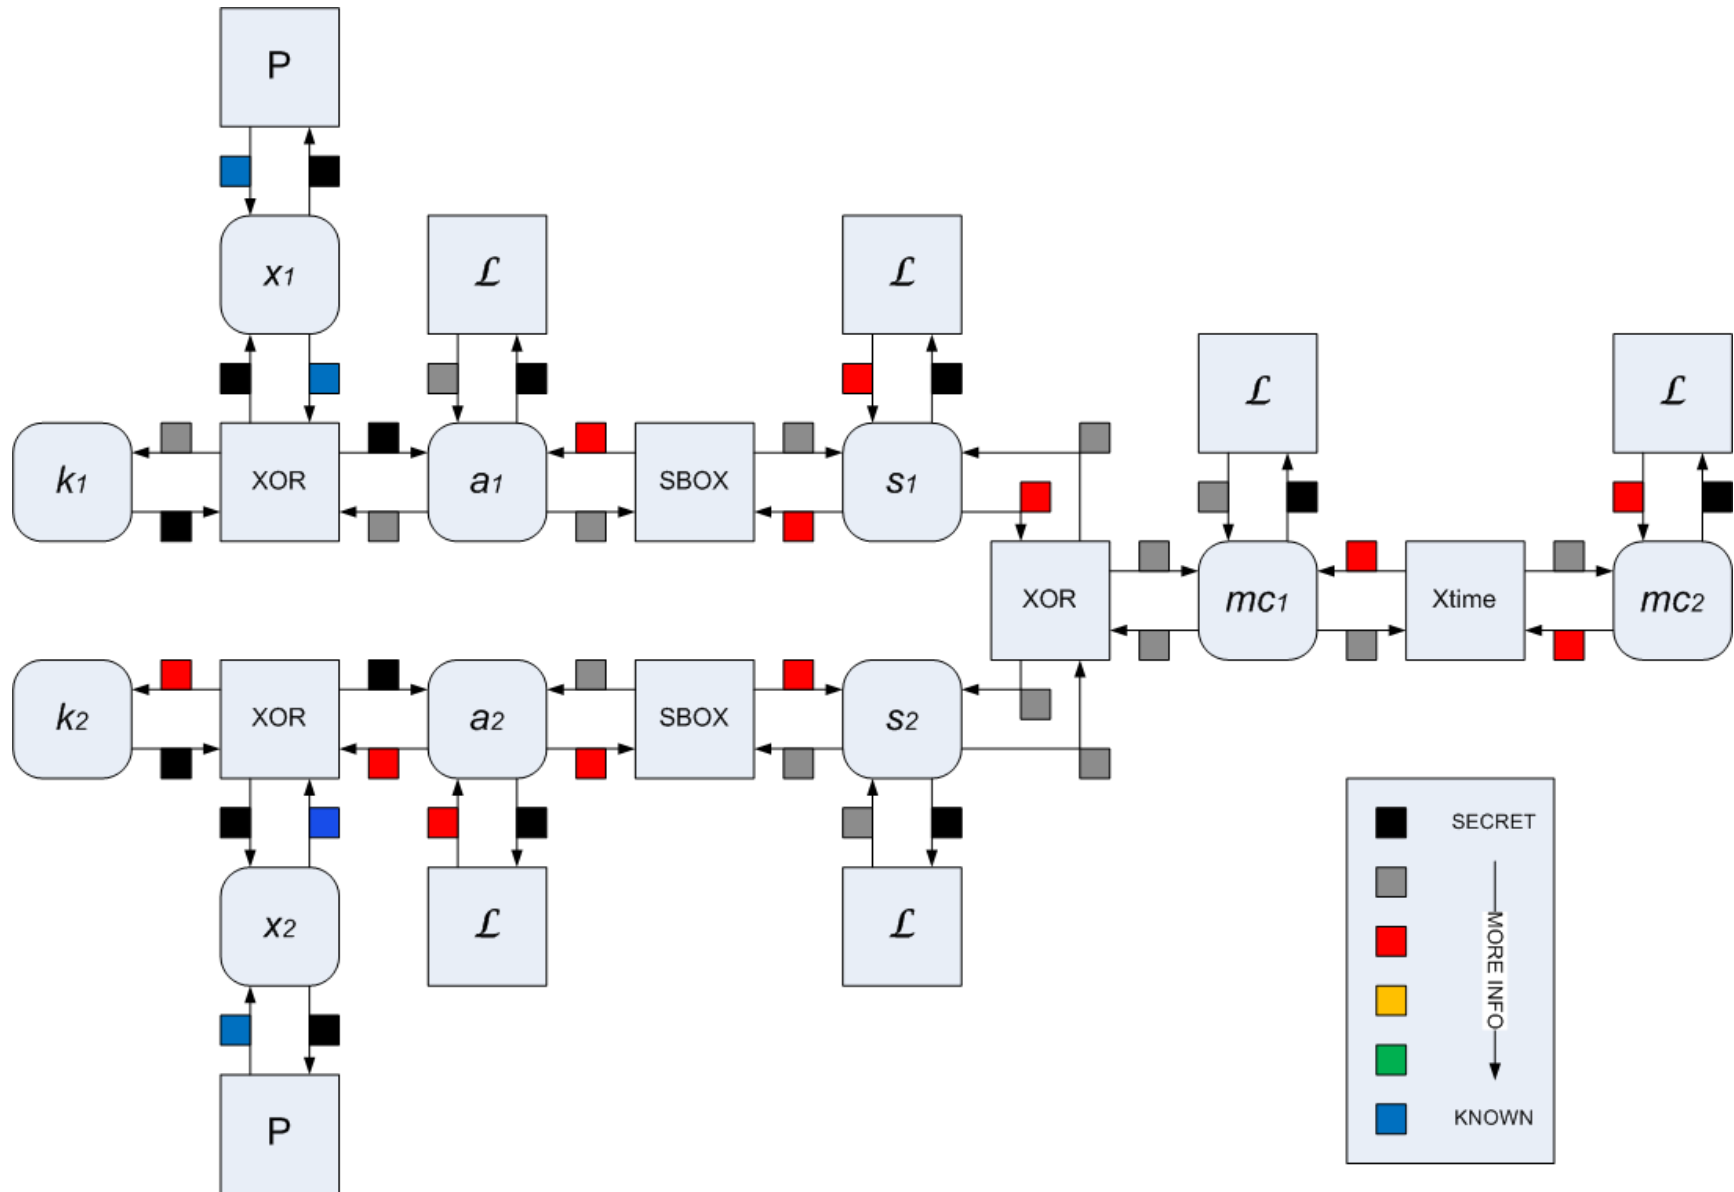
Example (IV): initialize the r 's (f to v)



Example (V): update the q 's (v to f)

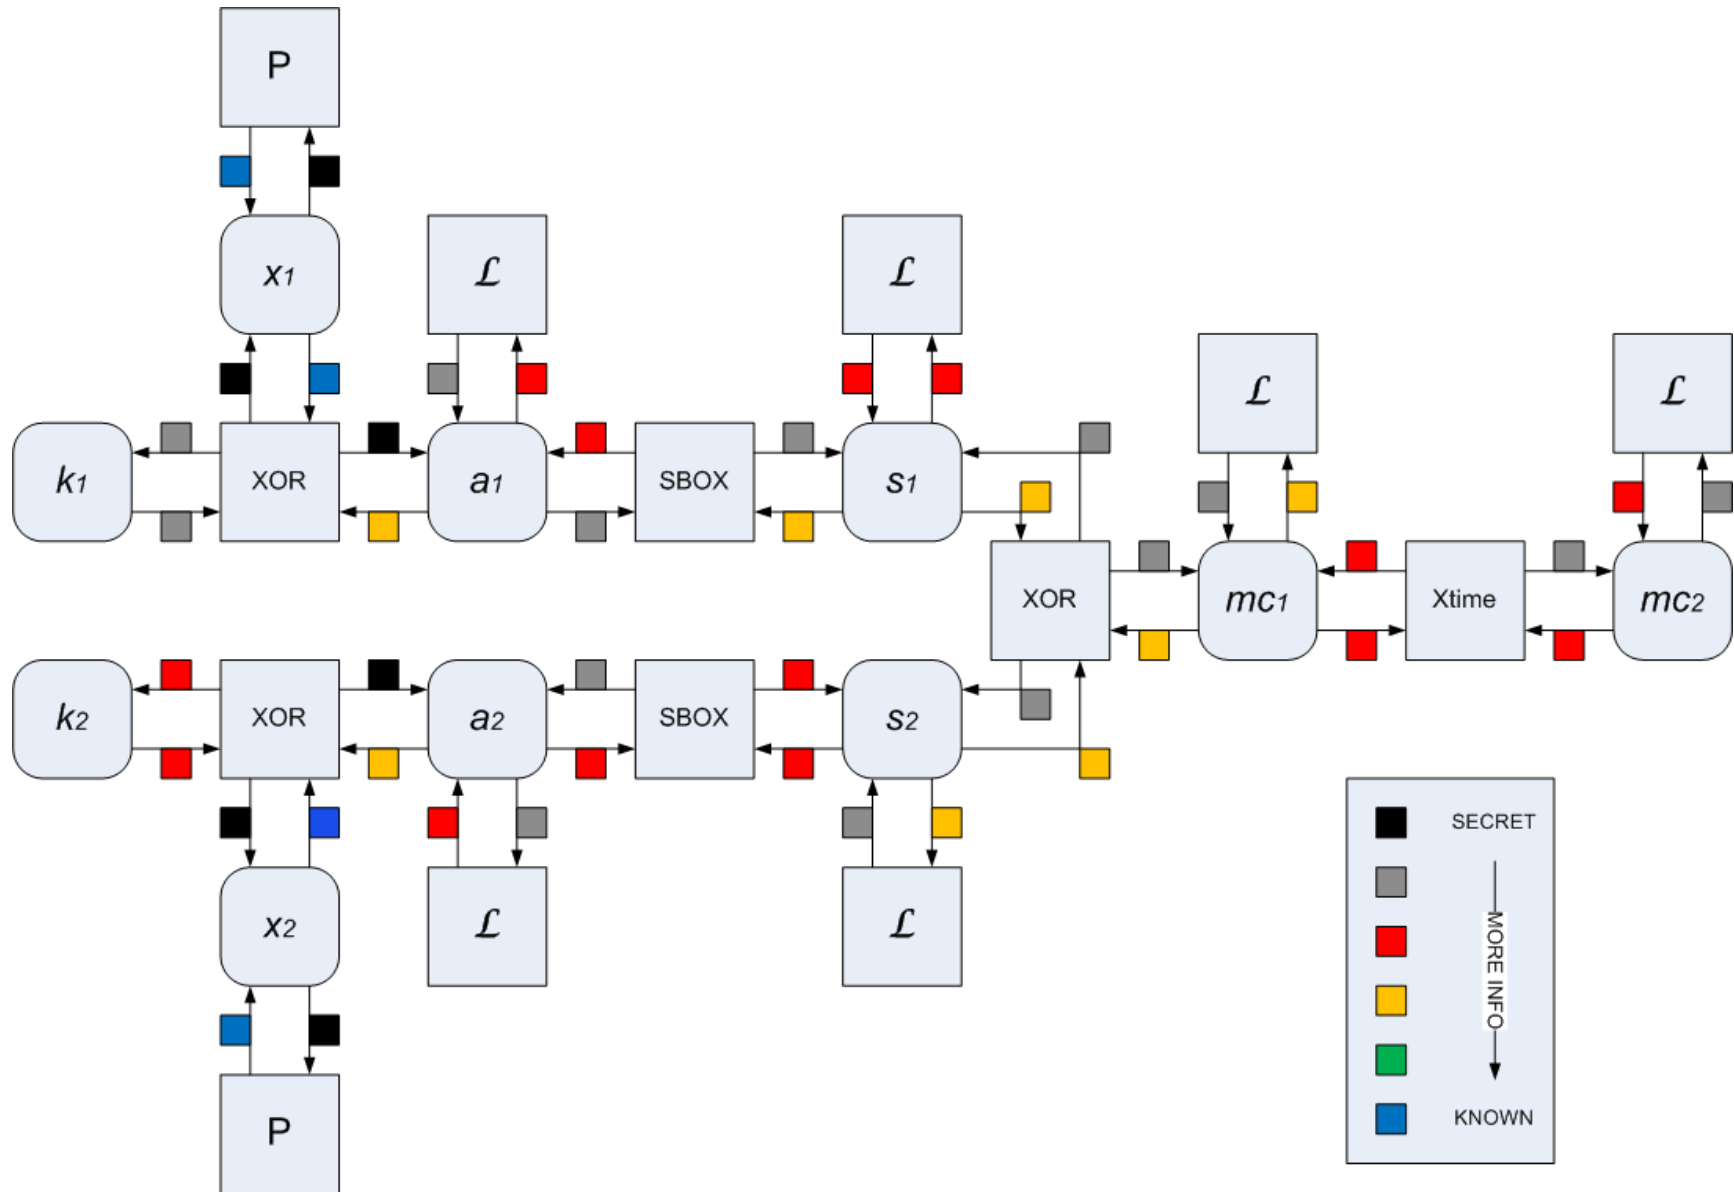


Example (VI): update the r 's (f to v)



Example (VII): update the q 's (v to f)

24



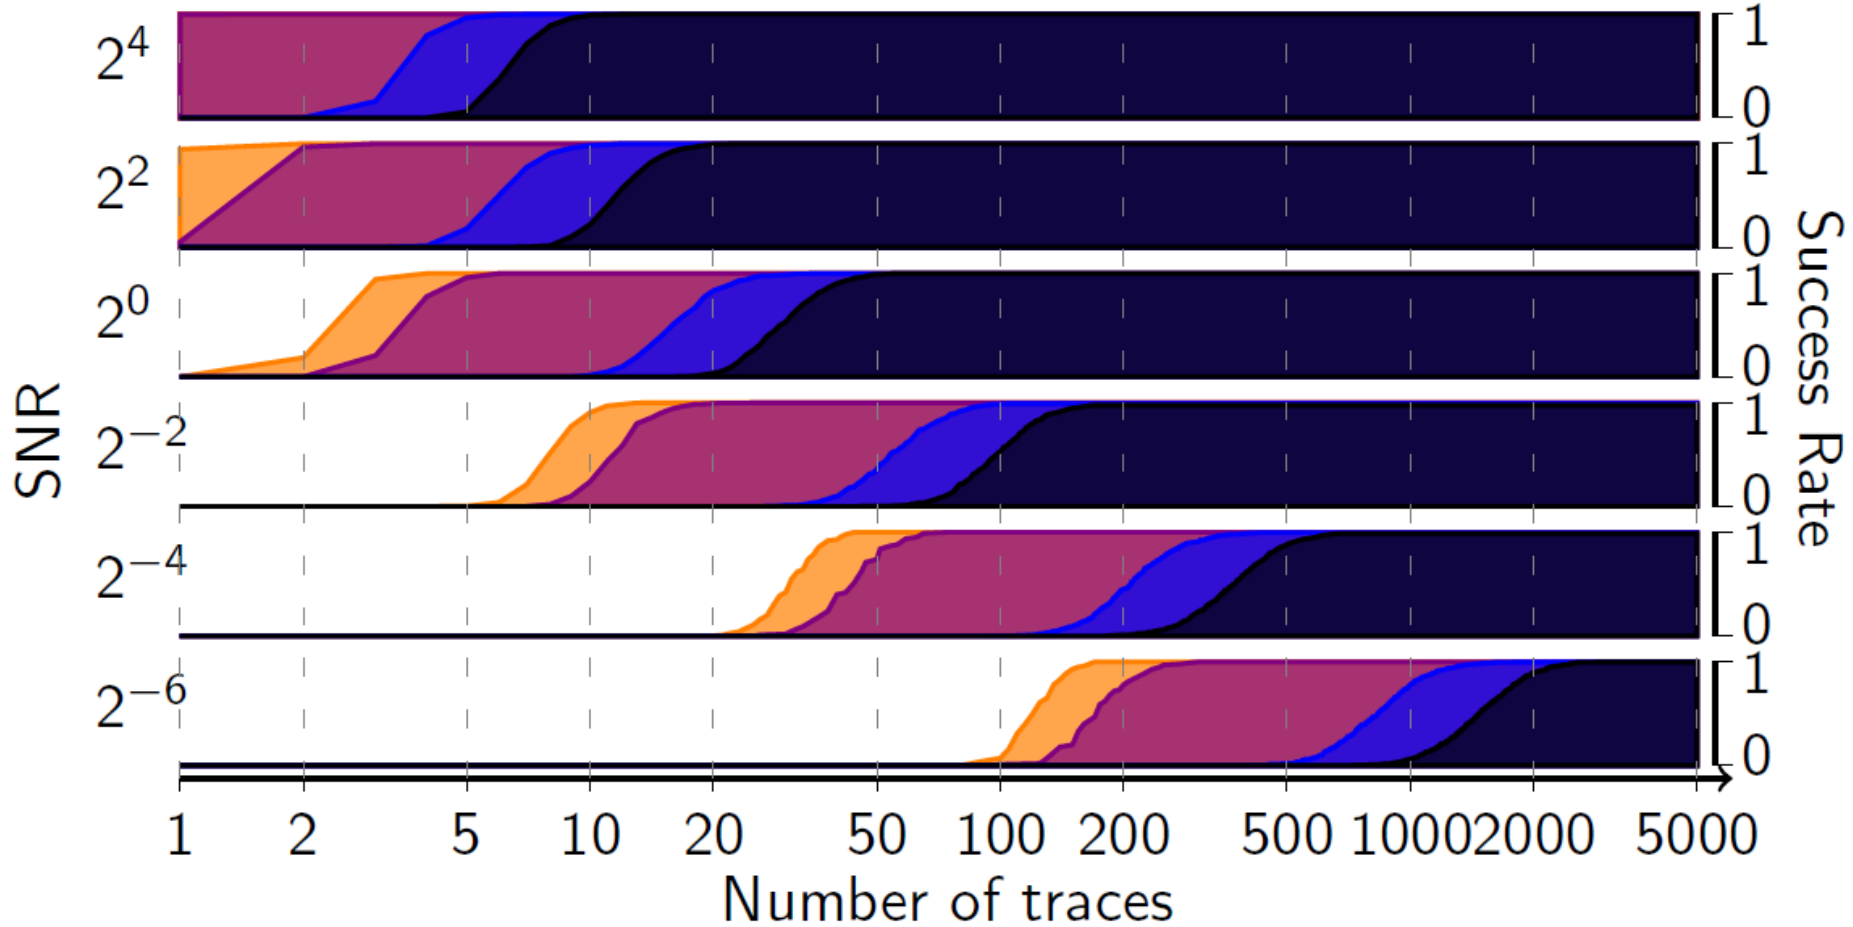
- Good news: any knowledge can be exploited
- e.g. (open source) AES furious assembly code

ASM code	Graph description	Factor graph
<pre>ld H1, Y+ eor ST11, H1 mov ZL, ST11 lpm ST11, Z</pre>	<pre>* _Xor AK[1,1]_0 ST[1,1]_0 K[1,1]_0 * _Sbox SB[1,1]_0 AK[1,1]_0</pre>	<p>A factor graph with four nodes: $K_{1,1}^0$, $ST_{1,1}^0$, $AK_{1,1}^0$, and $SB_{1,1}^0$. Edges connect $K_{1,1}^0$ to $AK_{1,1}^0$, $ST_{1,1}^0$ to $AK_{1,1}^0$, $AK_{1,1}^0$ to $SB_{1,1}^0$, and $SB_{1,1}^0$ to $AK_{1,1}^0$. Two operations are shown: XOR, which is connected to $K_{1,1}^0$ and $ST_{1,1}^0$, and SBOX, which is connected to $AK_{1,1}^0$ and $SB_{1,1}^0$.</p>
<pre>mov H3, ST11 eor H3, ST21 mov ZL, H3 lpm H3, Z</pre>	<pre>* _Xor MC[3,1]_0 SB[1,1]_0 SB[2,1]_0 * _Xtime XT[1,1]_0 MC[3,1]_0</pre>	<p>A factor graph with four nodes: $SB_{1,1}^0$, $SB_{2,1}^0$, $MC_{3,1}^0$, and $XT_{1,1}^0$. Edges connect $SB_{1,1}^0$ to $MC_{3,1}^0$, $SB_{2,1}^0$ to $MC_{3,1}^0$, $MC_{3,1}^0$ to $XT_{1,1}^0$, and $XT_{1,1}^0$ to $MC_{3,1}^0$. Two operations are shown: XOR, which is connected to $SB_{1,1}^0$ and $SB_{2,1}^0$, and XTIME, which is connected to $MC_{3,1}^0$ and $XT_{1,1}^0$.</p>
<pre>mov ZL, ST24 lpm H3, Z eor ST11, H3 eor ST11, H1</pre>	<pre>* _Sbox SK[1,1]_1 K[2,4]_0 _Xor XK[1,1]_1 SK[1,1]_1 K[1,1]_0 _XorCst K[1,1]_1 XK[1,1]_1 0x1</pre>	<p>A factor graph with five nodes: $K_{2,4}^0$, $SK_{1,1}^1$, $K_{1,1}^0$, $XK_{1,1}^1$, and $K_{1,1}^1$. Edges connect $K_{2,4}^0$ to $XK_{1,1}^1$, $SK_{1,1}^1$ to $XK_{1,1}^1$, $K_{1,1}^0$ to $XK_{1,1}^1$, $XK_{1,1}^1$ to $K_{1,1}^1$, and $K_{1,1}^1$ to $XK_{1,1}^1$. Three operations are shown: SBOX, connected to $K_{2,4}^0$ and $XK_{1,1}^1$; XOR, connected to $SK_{1,1}^1$ and $K_{1,1}^0$; and XORCST, connected to $XK_{1,1}^1$ and $K_{1,1}^1$.</p>

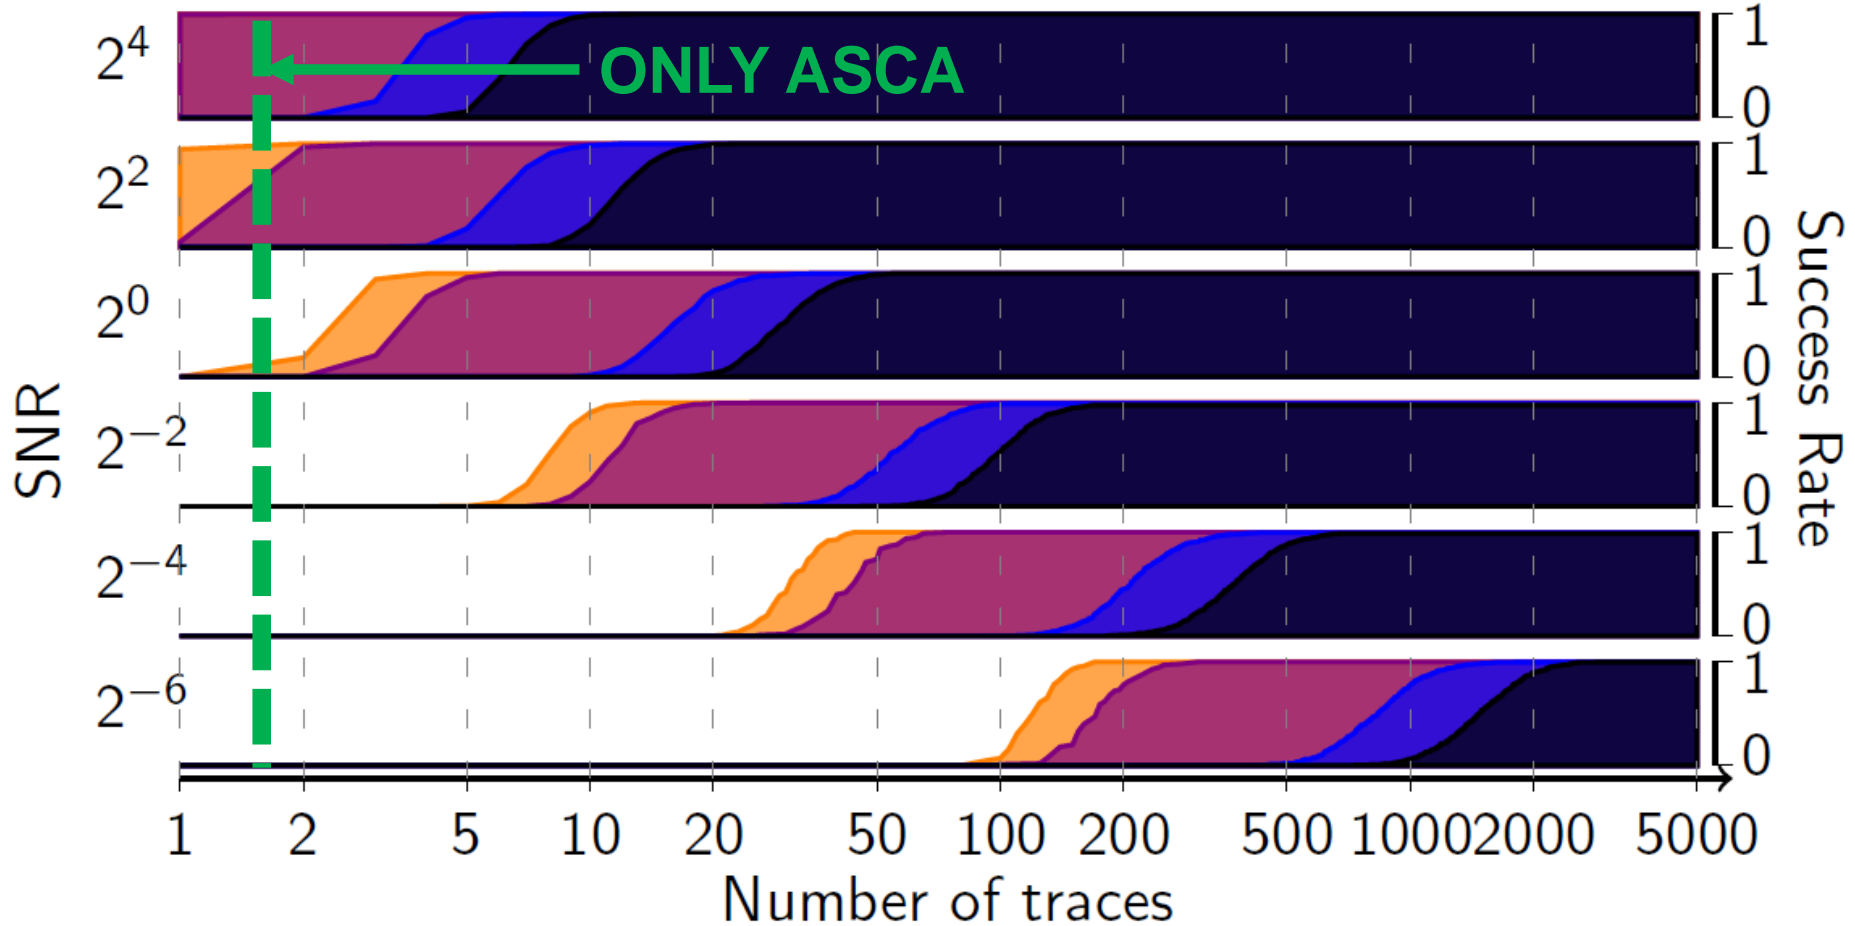
- Good news: any knowledge can be exploited
- e.g. (open source) AES furious assembly code

ASM code	Graph description	Factor graph
<pre>ld H1, Y+ eor ST11, H1 mov ZL, ST11 lpm ST11, Z</pre>	<pre>* _Xor AK[1,1]_0 ST[1,1]_0 K[1,1]_0 * _Sbox SB[1,1]_0 AK[1,1]_0</pre>	
<pre>mov H3, ST11 eor H3, ST21 mov ZL, H3 lpm H3, Z</pre>	<pre>* _Xor MC[3,1]_0 SB[1,1]_0 SB[2,1]_0 * _Xtime XT[1,1]_0 MC[3,1]_0</pre>	
<pre>mov ZL, ST24 lpm H3, Z eor ST11, H3 eor ST11, H1</pre>	<pre>* _Sbox SK[1,1]_1 K[2,4]_0 _Xor XK[1,1]_1 SK[1,1]_1 K[1,1]_0 _XorCst K[1,1]_1 XK[1,1]_1 0x1</pre>	

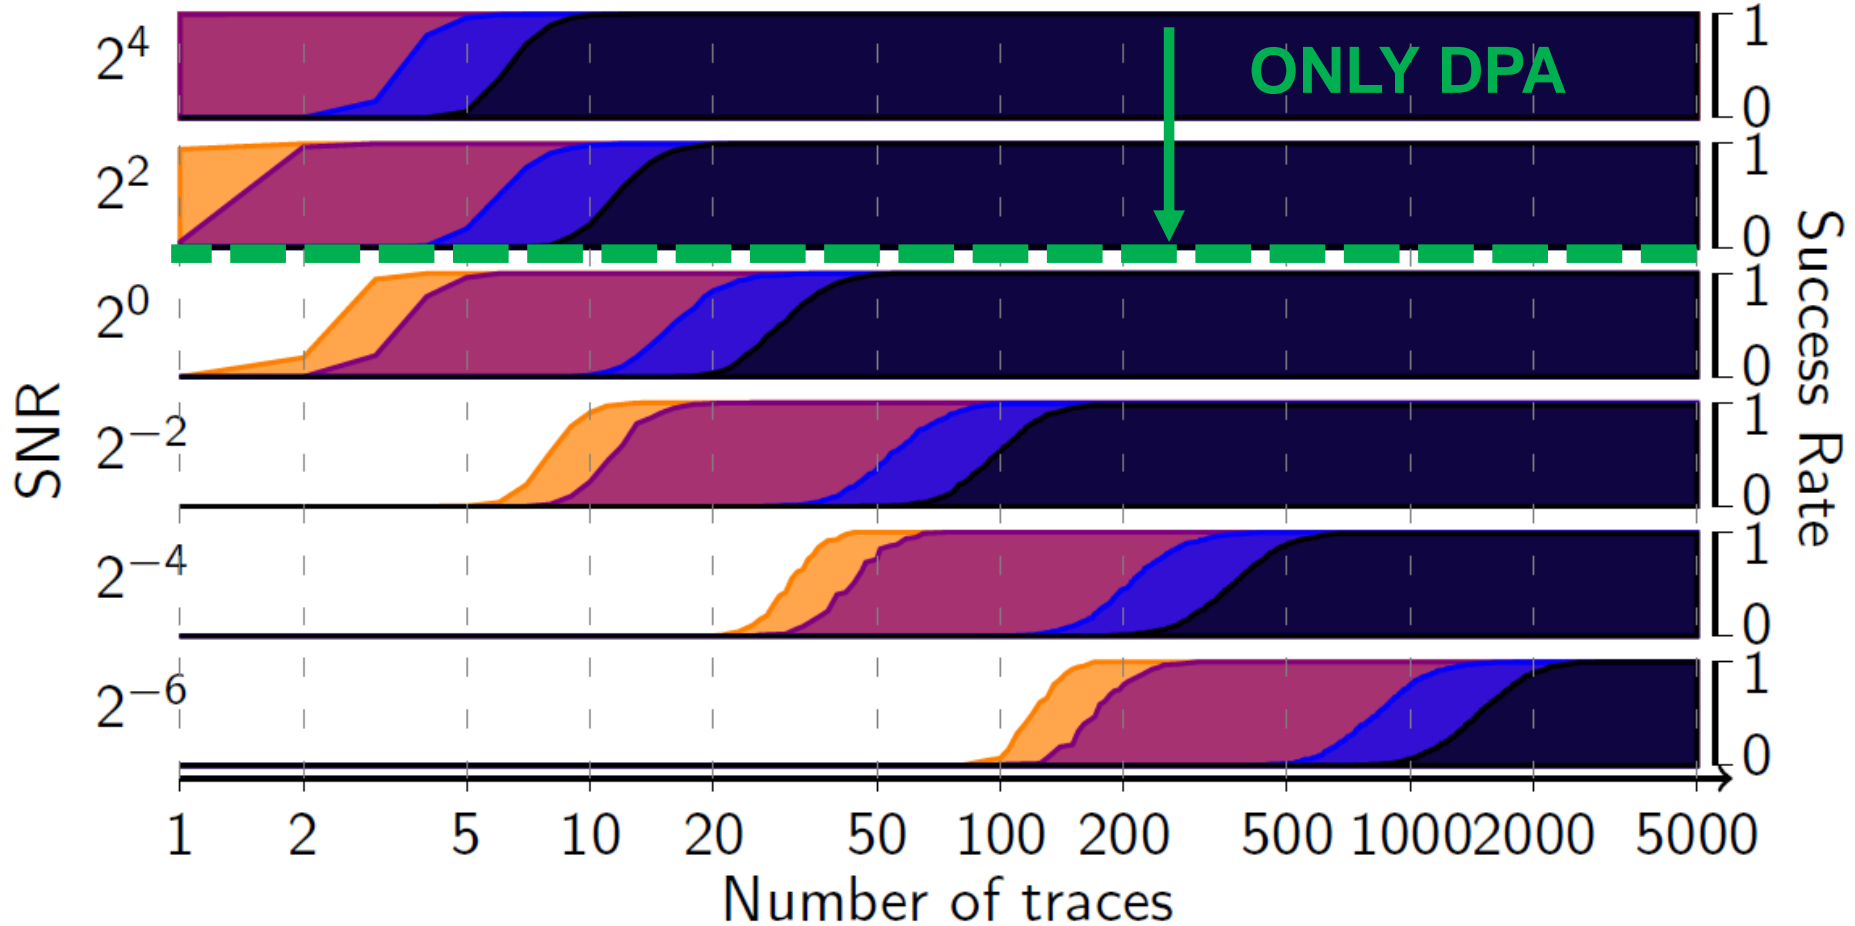
- Simulated HW leakages with variable noise



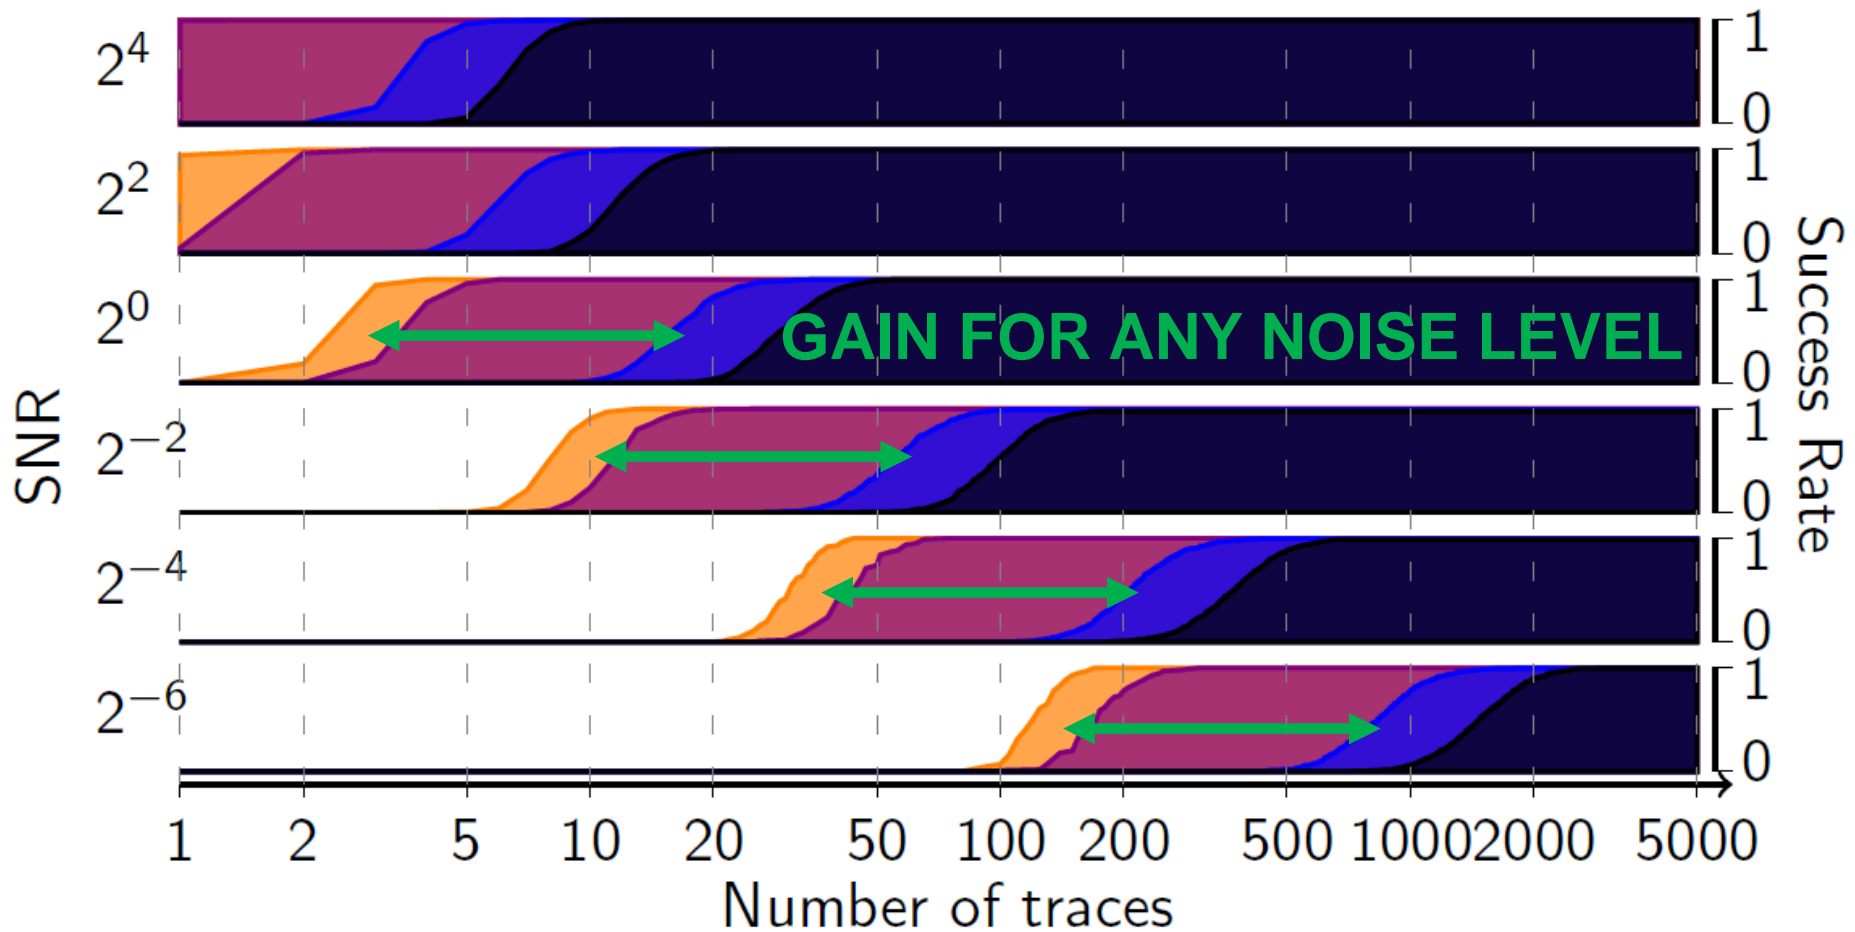
- univariate TA, SBOX output
- bivariate TA, SBOX input and output
- SASCA attack, no key schedule leakages
- SASCA attack, all intermediate values



- univariate TA, SBOX output
- bivariate TA, SBOX input and output
- SASCA attack, no key schedule leakages
- SASCA attack, all intermediate values



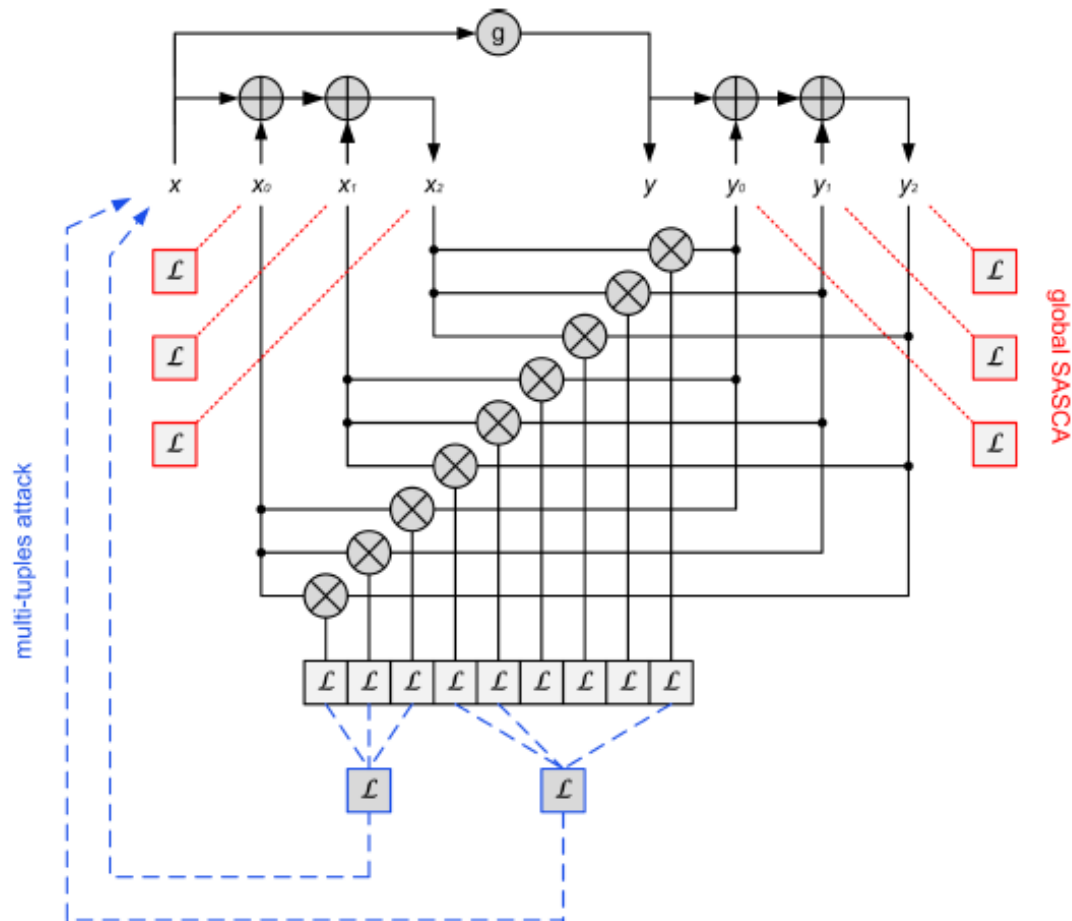
- univariate TA, SBOX output
- bivariate TA, SBOX input and output
- SASCA attack, no key schedule leakages
- SASCA attack, all intermediate values



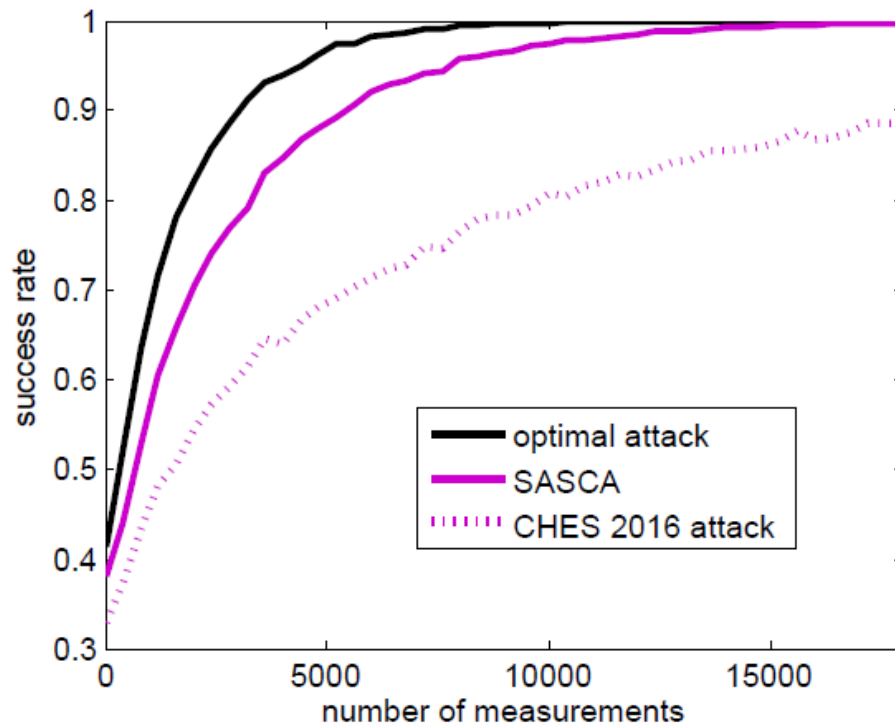
- univariate TA, SBOX output
- bivariate TA, SBOX input and output
- SASCA attack, no key schedule leakages
- SASCA attack, all intermediate values

- Also works in practice (Asiacrypt 2015)

- Also works in practice (Asiacrypt 2015)
- And against masked implementations



- Also works in practice (Asiacrypt 2015)
- And against masked implementations
 - Improvement of the CHES 2016 horizontal attacks



Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- **Post-processing**
 - Key enumeration, **rank estimation**
- Future trends
 - Security without obscurity
 - IT metrics & (tight) proofs

- Enough measurements \Rightarrow direct key recovery

- Enough measurements \Rightarrow direct key recovery
- Not enough measurements but enough computational power \Rightarrow key enumeration

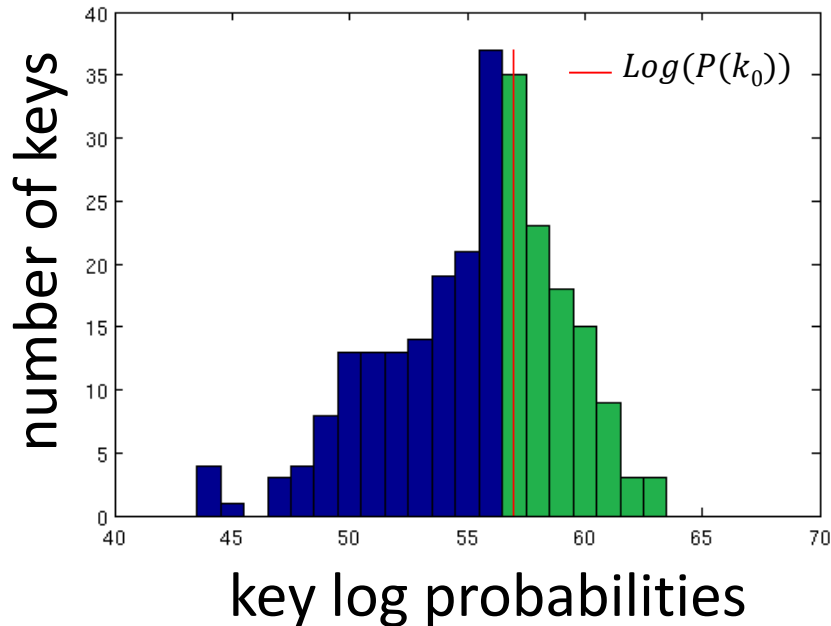
- Enough measurements \Rightarrow direct key recovery
- Not enough measurements but enough computational power \Rightarrow key enumeration
- But what to do if it is not enough?
 - Key rank estimation (requires key knowledge)
 - \Rightarrow Only possible for evaluators

- Enough measurements \Rightarrow direct key recovery
- Not enough measurements but enough computational power \Rightarrow key enumeration
- But what to do if it is not enough?
 - Key rank estimation (requires key knowledge)
 \Rightarrow Only possible for evaluators
- Note: only optimal with probabilities (to combine the information of \neq S-boxes)

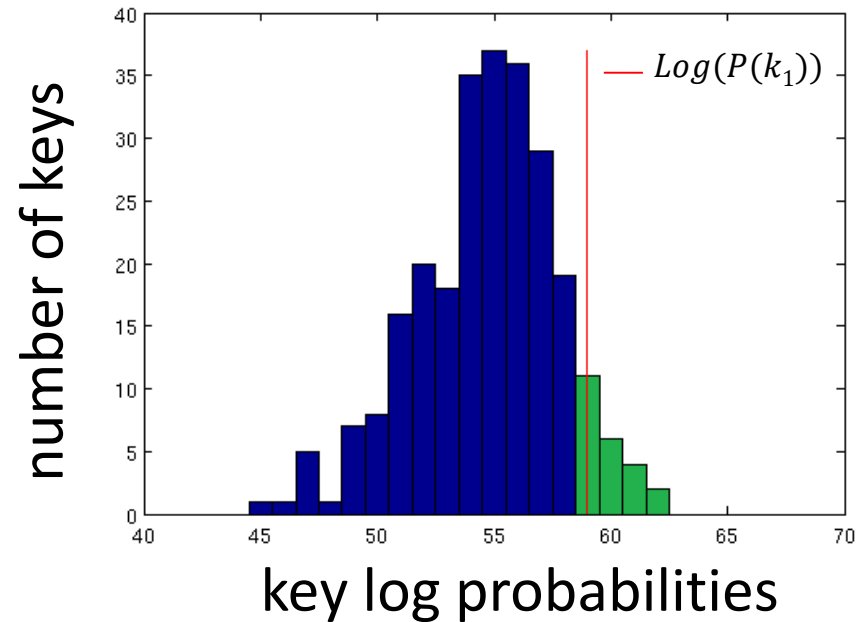
- Thanks to T-systems people (Glowacz, Schueth)

- Thanks to T-systems people (Glowacz, Schueth)
- Representation with histograms

S-box 0

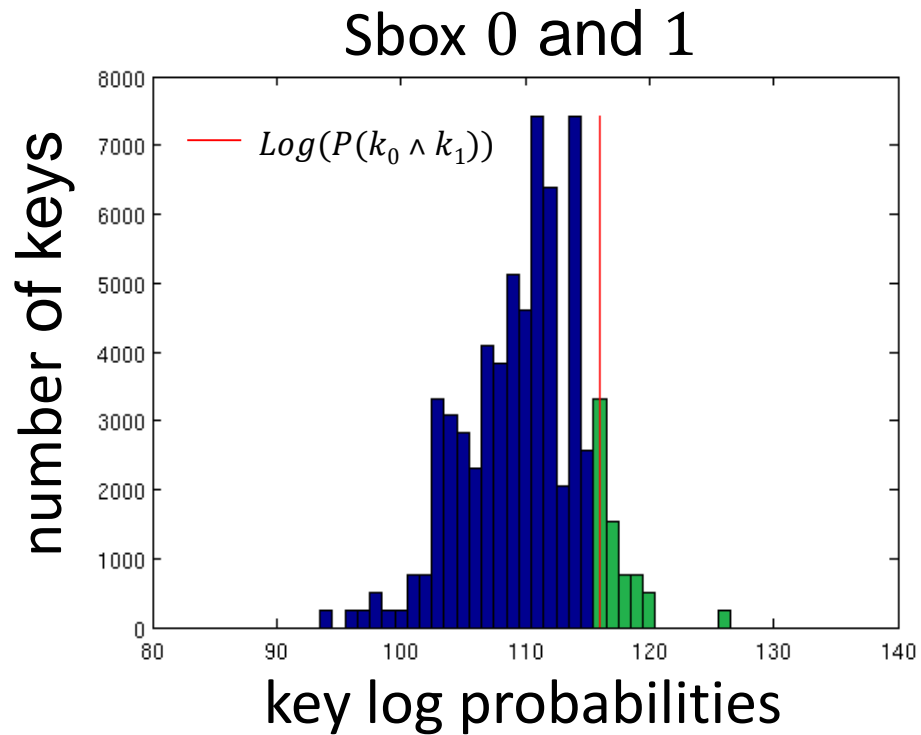


S-box 1

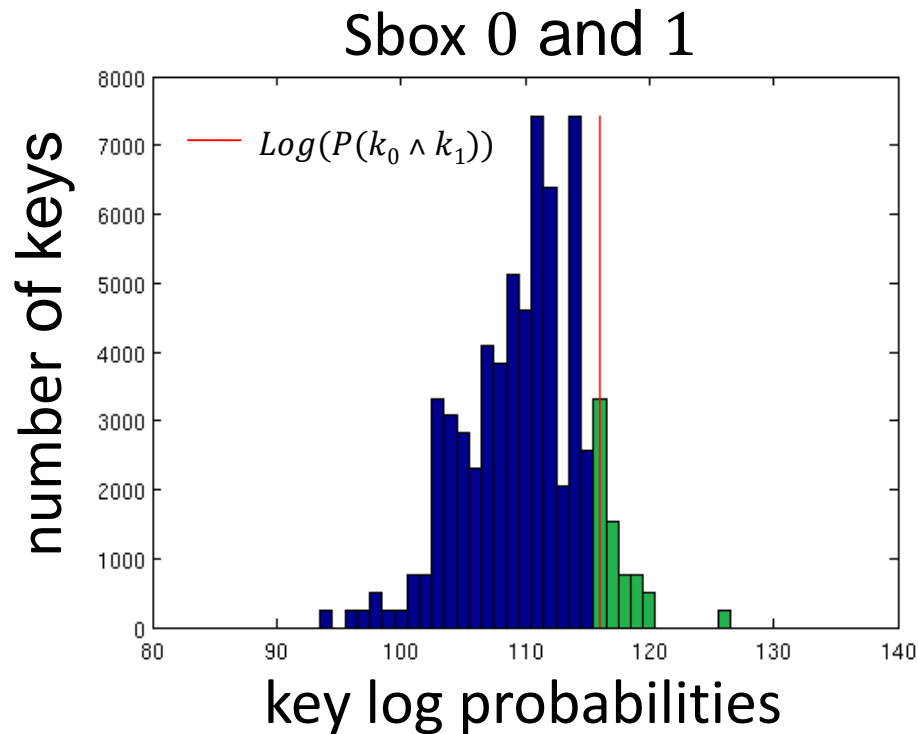


- The rank is the number of key in the green zone

- Combination with convolution



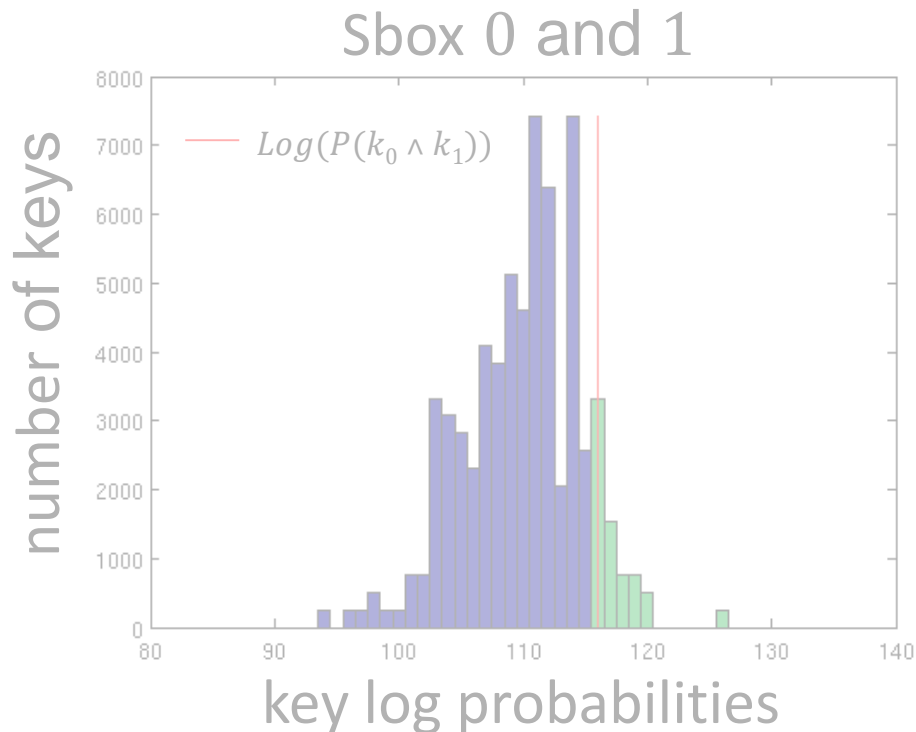
- Combination with convolution



... and keep track of the error that depends on:

- The number of bins
- The number of conv.

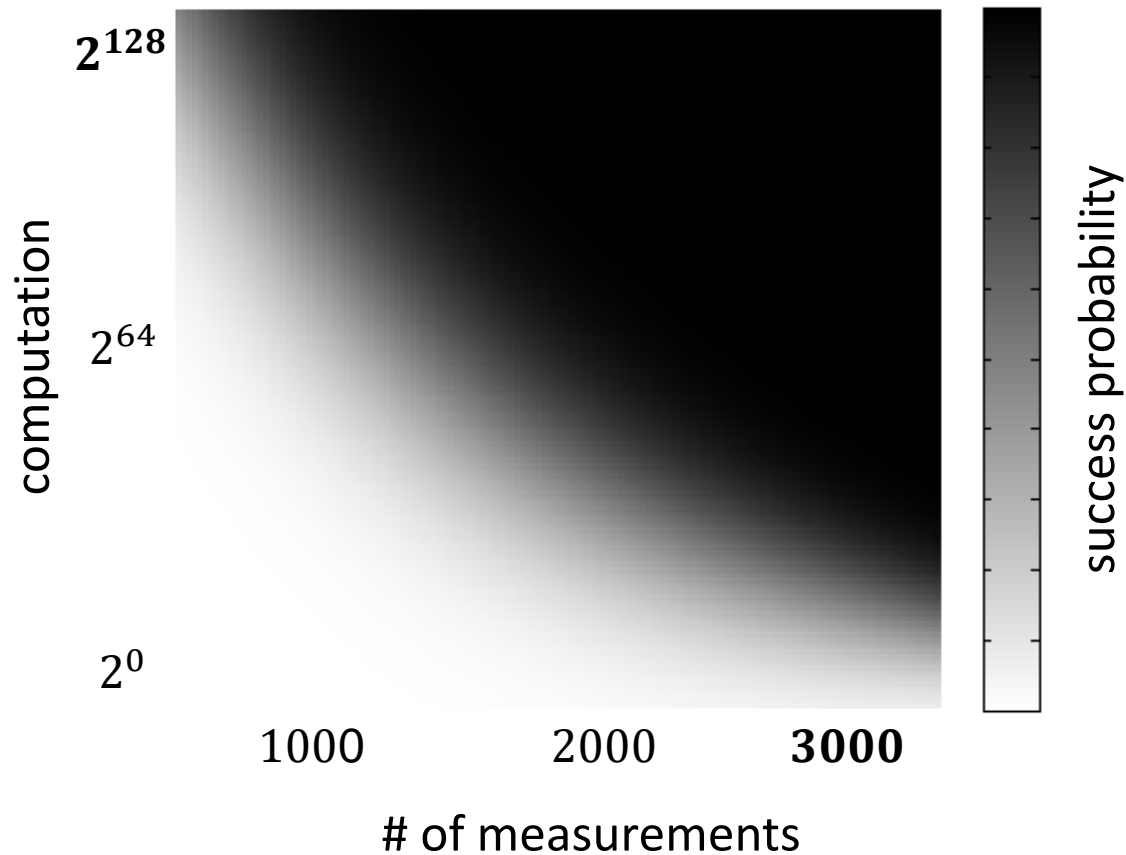
- Combination with convolution



... and keep track of the error that depends on:

- The number of bins
 - The number of conv.
- Just iterating this gives the key rank accurately
 - e.g., < 1 bits in < 1 sec. for a 128-bit key

- Security graph (IMO the sound outcome of an evaluation)



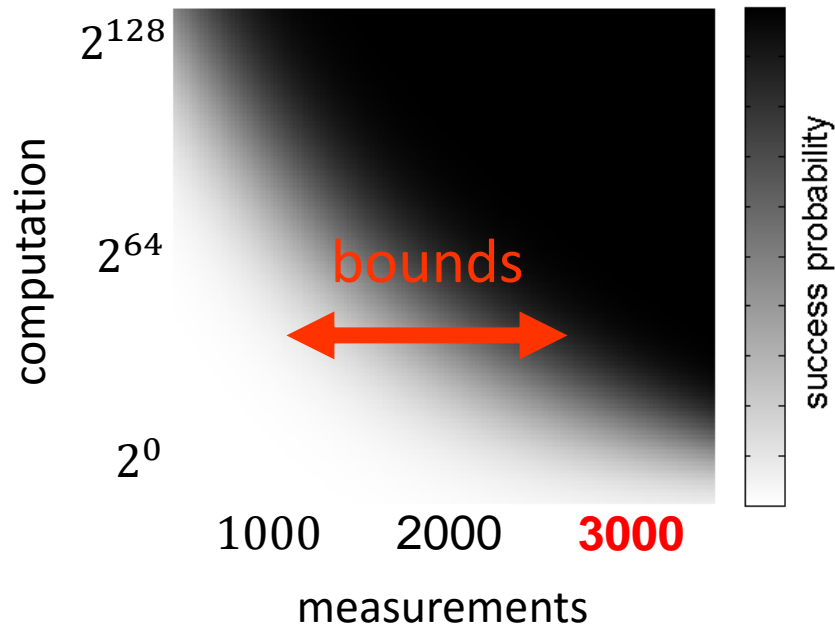
Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- **Future trends**
 - **Security without obscurity**
 - IT metrics & (tight) proofs

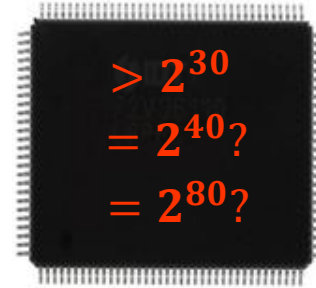
standard practice



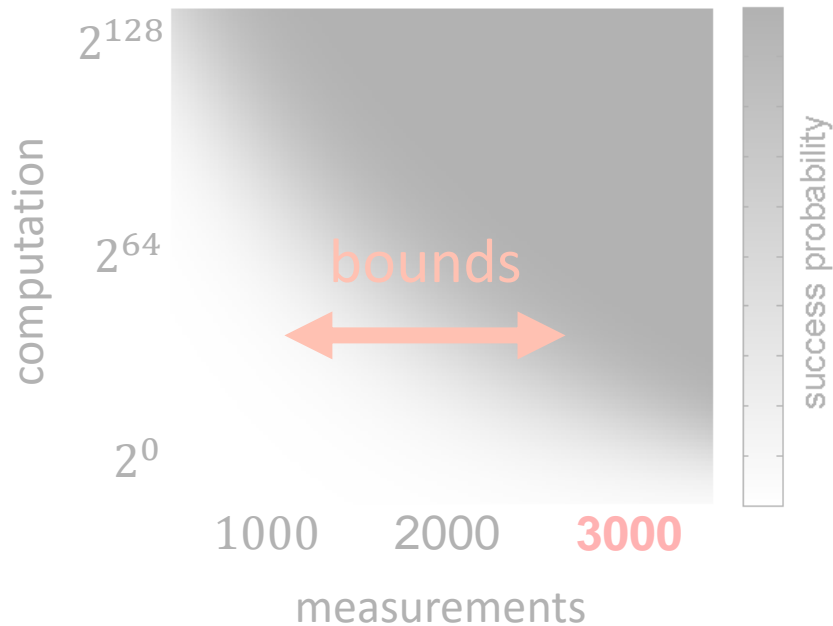
attack-based evaluations



standard practice



attack-based evaluations



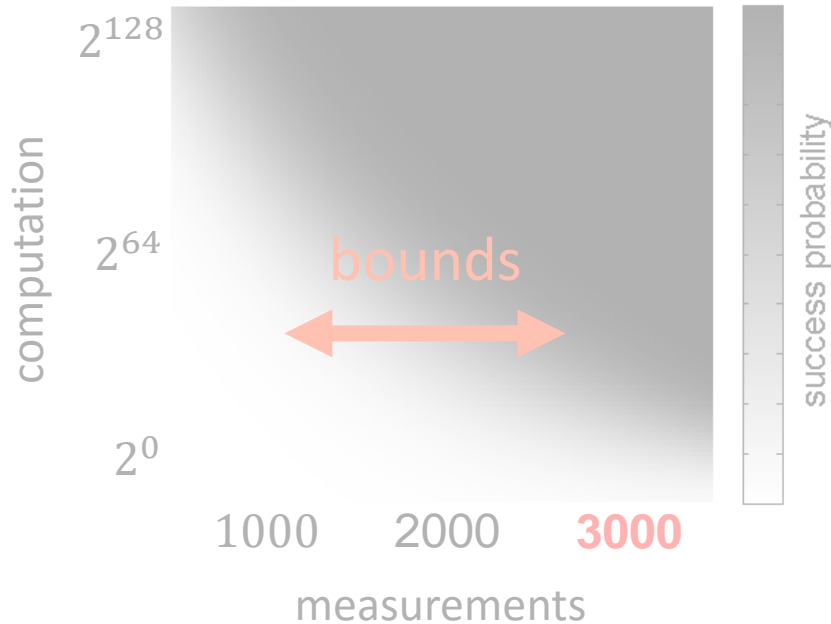
standard practice



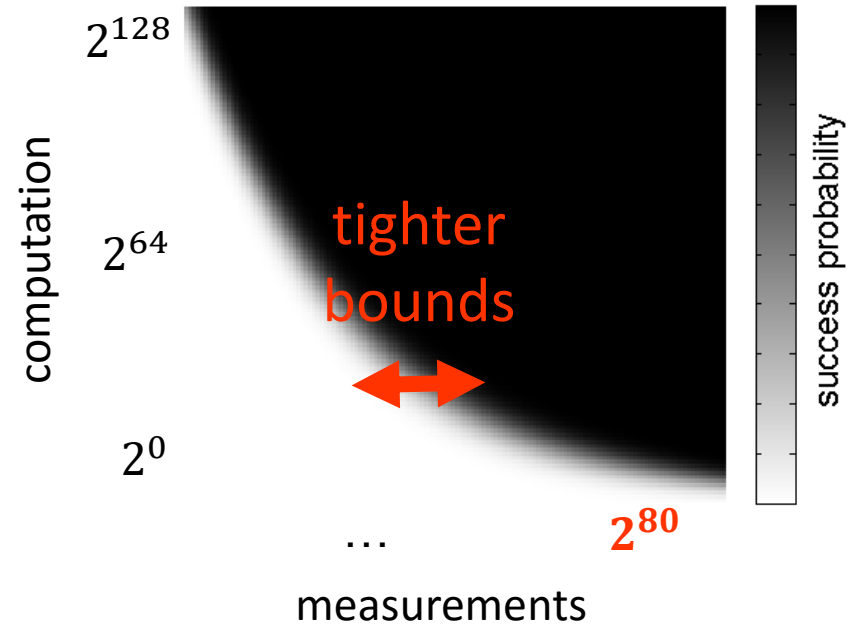
transparency
helps evaluations



attack-based evaluations



proof-based evaluations



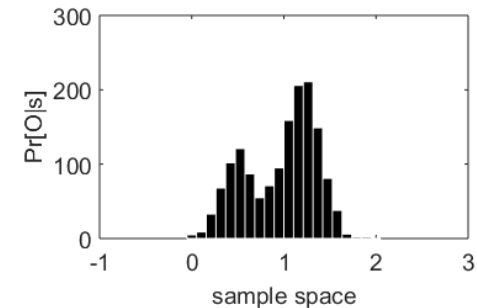
Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- **Future trends**
 - Security without obscurity
 - **IT metrics & (tight) proofs**

- Discrete observations:

$$MI(S; O) = H[S] + \sum_{s \in S} \Pr[s] \cdot \sum_{o \in O} \Pr[o|s] \cdot \log_2(\Pr[s|o])$$

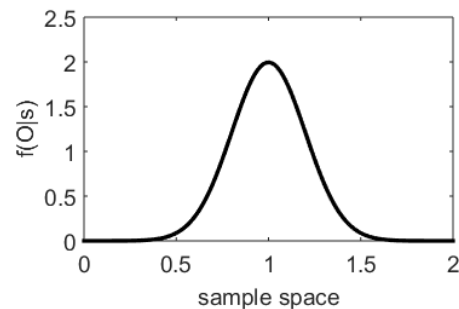
- With: $\Pr[s|o] = \frac{\Pr[o|s]}{\sum_{s^*} \Pr[o|s^*]}$



- Continuous observations:

$$MI(S; O) = H[S] + \sum_{s \in S} \Pr[s] \cdot \int_0 f(o|s) \cdot \log_2(\Pr[s|o]) \, do$$

- With: $\Pr[s|o] = \frac{f(o|s)}{\sum_{s^*} f(o|s^*)}$



- Continuous observations:

$$MI(S; O) = H[S] + \sum_{s \in S} \Pr[s] \cdot \int_o f(o|s) \cdot \log_2(\Pr[s|o]) \, do$$

- With: $\Pr[s|o] = \frac{f(o|s)}{\sum_{s^*} f(o|s^*)}$

- Intuition: average amount of information on S that is gained by observing a sample o

- Higher $MI(S; O)$ asymptotically implies higher Bayesian classification success rate:

$$SR(n) = \Pr_O \left[\operatorname{argmax}_{s^*} \Pr[s^* | o_1] \cdot \Pr[s^* | o_2] \cdots \Pr[s^* | o_n] = s \right]$$

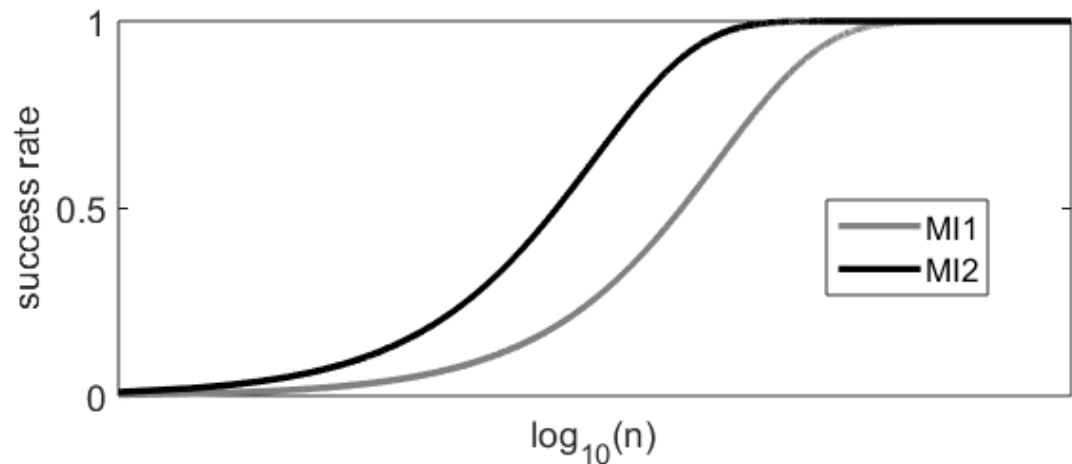
- With n observations used for classification

- Higher $MI(S; O)$ asymptotically implies higher Bayesian classification success rate:

$$SR(n) = \Pr_O \left[\operatorname{argmax}_{s^*} \Pr[s^* | o_1] \cdot \Pr[s^* | o_2] \cdots \Pr[s^* | o_n] = s \right]$$

- With n observations used for classification

Example, with
 $MI2 > MI1$



- In practice, the true distributions $\text{Pr}_{\text{real}}[o|s]$ and $f_{\text{real}}(o|s)$ are generally unknown
 - So they can be sampled (i.e., measured)
 - But they cannot be computed exactly

- In practice, the true distributions $\text{Pr}_{\text{real}}[o|s]$ and $f_{\text{real}}(o|s)$ are generally unknown
 - So they can be sampled (i.e., measured)
 - But they cannot be computed exactly
- Hence, for the classification we use statistical models $\text{Pr}_{\text{model}}[o|s]$ and $f_{\text{model}}(o|s)$
 - But these models can suffer from both estimation errors and assumption errors

- In practice, the true distributions $\text{Pr}_{\text{real}}[o|s]$ and $f_{\text{real}}(o|s)$ are generally unknown
 - So they can be sampled (i.e., measured)
 - But they cannot be computed exactly
 - Hence, for the classification we use statistical models $\text{Pr}_{\text{model}}[o|s]$ and $f_{\text{model}}(o|s)$
 - But these models can suffer from both estimation errors and assumption errors
- ⇒ How to be sure that the model is good?*

- Exactly reflects this practical challenge
- For example in the continuous case:

$$PI(S; O) = H[S] + \sum_{s \in S} \Pr[s] \cdot \int_0 f_{\text{real}}(o|s) \cdot \log_2(\Pr_{\text{model}}[s|o]) do$$

- Exactly reflects this practical challenge
- For example in the continuous case:

$$PI(S; O) = H[S] + \sum_{s \in S} \Pr[s] \cdot \int_0 f_{\text{real}}(o|s) \cdot \log_2(\Pr_{\text{model}}[s|o]) do$$

- Intuition: average amount of information on S that is gained by observing a sample o , biased by the model estimation/assumption errors

- Exactly reflects this practical challenge
- For example in the continuous case:

$$PI(S; O) = H[S] + \sum_{s \in S} \Pr[s] \cdot \int_0 f_{\text{real}}(o|s) \cdot \log_2(\Pr_{\text{model}}[s|o]) do$$

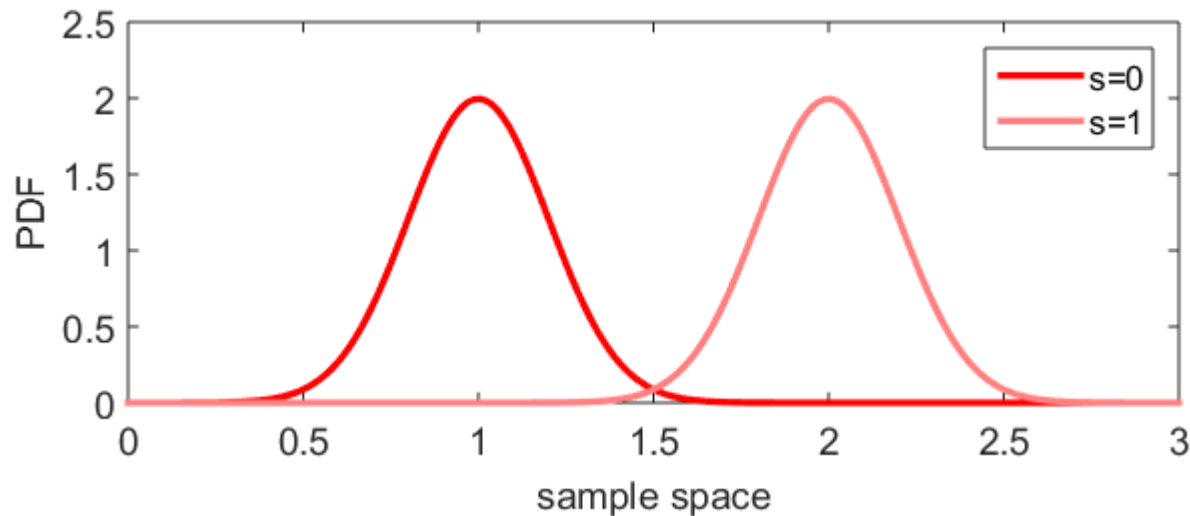
- Intuition: average amount of information on S that is gained by observing a sample o , biased by the model estimation/assumption errors
- $PI(S; O)$ is a statistical distance between the real and modeled distributions (i.e., a measure of how well a model “explains” real observations)

- Say you have a set of $2m$ observations per s

- Say you have a set of $2m$ observations per s

1. Use m observations to build a model for each s :

$$\hat{f}_{\text{model}}(o|s) \stackrel{m}{\leftarrow} f_{\text{real}}(o|s)$$



- Say you have a set of $2m$ observations per s

1. Use m observations to build a model for each s :

$$\hat{f}_{\text{model}}(o|s) \stackrel{m}{\leftarrow} f_{\text{real}}(o|s)$$

2. Use the other m observations to test the model:

$$\hat{\text{PI}}(S; O) = H[S] + \sum_{s \in S} \text{Pr}[s] \cdot \sum_{o' \stackrel{m}{\leftarrow} f_{\text{real}}(o|s)} \frac{1}{m} \cdot \log_2(\hat{\text{Pr}}_{\text{model}}[s|o'])$$

- Say you have a set of $2m$ observations per s

1. Use m observations to build a model for each s :

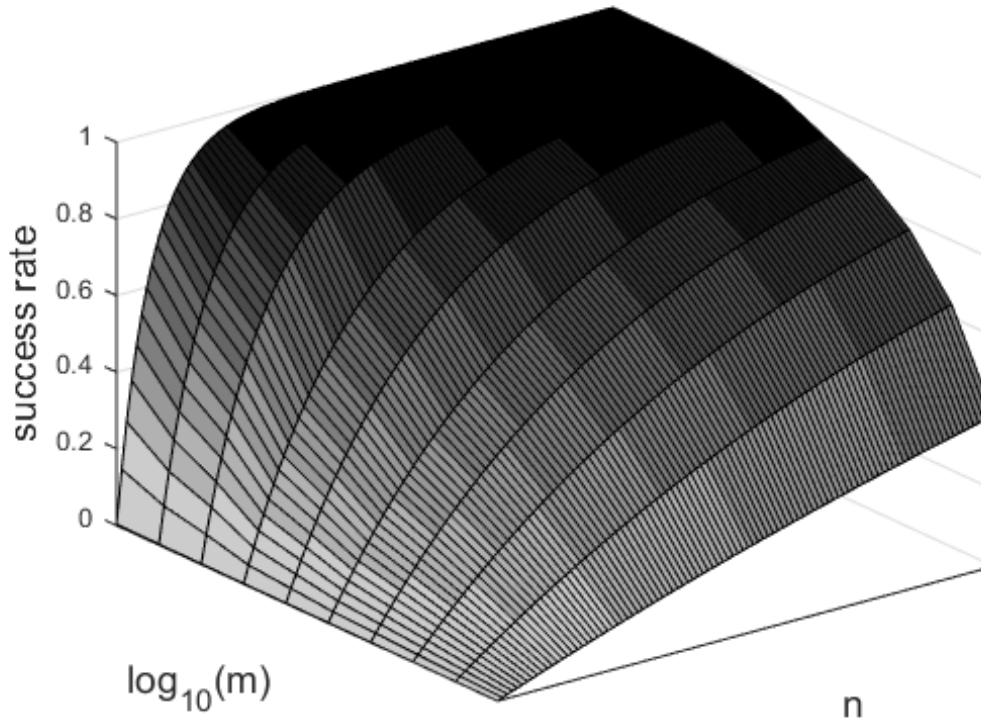
$$\hat{f}_{\text{model}}(o|s) \stackrel{m}{\leftarrow} f_{\text{real}}(o|s)$$

2. Use the other m observations to test the model:

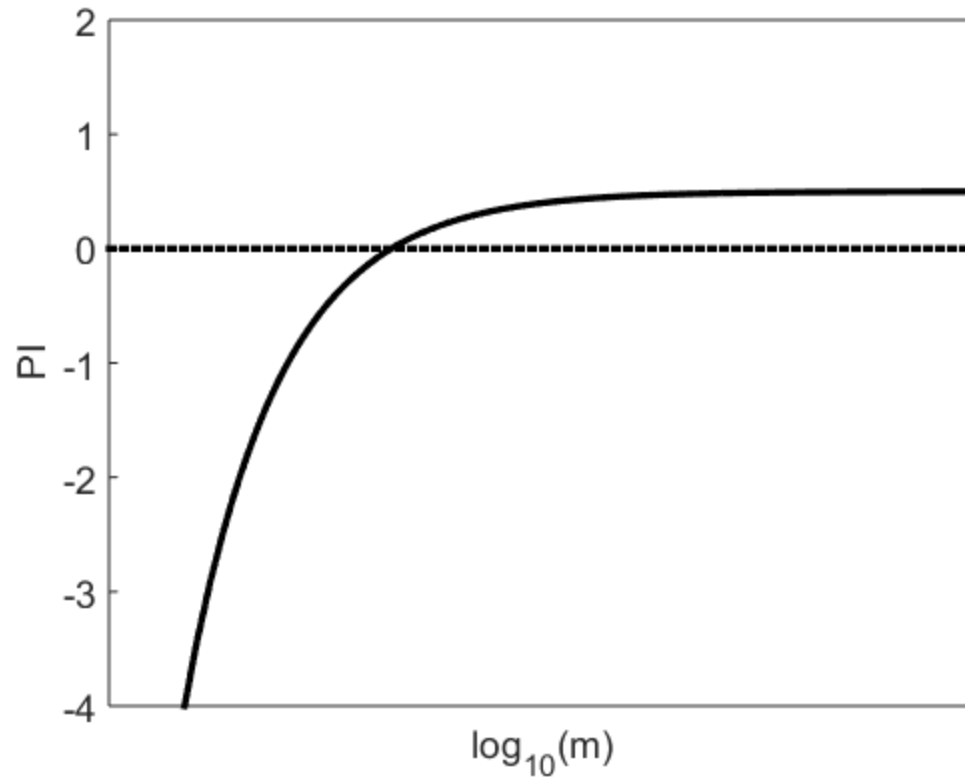
$$\hat{\text{PI}}(S; O) = H[S] + \sum_{s \in S} \text{Pr}[s] \cdot \sum_{o' \stackrel{m}{\leftarrow} f_{\text{real}}(o|s)} \frac{1}{m} \cdot \log_2(\hat{\text{Pr}}_{\text{model}}[s|o'])$$

(More efficient use of observations with cross-validation)

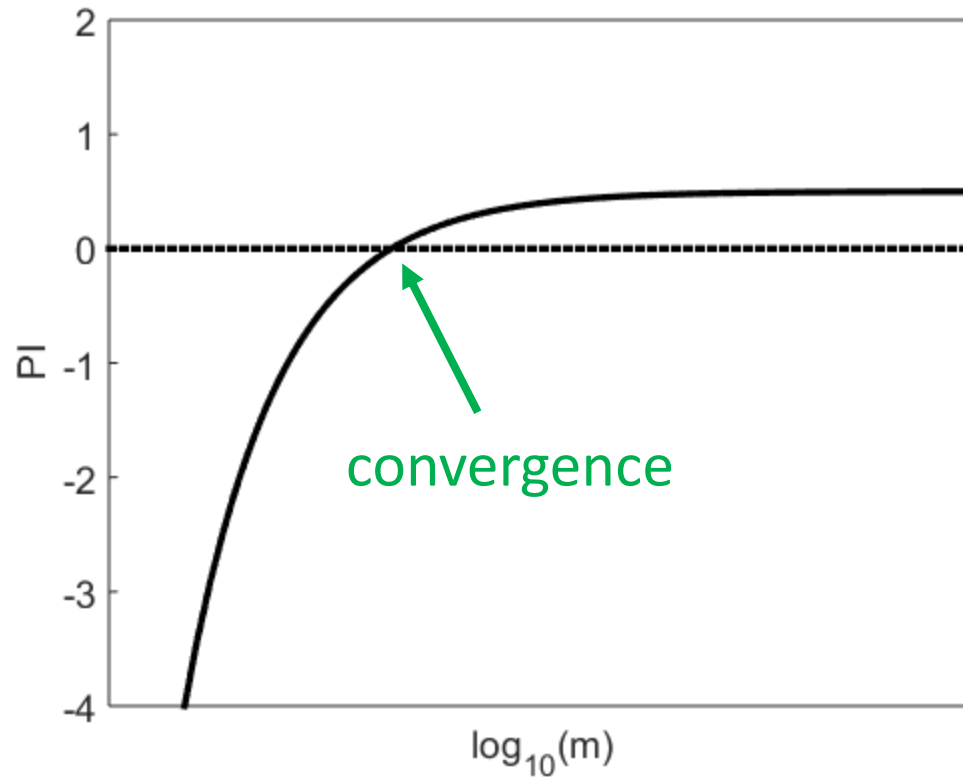
- Directly estimating model *convergence* and *informativeness* based on the SR is expensive



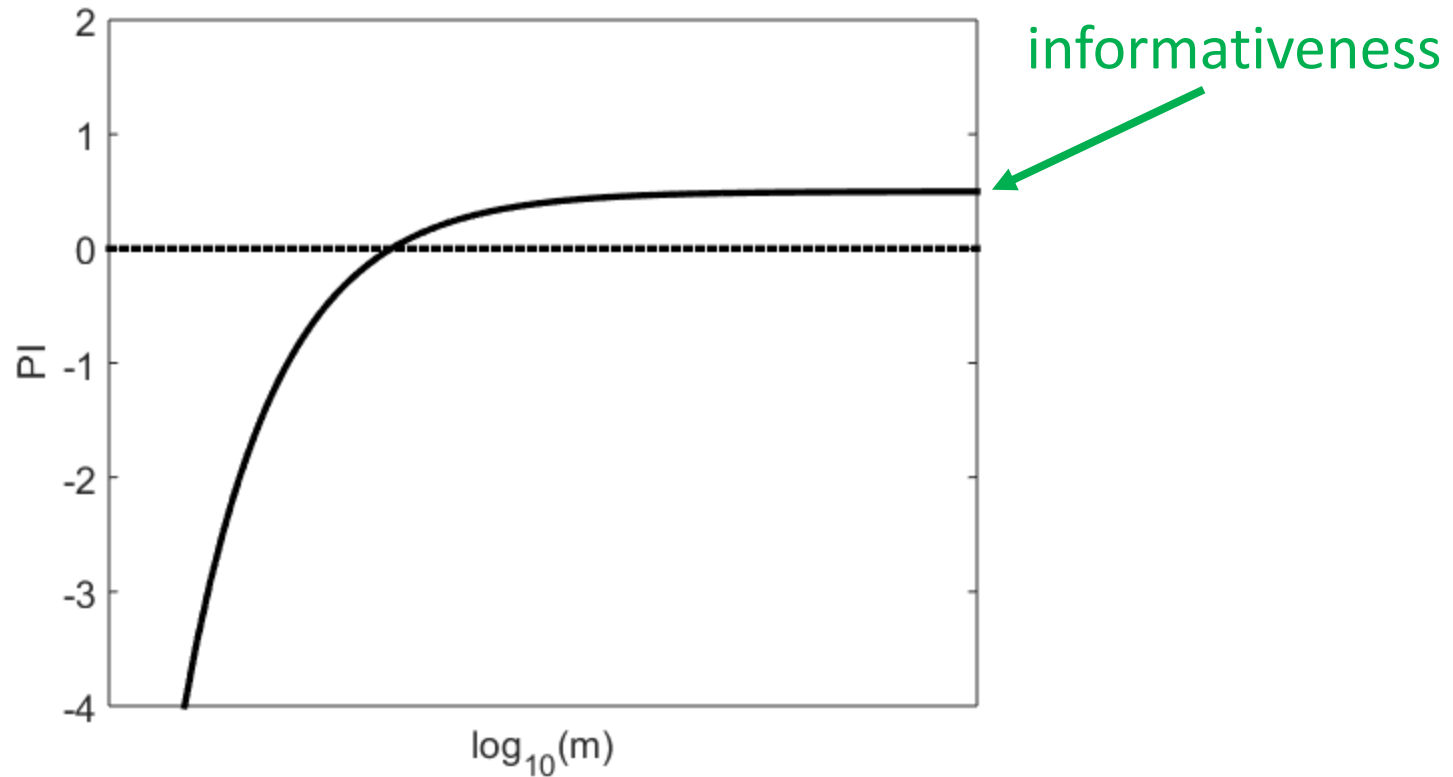
- The PI curve allows “getting rid of” the n axis
⇒ it is faster to estimate than the SR surface



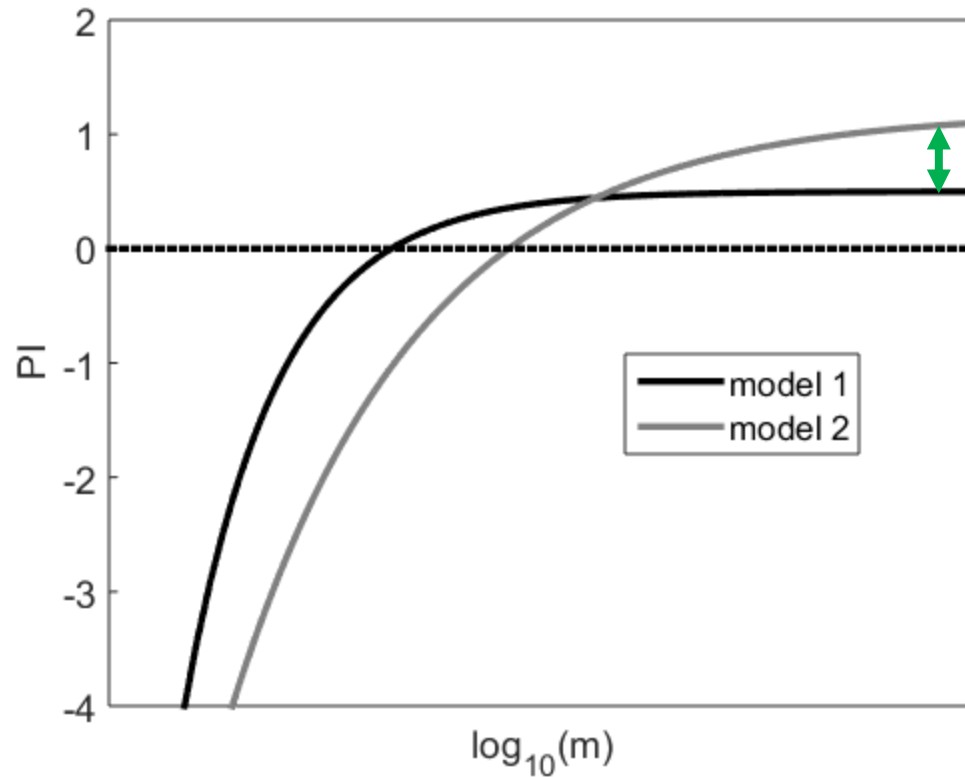
- The PI curve allows “getting rid of” the n axis
⇒ it is faster to estimate than the SR surface



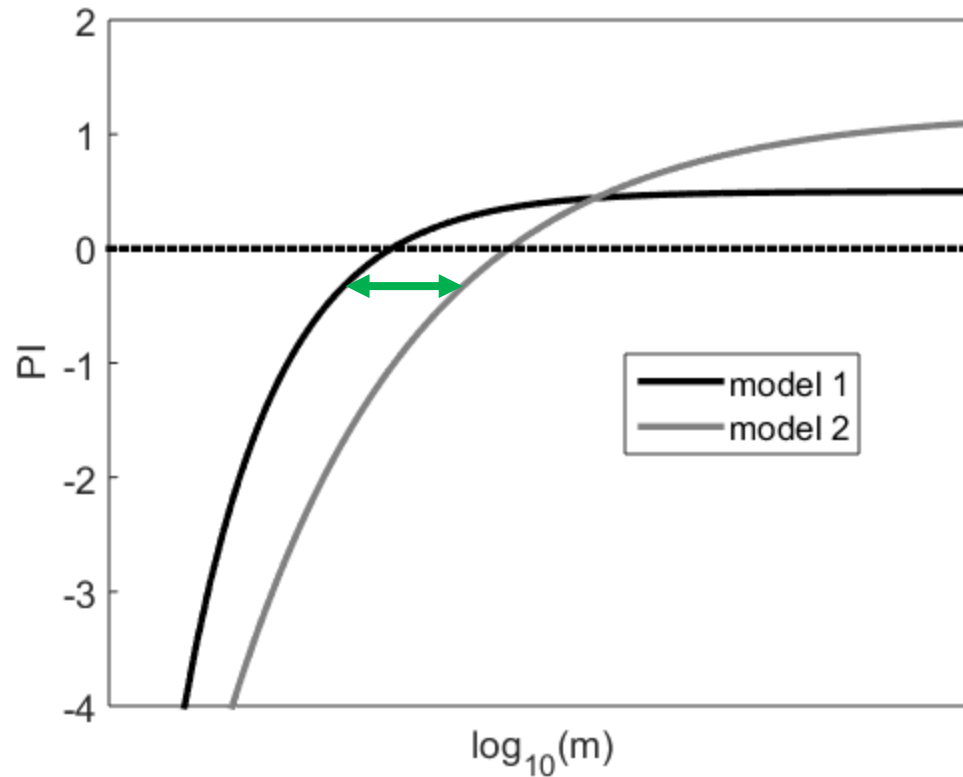
- The PI curve allows “getting rid of” the n axis
⇒ it is faster to estimate than the SR surface



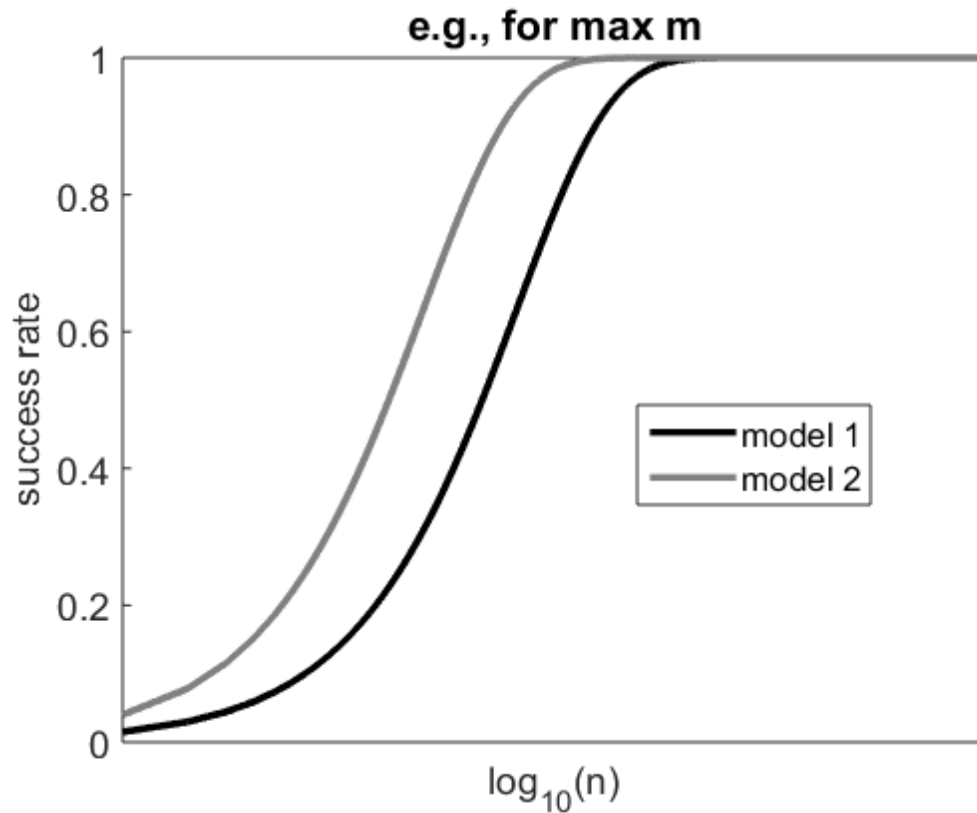
- Allows sound & efficient comparison of models
 - e.g., model 2 more informative than model 1



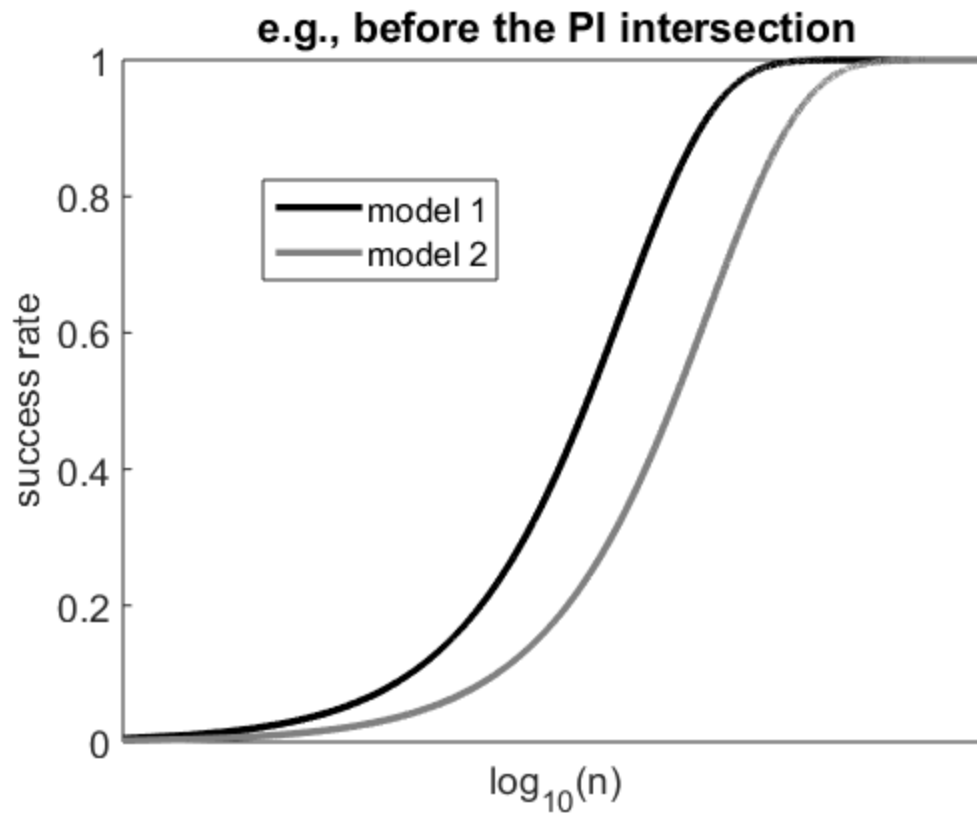
- Allows sound & efficient comparison of models
 - and model 1 converges faster than model 2



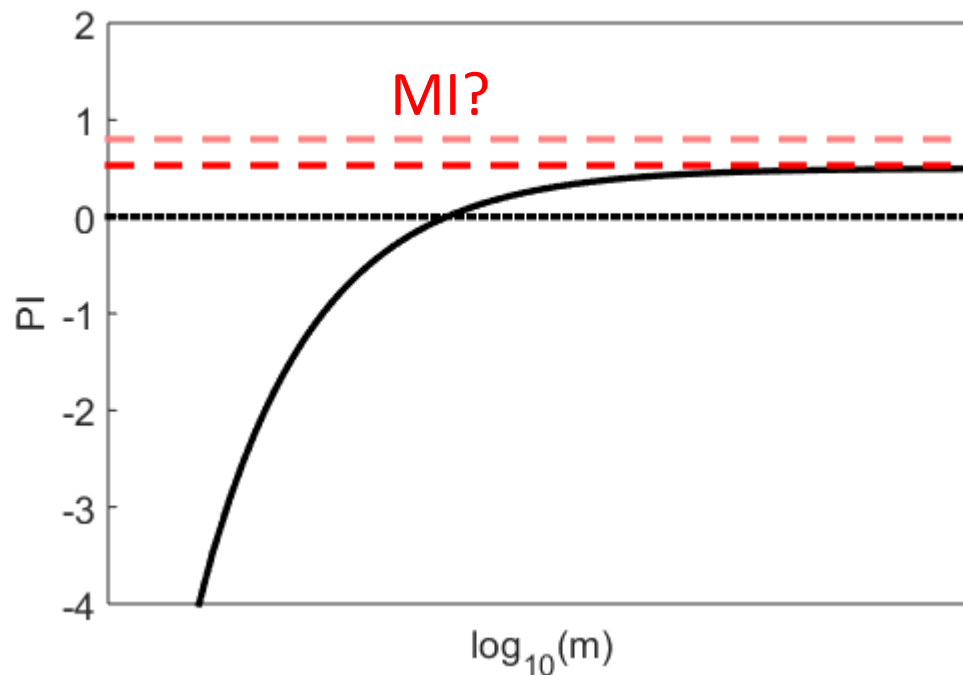
- For a given m , one can always compute the success rate curves to gain concrete intuition



- For a given m , one can always compute the success rate curves to gain concrete intuition

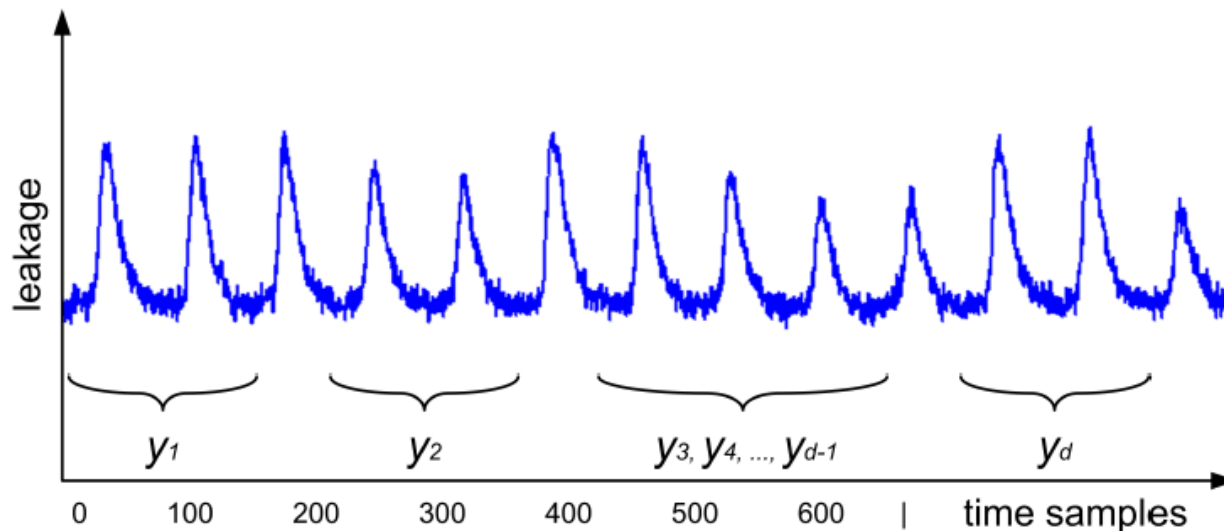


- Remaining question: how far are we from the MI?



- *Leakage certification* allows answering this!

- 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

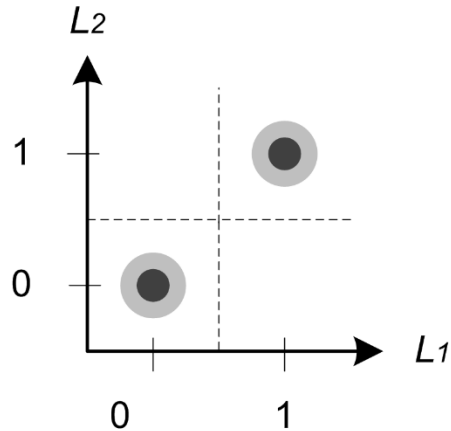
⇒ Composable (from gadgets to circuits)

- 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: $\text{SR} \leq 1 - \left(1 - \text{MI}(Z_i; L_{Z_i})\right)^d)^m$

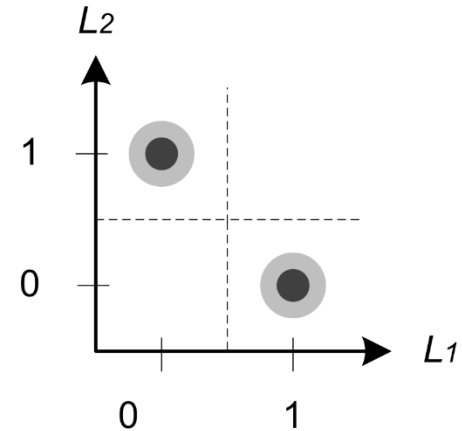
- Assume leakage variables $L_{Z_i} = \delta(Z_i) + N$ s.t.
 - $\text{MI}(Z_i; L_{Z_i}) \leq \frac{c}{d}$ (multiplications)
 - The leakages of the shares are independent
- For a masking scheme with d shares
- And an adversary using m measurements
- Then: $\text{SR} \leq 1 - \left(1 - \text{MI}(Z_i; L_{Z_i})\right)^d)^m$

- 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: $\text{SR} \leq 1 - \left(1 - \text{MI}(Z_i; L_{Z_i})^d\right)^m$
- 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)

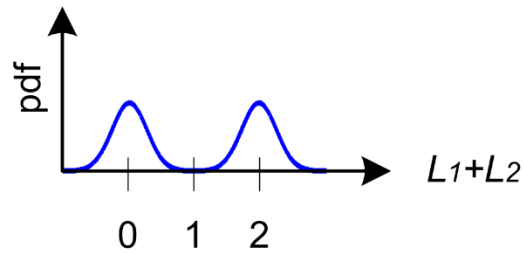
- 1-bit, 2-shares example



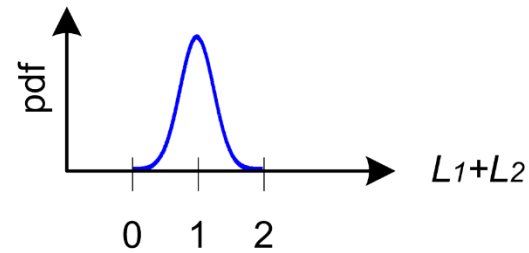
(a) $Z = 0$, serial.



(b) $Z = 1$, serial.



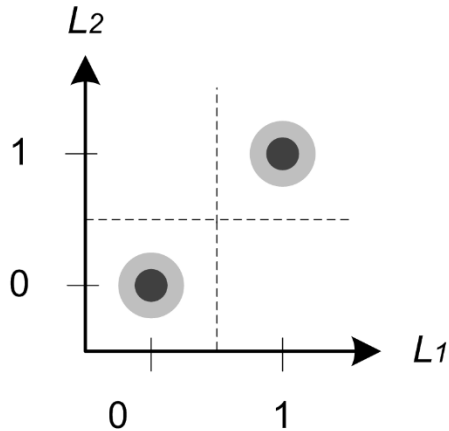
(c) $Z = 0$, parallel.



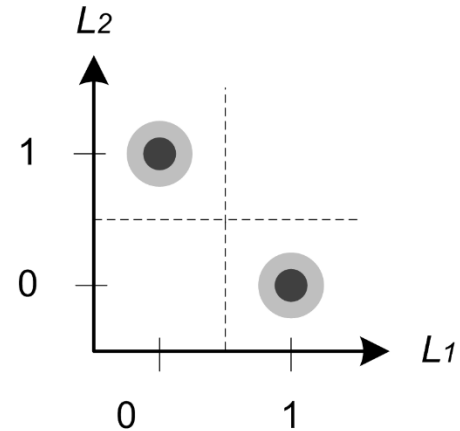
(d) $Z = 1$, parallel.

- 1-bit, 2-shares example

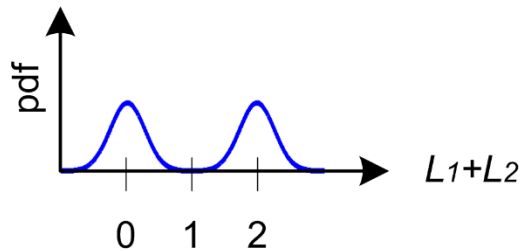
key-independent means



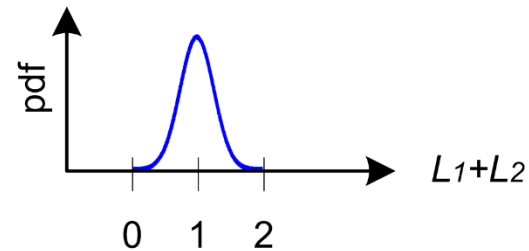
(a) $Z = 0$, serial.



(b) $Z = 1$, serial.

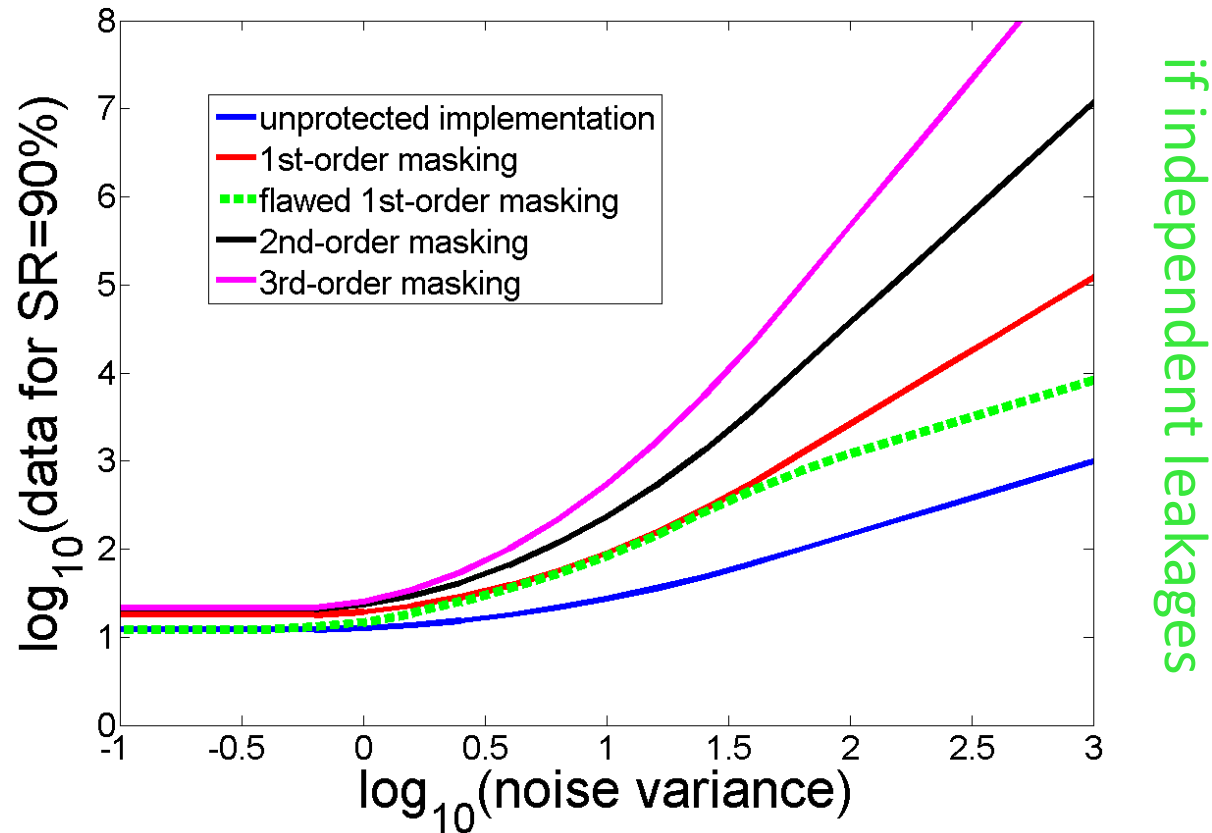


(c) $Z = 0$, parallel.



(d) $Z = 1$, parallel.

- Slope of the IT curves = d (i.e., security order)
 - e.g., for information leakage of an encoding



- 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

- And that t-tests become irrelevant with large #dim.

Outline

- Introduction
 - Side-channel analysis (attack steps)
 - Heuristic vs. optimal separation
- Measurement & preprocessing
 - Filtering, leakage/POI detection, dimension. reduction
- Predictions & modeling
 - Profiled vs. non-profiled separation, leakage certification
- Exploitation
 - Soft Analytical Side-Channel Attacks
- Post-processing
 - Key enumeration, rank estimation
- Future trends
 - Security without obscurity
 - Exploiting (tight) proofs

Wrapping up!

- For some parts, verifiably fair (i.e., close to worst-case) security evaluations are possible
 - But measurements & preprocessing remain essentially based on engineering knowledge
 - & there remain challenges for highly multivariate and (very) high-order side-channel attacks

- For some parts, verifiably fair (i.e., close to worst-case) security evaluations are possible
 - But measurements & preprocessing remain essentially based on engineering knowledge
 - & there remain challenges for highly multivariate and (very) high-order side-channel attacks
- **Transparency is needed for high security**
 - e.g., HW with 2^{80} security should be open source

- For some parts, verifiably fair (i.e., close to worst-case) security evaluations are possible
 - But measurements & preprocessing remain essentially based on engineering knowledge
 - & there remain challenges for highly multivariate and (very) high-order side-channel attacks
- Transparency is needed for high security
 - e.g., HW with 2^{80} security should be open source
- **First focus should be on understanding (adv.'s practicality comes only afterwards)**
 - e.g., thing about linear cryptanalysis

- Effective countermeasures against side-channel attacks always combine sound hardware assumptions & mathematical amplification
 - ⇒ Empirically verifiable (falsifiable) assumptions
 - ⇒ Systematic ways to deal with hardware defaults (or constructions that are less demanding)
 - ⇒ Tight proofs in (reasonably) realistic models

- Effective countermeasures against side-channel attacks always combine sound hardware assumptions & mathematical amplification
 - ⇒ Empirically verifiable (falsifiable) assumptions
 - ⇒ Systematic ways to deal with hardware defaults (or constructions that are less demanding)
 - ⇒ Tight proofs in (reasonably) realistic models
- **Tools, formal methods, design automation**

- Effective countermeasures against side-channel attacks always combine sound hardware assumptions & mathematical amplification
 - ⇒ Empirically verifiable (falsifiable) assumptions
 - ⇒ Systematic ways to deal with hardware defaults (or constructions that are less demanding)
 - ⇒ Tight proofs in (reasonably) realistic models
- Tools, formal methods, design automation
- **We need both theoretical works to lay out foundations & experimental case studies**

THANKS

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

<http://perso.uclouvain.be/fstandae/PUBLIS/183.pdf>