“Noisy” vs. “Bounded” Leakage

João Ribeiro

Inst. Telecomunicações & Técnico - U Lisboa

Based on joint work with

Gianluca Brian
Sapienza (now ETHZ)

Antonio Faonio
EURECOM

Maciej Obremski
NUS

Lawrence Roy
Aarhus

Maciej Skórski
Luxembourg

François-Xavier Standaert
UC Louvain

Mark Simkin
Aarhus/EF

Daniele Venturi
Sapienza

Side-channel attacks

Attacks on cryptographic schemes exploiting physical hardware quirks.

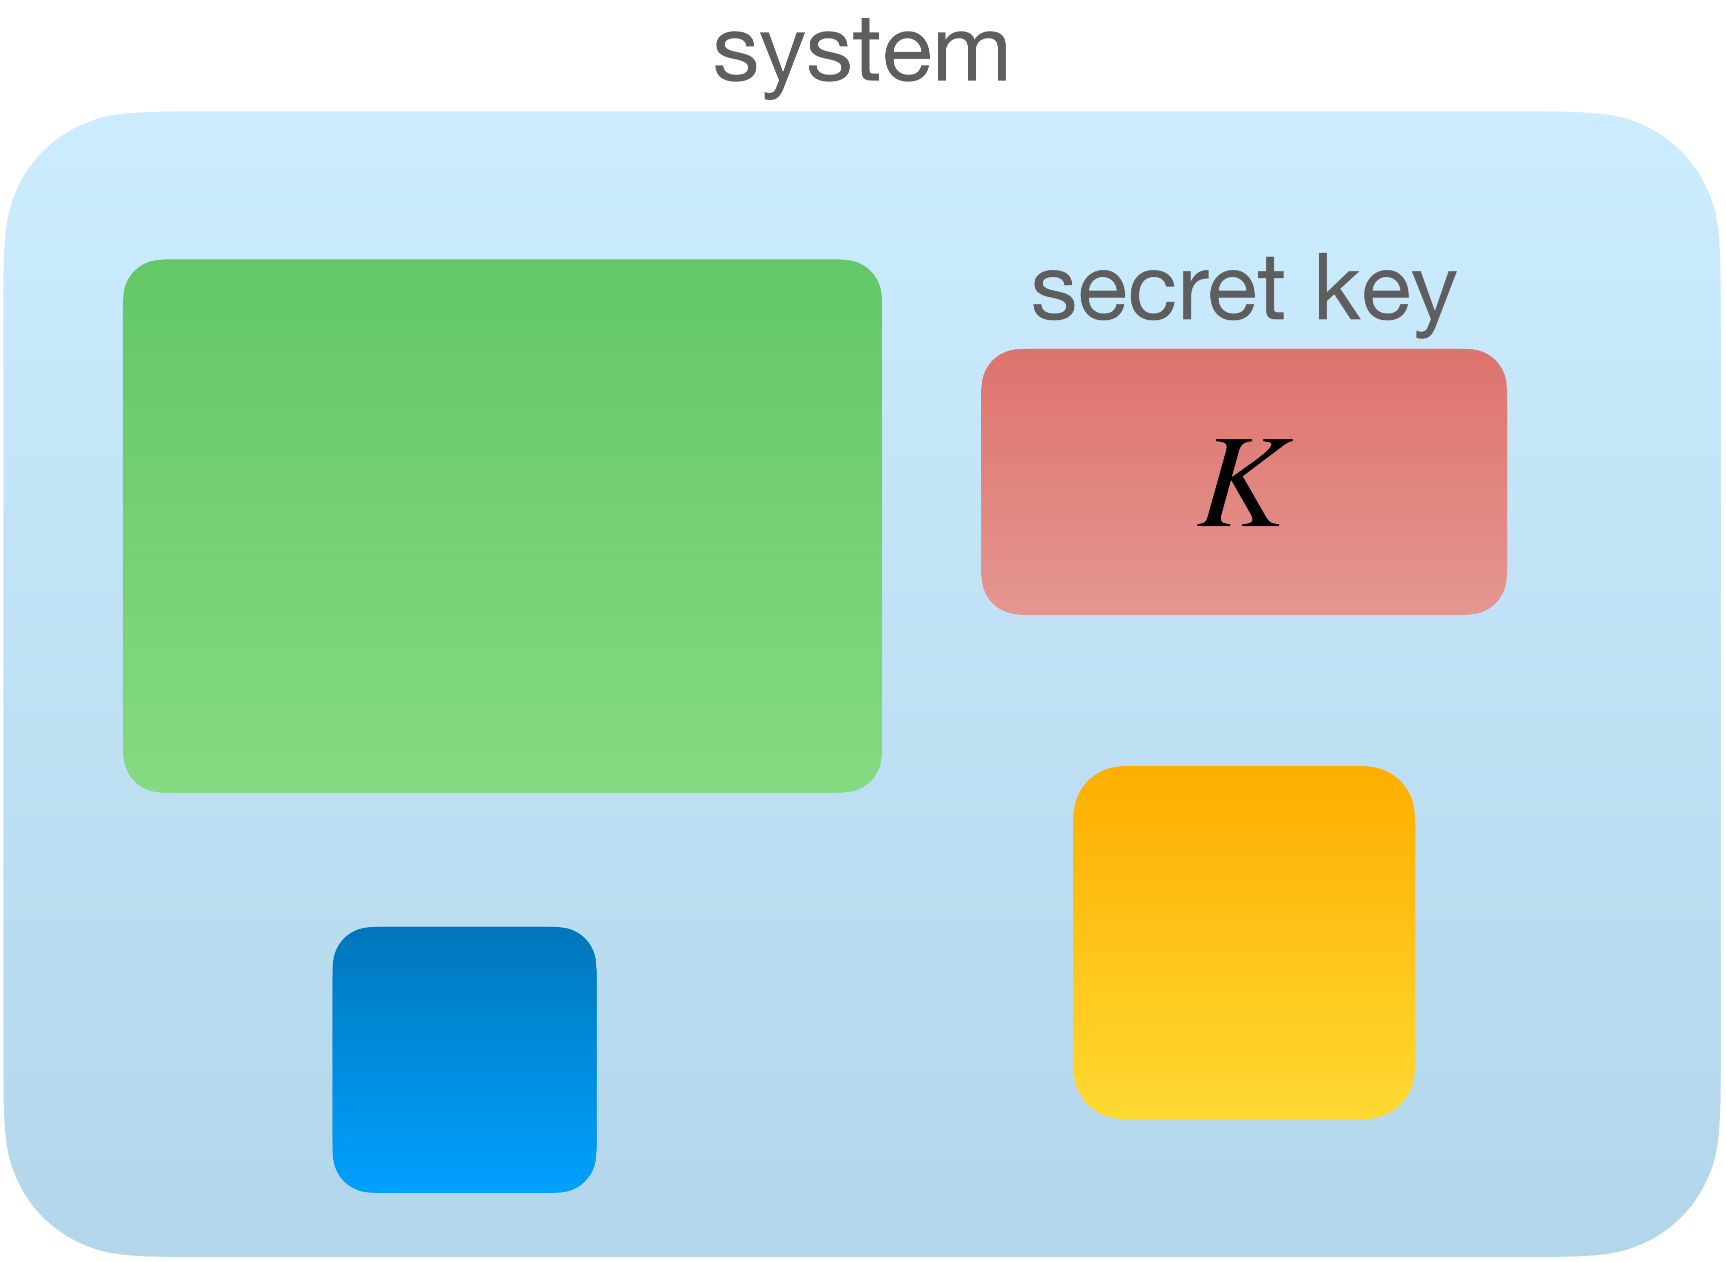

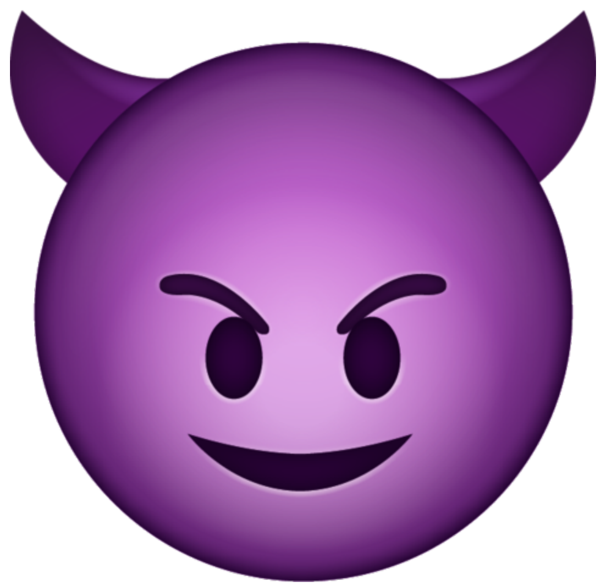

Side-channel attacks

Attacks on cryptographic schemes exploiting physical hardware quirks.

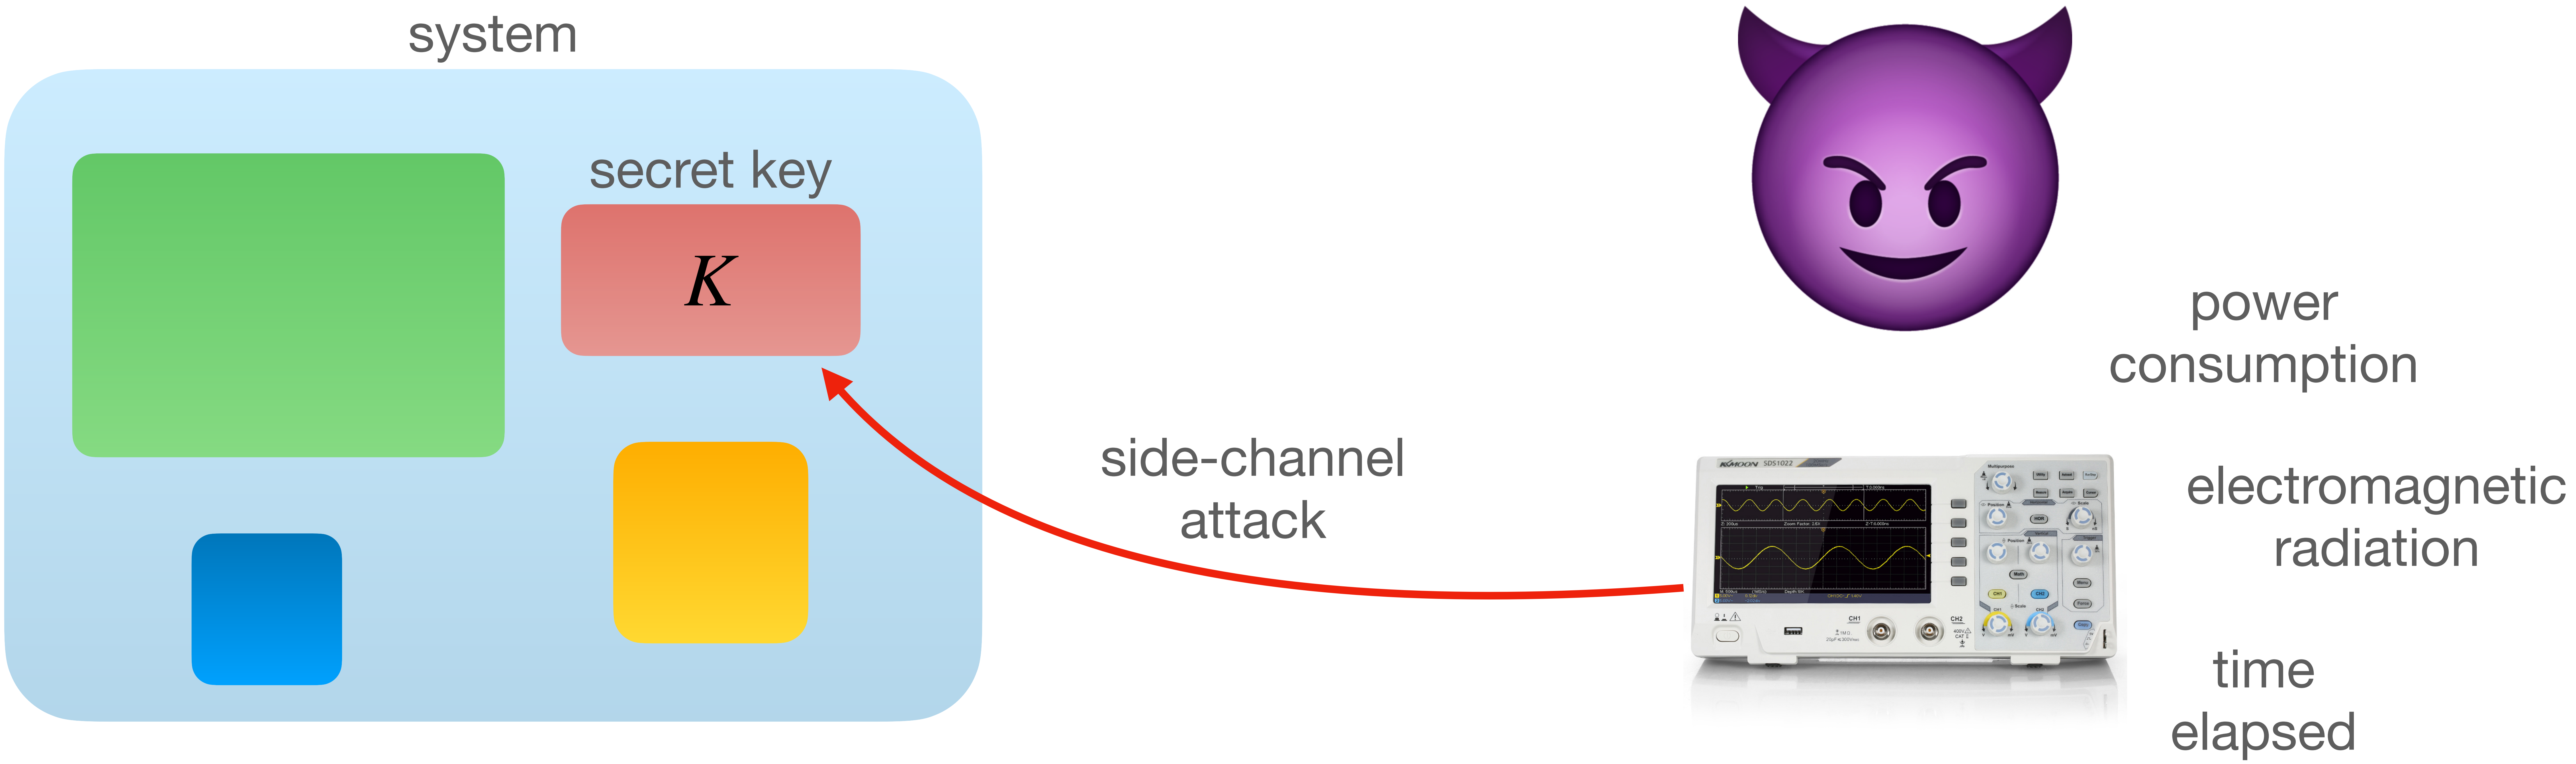

Side-channel attacks

Attacks on cryptographic schemes exploiting physical hardware quirks.

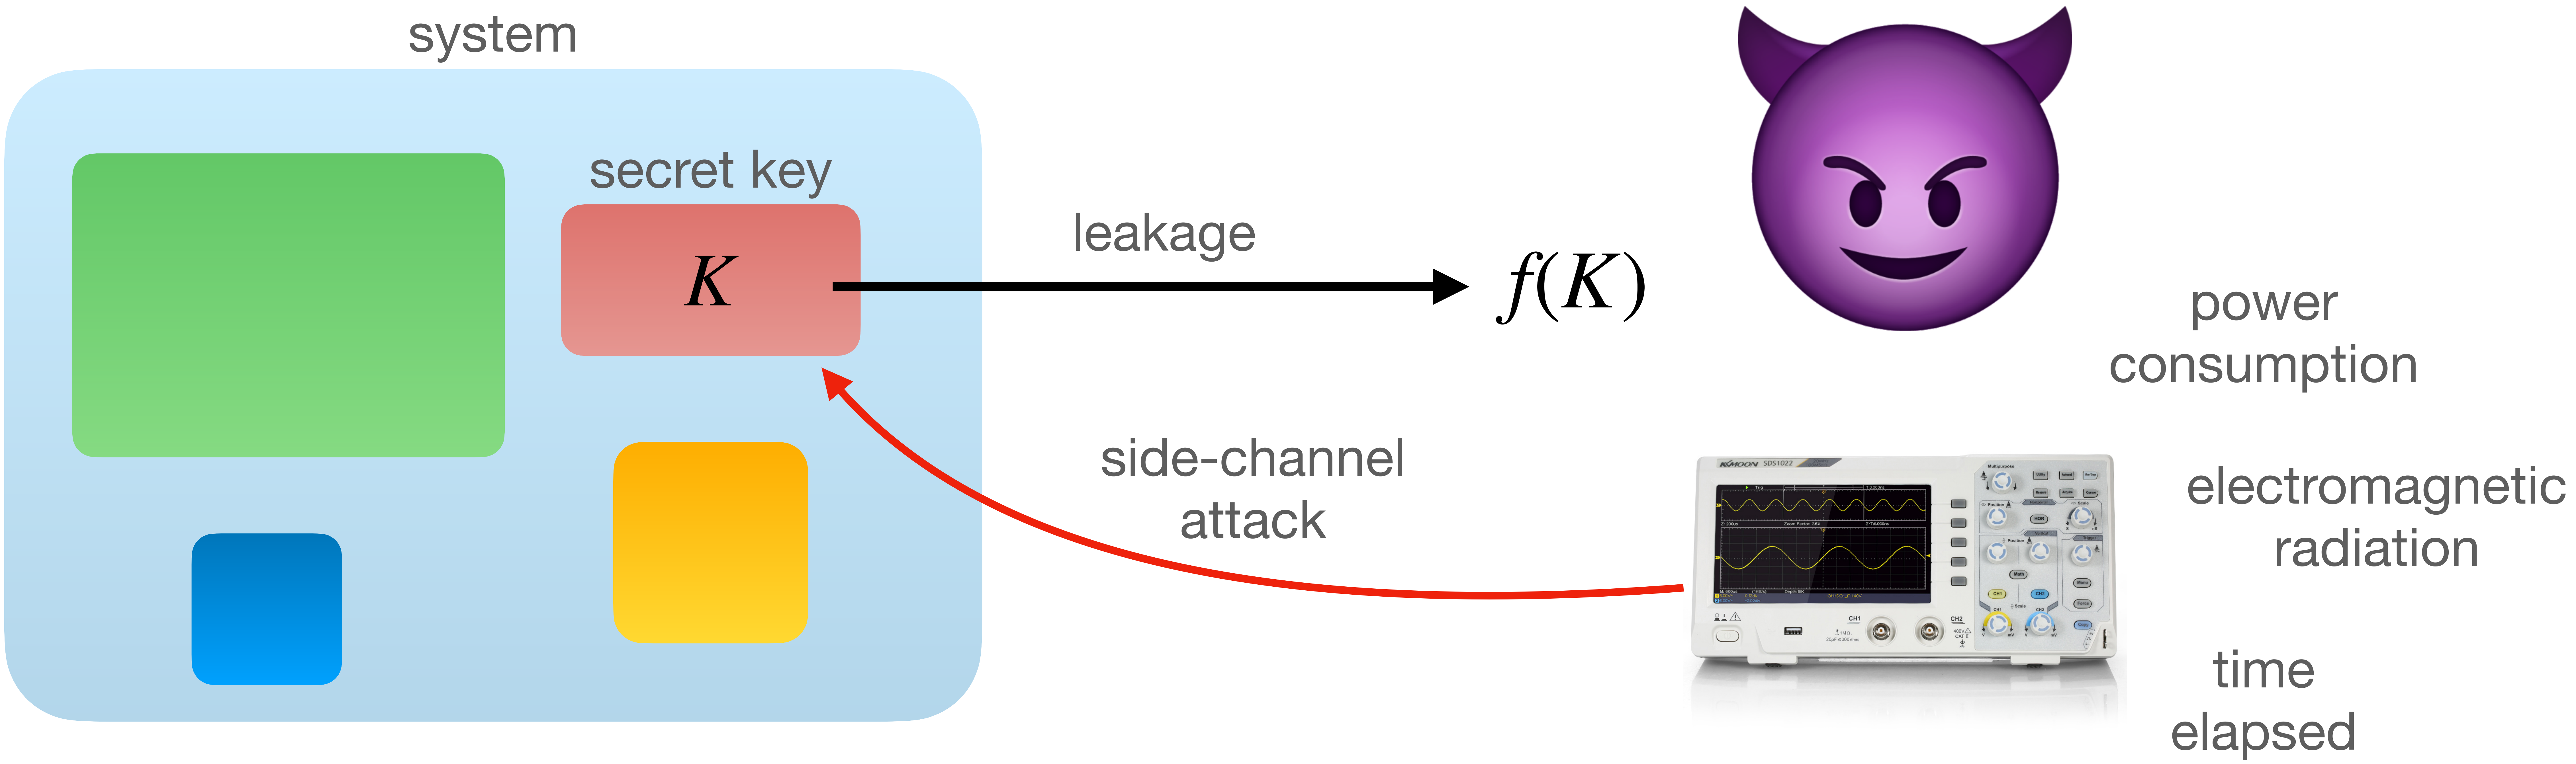

Side-channel attacks

Attacks on cryptographic schemes exploiting physical hardware quirks.

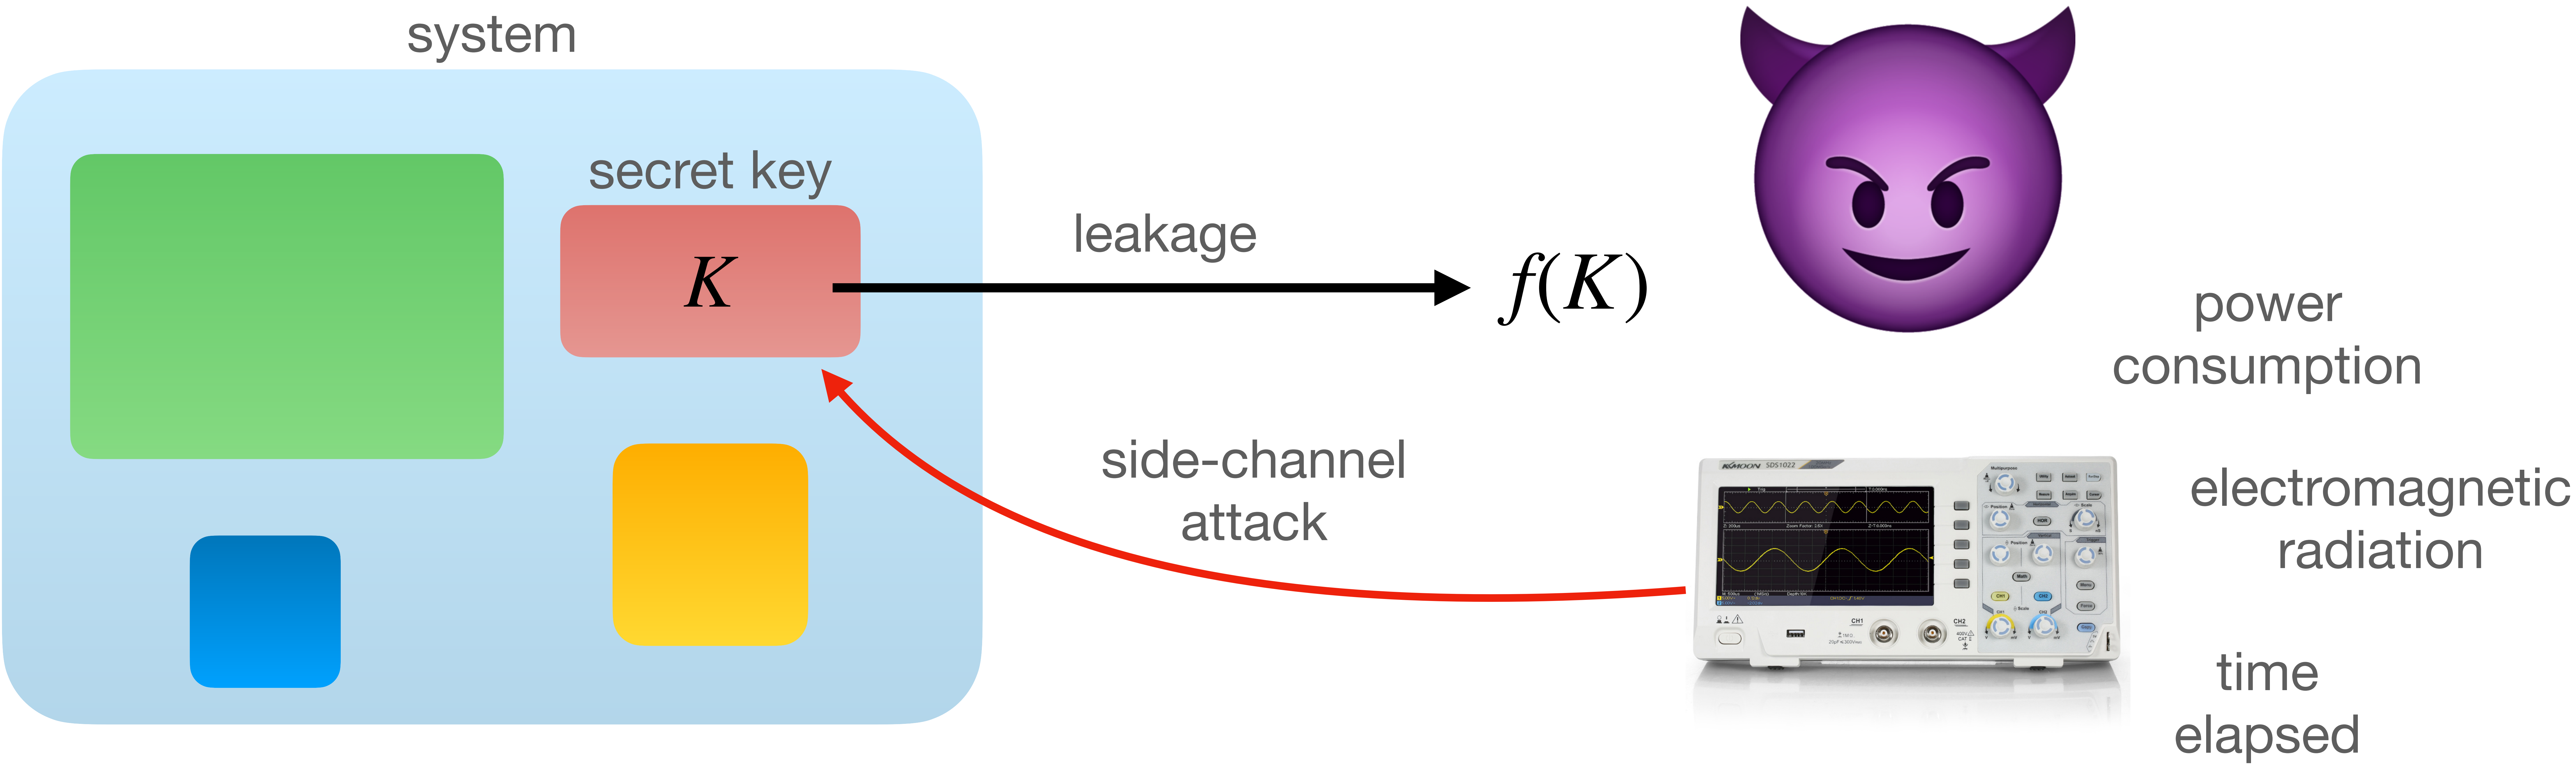

Leakage-resilience: System should remain secure even when adversary is able to mount a wide class of side-channel attacks.

Side-channel attacks can be cheap!

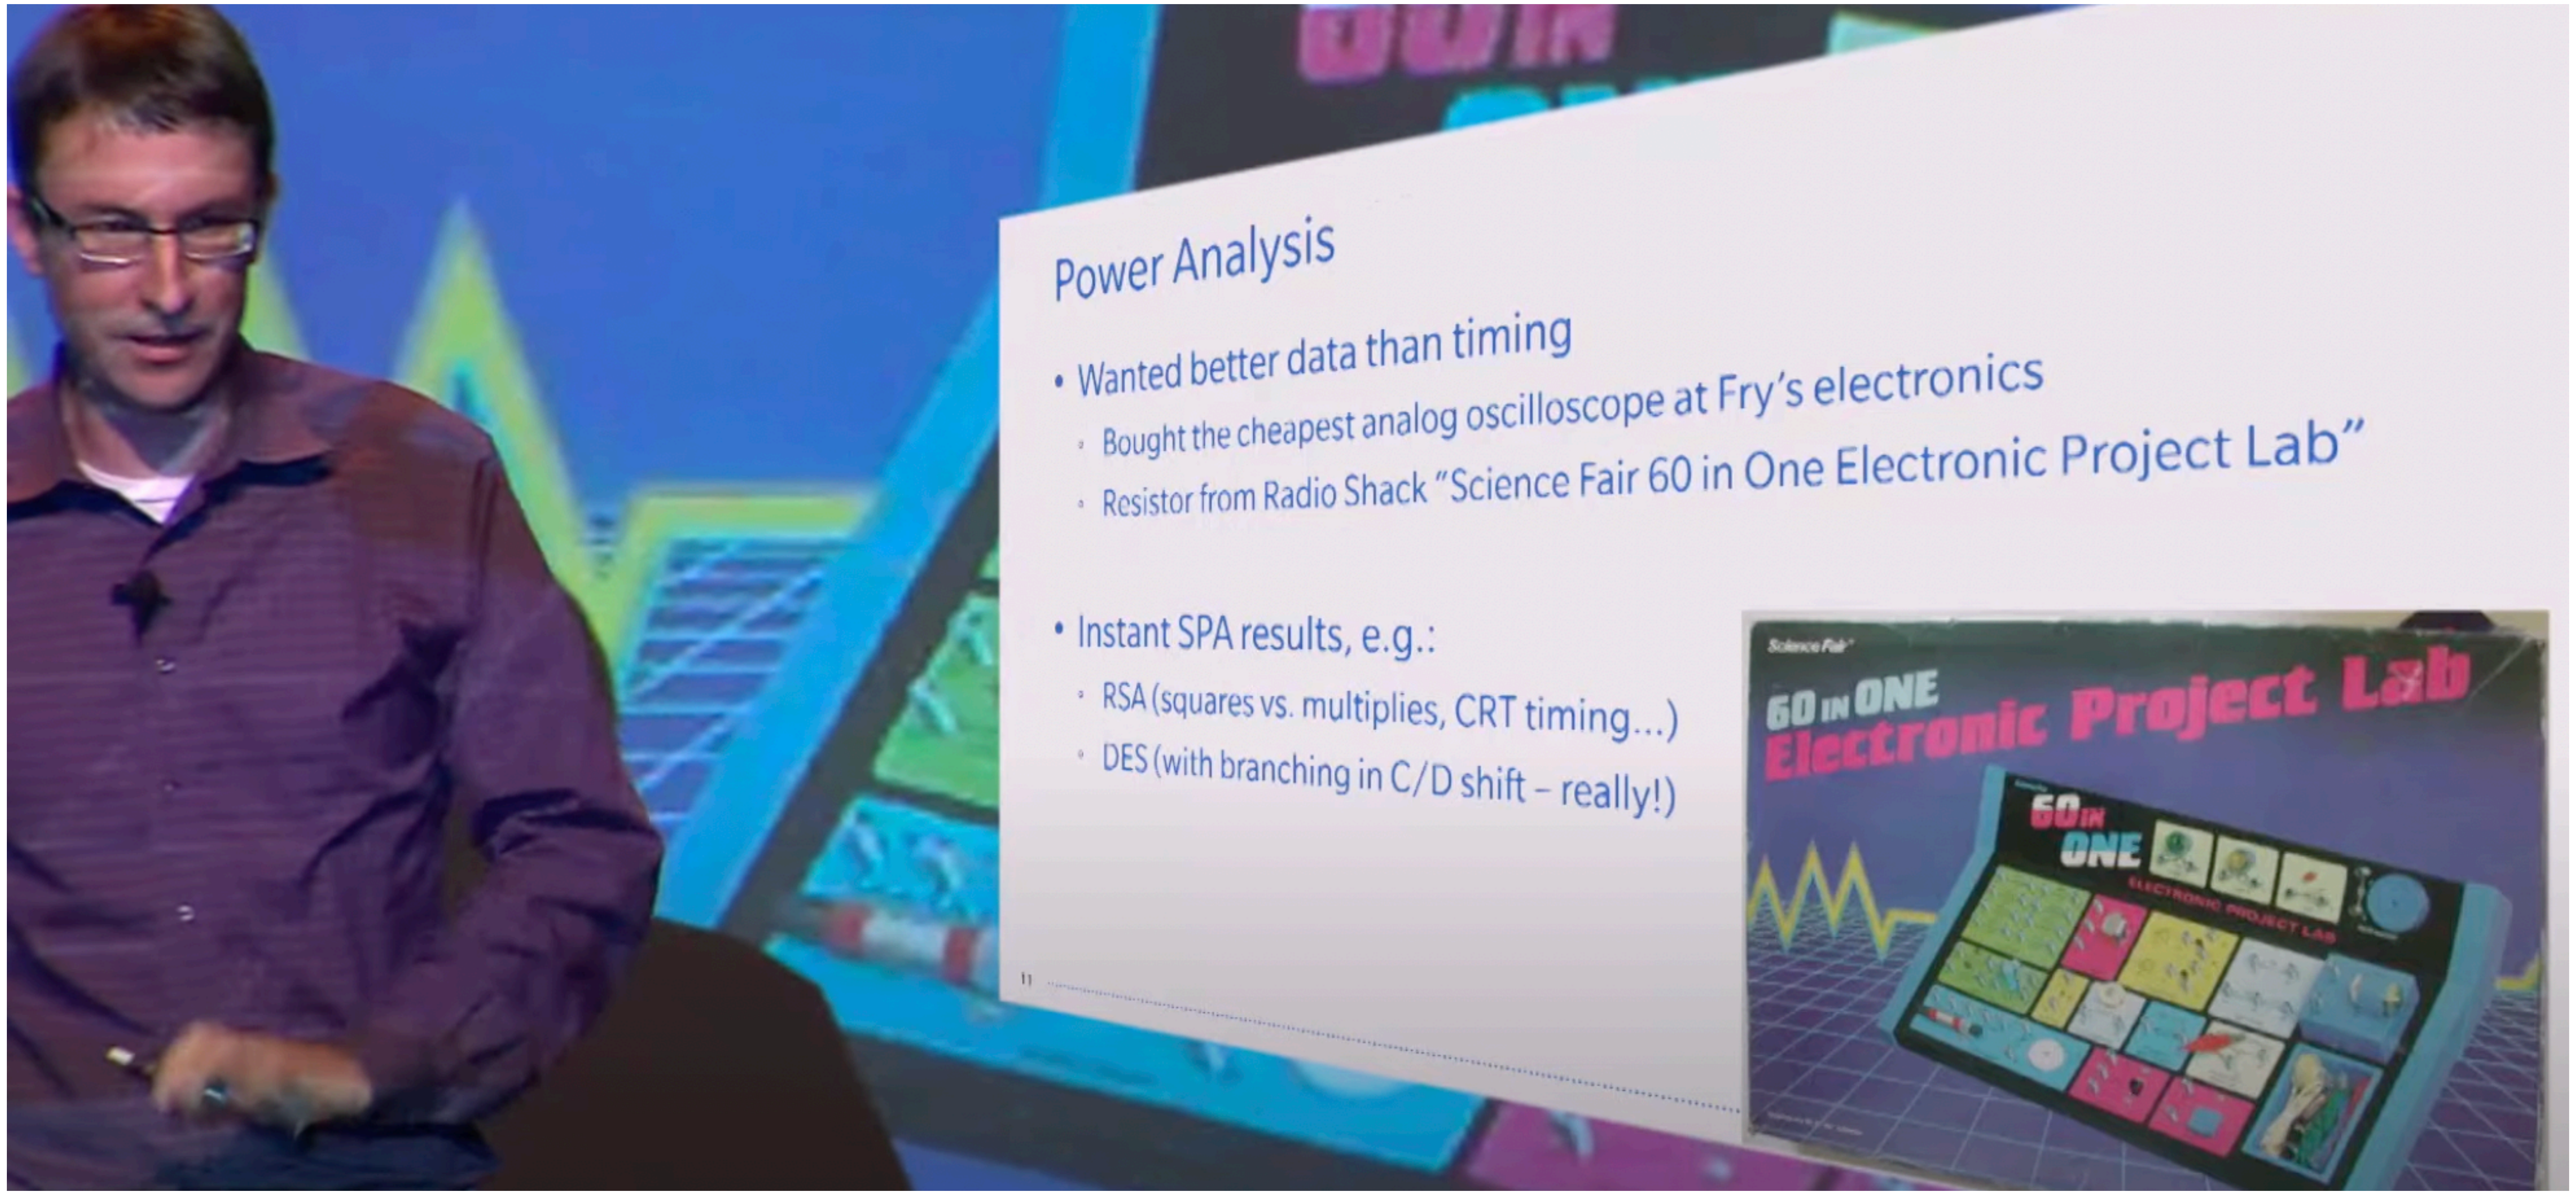

Paul Kocher — Obvious in hindsight: From side-channel attacks to the security challenges ahead
Invited talk at CRYPTO/CHES 2016
https://www.youtube.com/watch?v=6lt7ExN6Kw4

Bounded leakage

The most studied leakage model in theoretical cryptography.

Bounded leakage

The most studied leakage model in theoretical cryptography.

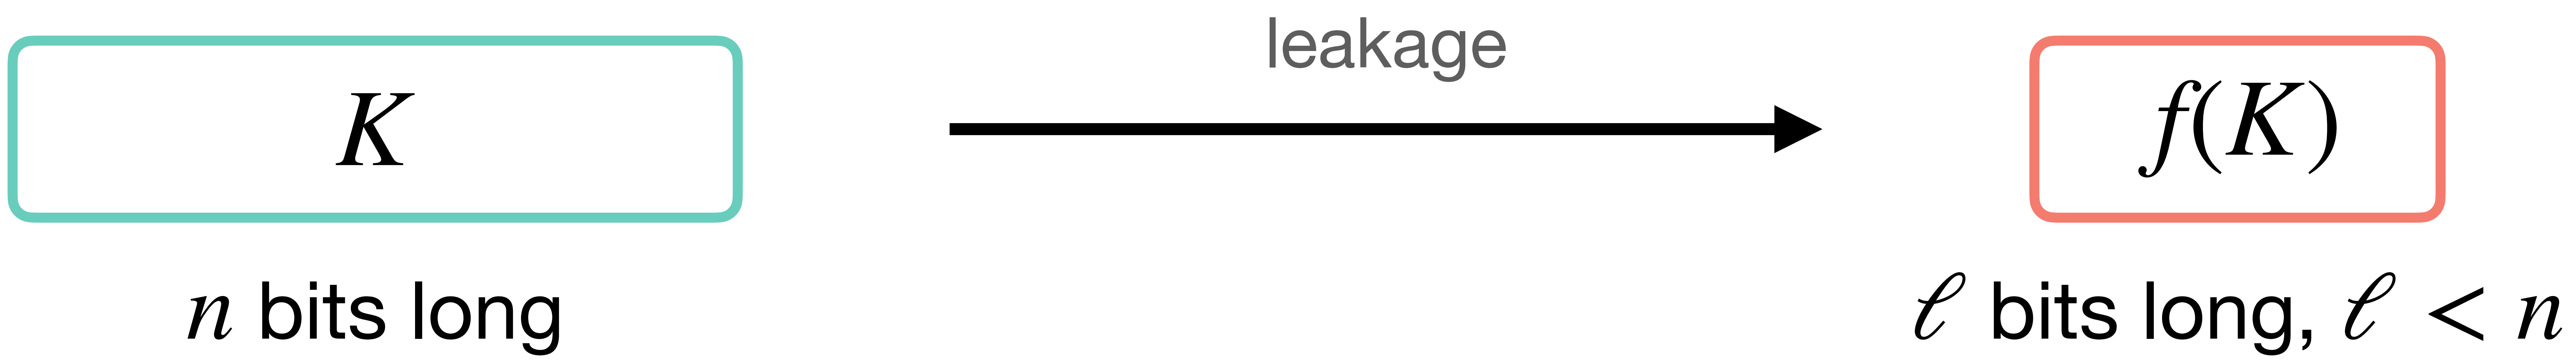

Example: keylength $n = 512$ bits, leakage length $\ell = 256$ bits

f can be **any** function with 256-bit output!

Bounded leakage

The most studied leakage model in theoretical cryptography.

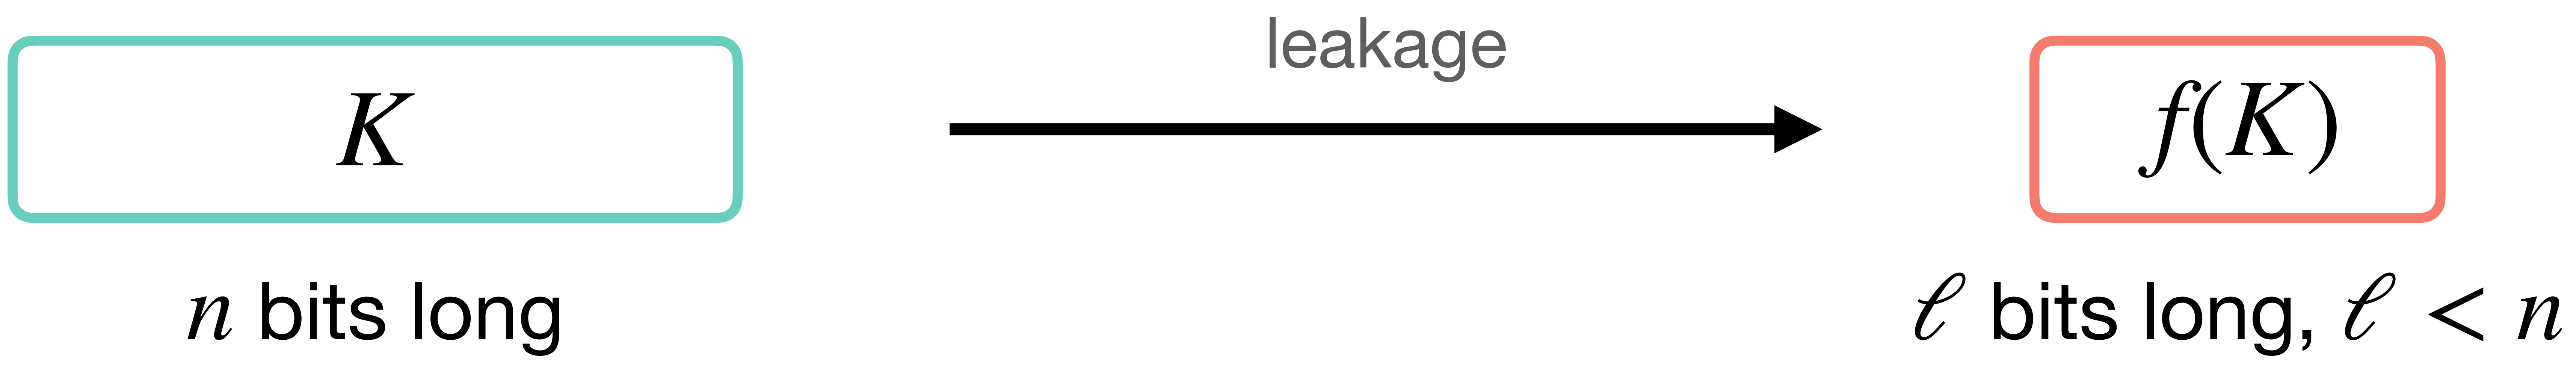

Example: keylength $n = 512$ bits, leakage length $\ell = 256$ bits

f can be **any** function with 256-bit output!

We know many cryptographic schemes with great “bounded leakage-resilience” guarantees.

Real-world leakage

Real-world side-channel attacks produce a lot of data, **but it is noisy!**

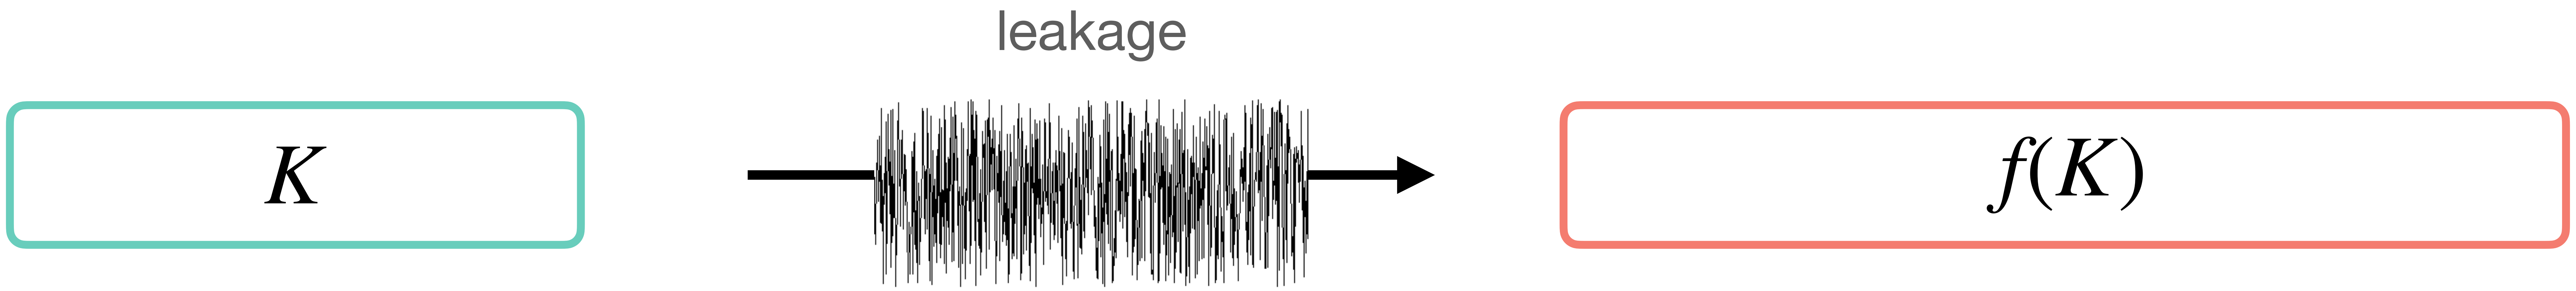

Real-world leakage

Real-world side-channel attacks produce a lot of data, **but it is noisy!**

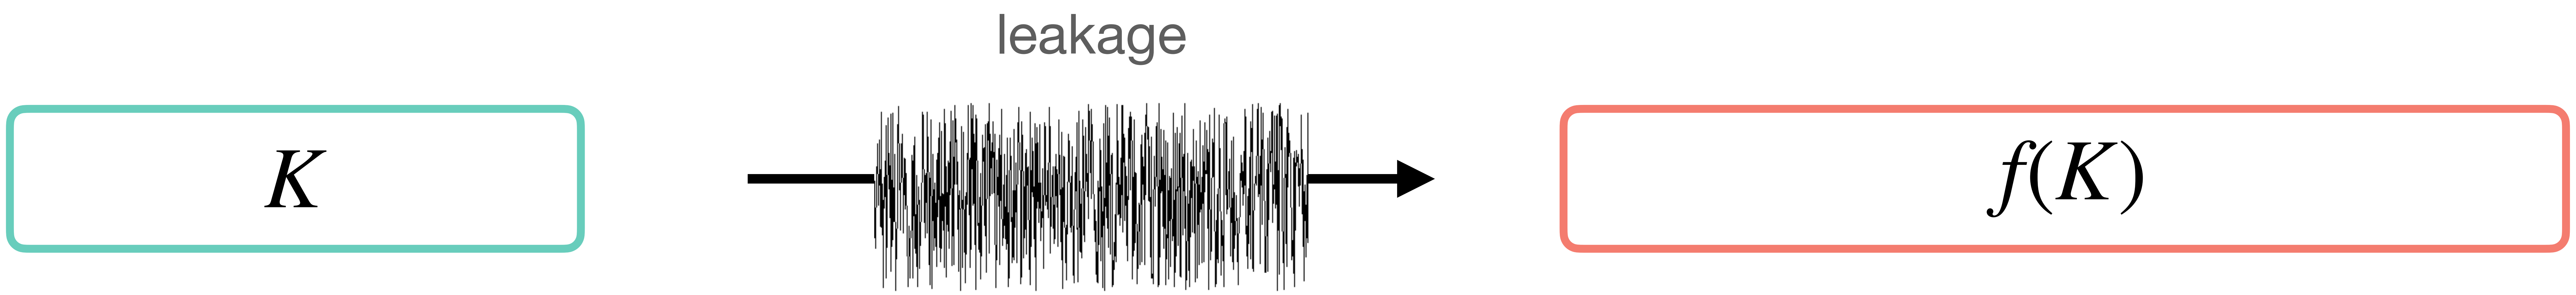

Several different measures of “noise” out there.

Real-world leakage

Real-world side-channel attacks produce a lot of data, **but it is noisy!**

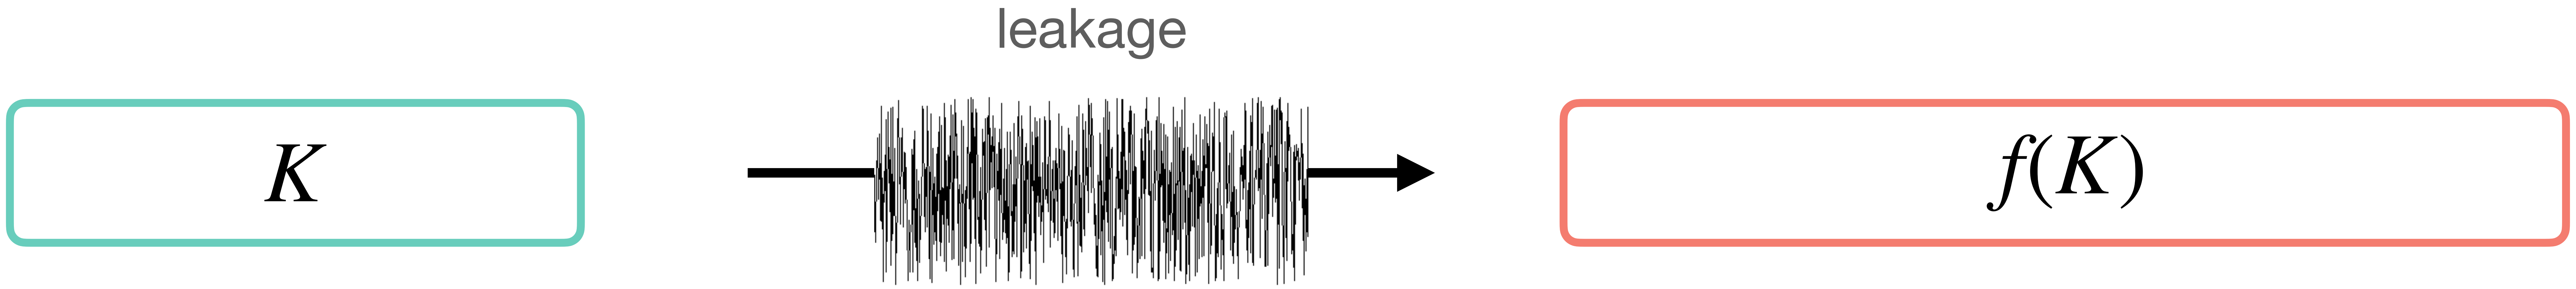

Several different measures of “noise” out there.

Popular noise measure: mutual information between K and $f(K)$.

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

We would like to find a noisy leakage model that:

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

We would like to find a noisy leakage model that:

1. Can be “simulated” effectively by bounded leakage.

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

We would like to find a noisy leakage model that:

1. Can be “simulated” effectively by bounded leakage.
2. Captures real-world side-channel attacks with good parameters.

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

We would like to find a noisy leakage model that:

1. Can be "simulated" effectively by bounded leakage.
2. Captures real-world side-channel attacks with good parameters.
    A. Practitioner infers leakage distributions induced by attacks on specific device;

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

We would like to find a noisy leakage model that:

1. Can be “simulated” effectively by bounded leakage.
2. Captures real-world side-channel attacks with good parameters.
    A. Practitioner infers leakage distributions induced by attacks on specific device;
    B. Can check if leakage falls into noisy leakage model (hopefully often the case!);

THE question

Does resilience to bounded leakage attacks imply non-trivial resilience to real world side-channel attacks?

Open-ended, depends on the noisy leakage model.

We would like to find a noisy leakage model that:

1. Can be “simulated” effectively by bounded leakage.
2. Captures real-world side-channel attacks with good parameters.
   A. Practitioner infers leakage distributions induced by attacks on specific device;
   B. Can check if leakage falls into noisy leakage model (hopefully often the case!);
   C. Readily derives useful concrete security guarantees.

The leakage simulation paradigm

Secret X, randomized leakage $Z = f(X)$

Ideal world

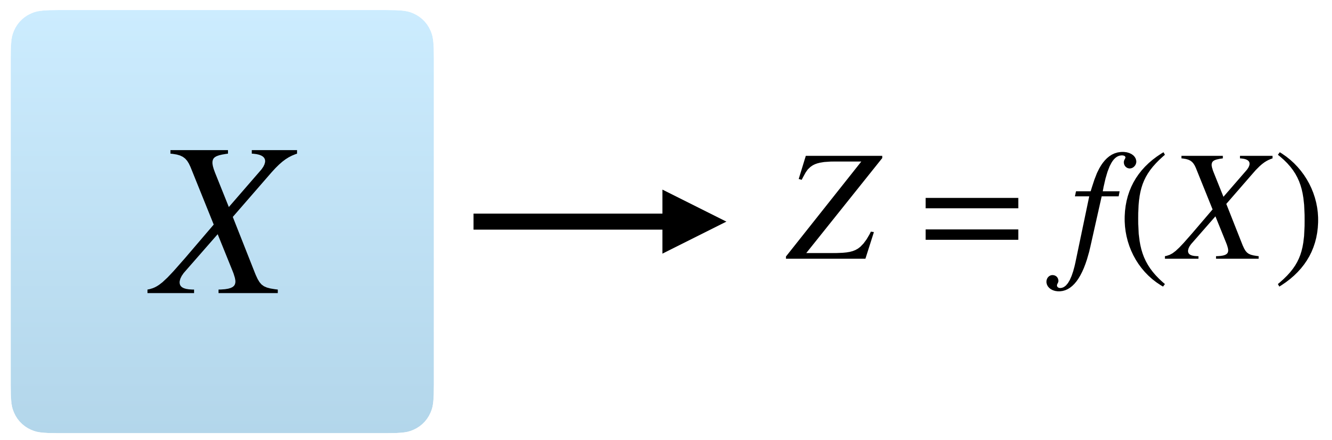

Real world

The leakage simulation paradigm

Secret X, randomized leakage $Z = f(X)$

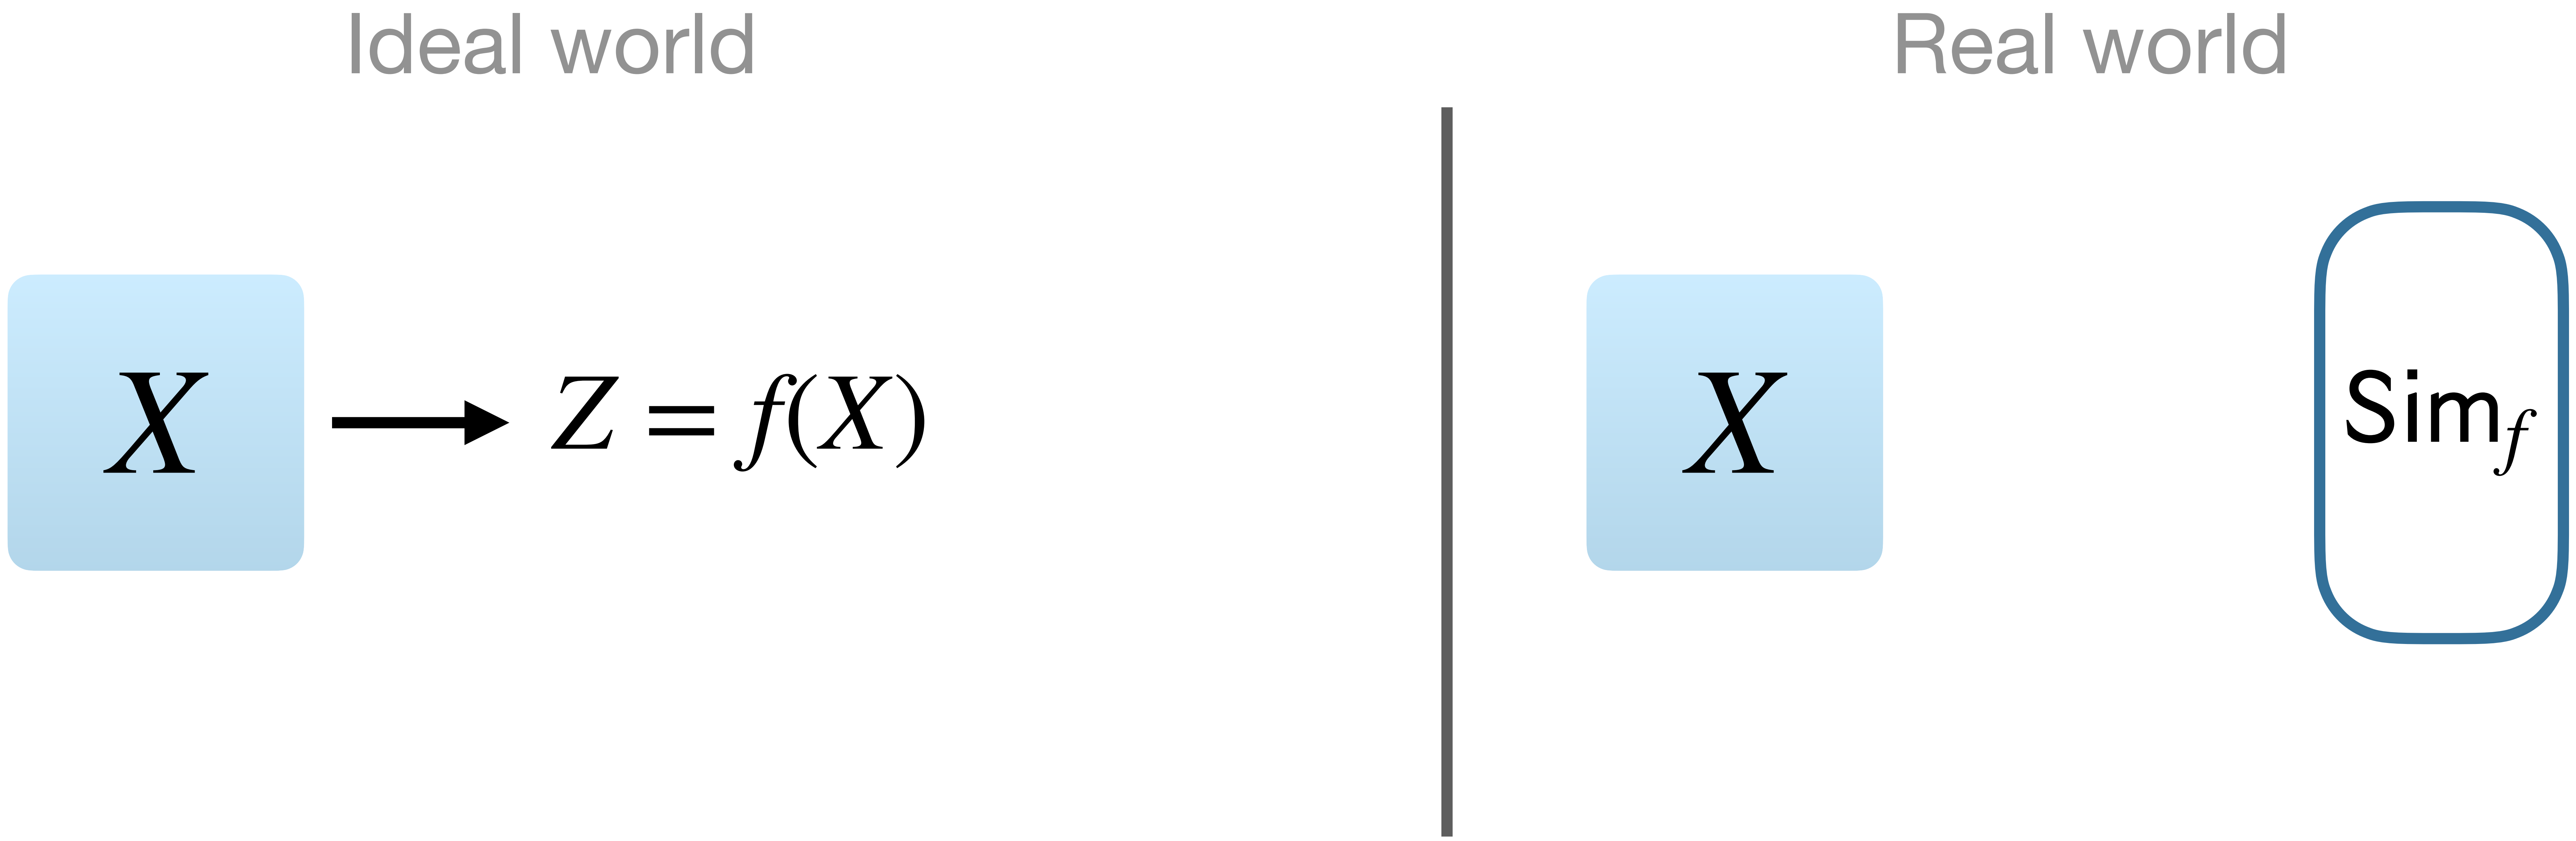

The leakage simulation paradigm

Secret X, randomized leakage $Z = f(X)$

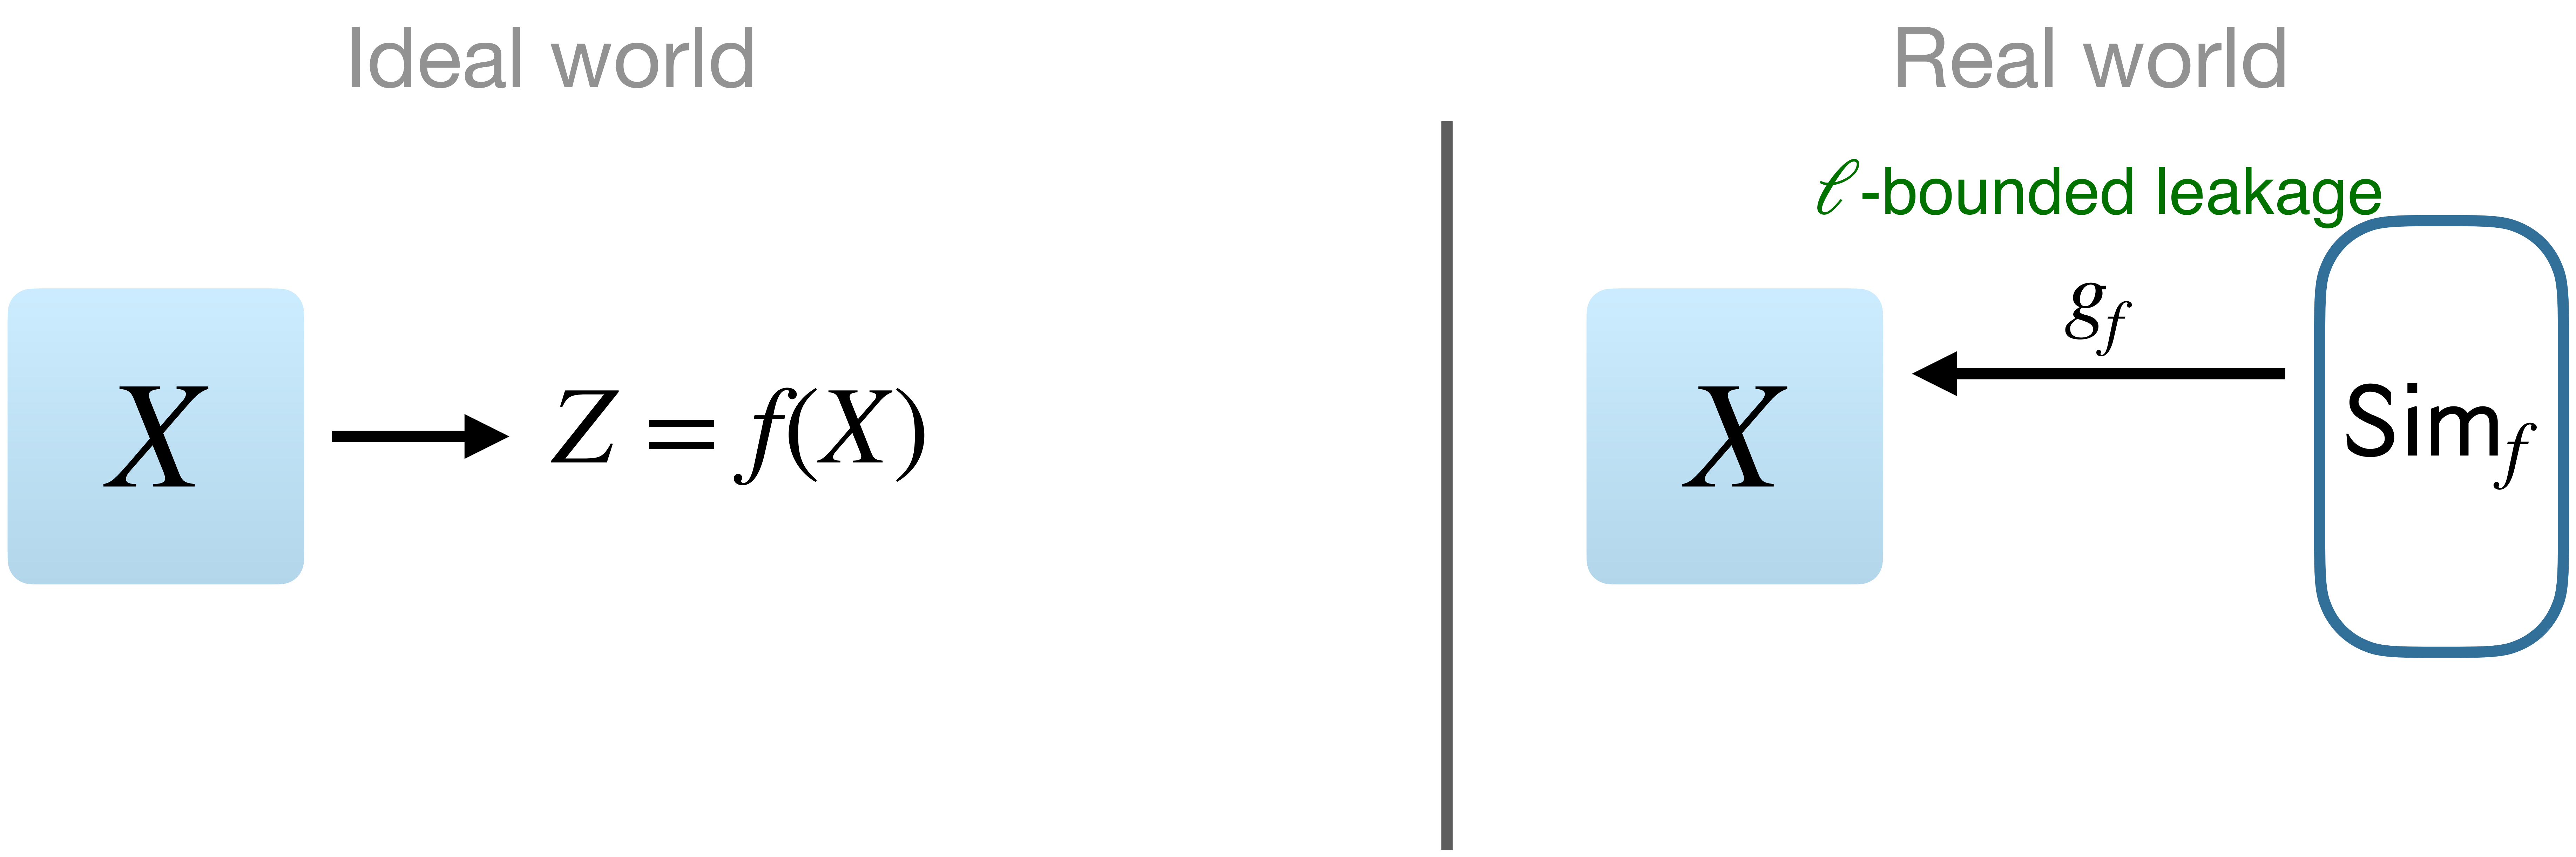

The leakage simulation paradigm

Secret X, randomized leakage $Z = f(X)$

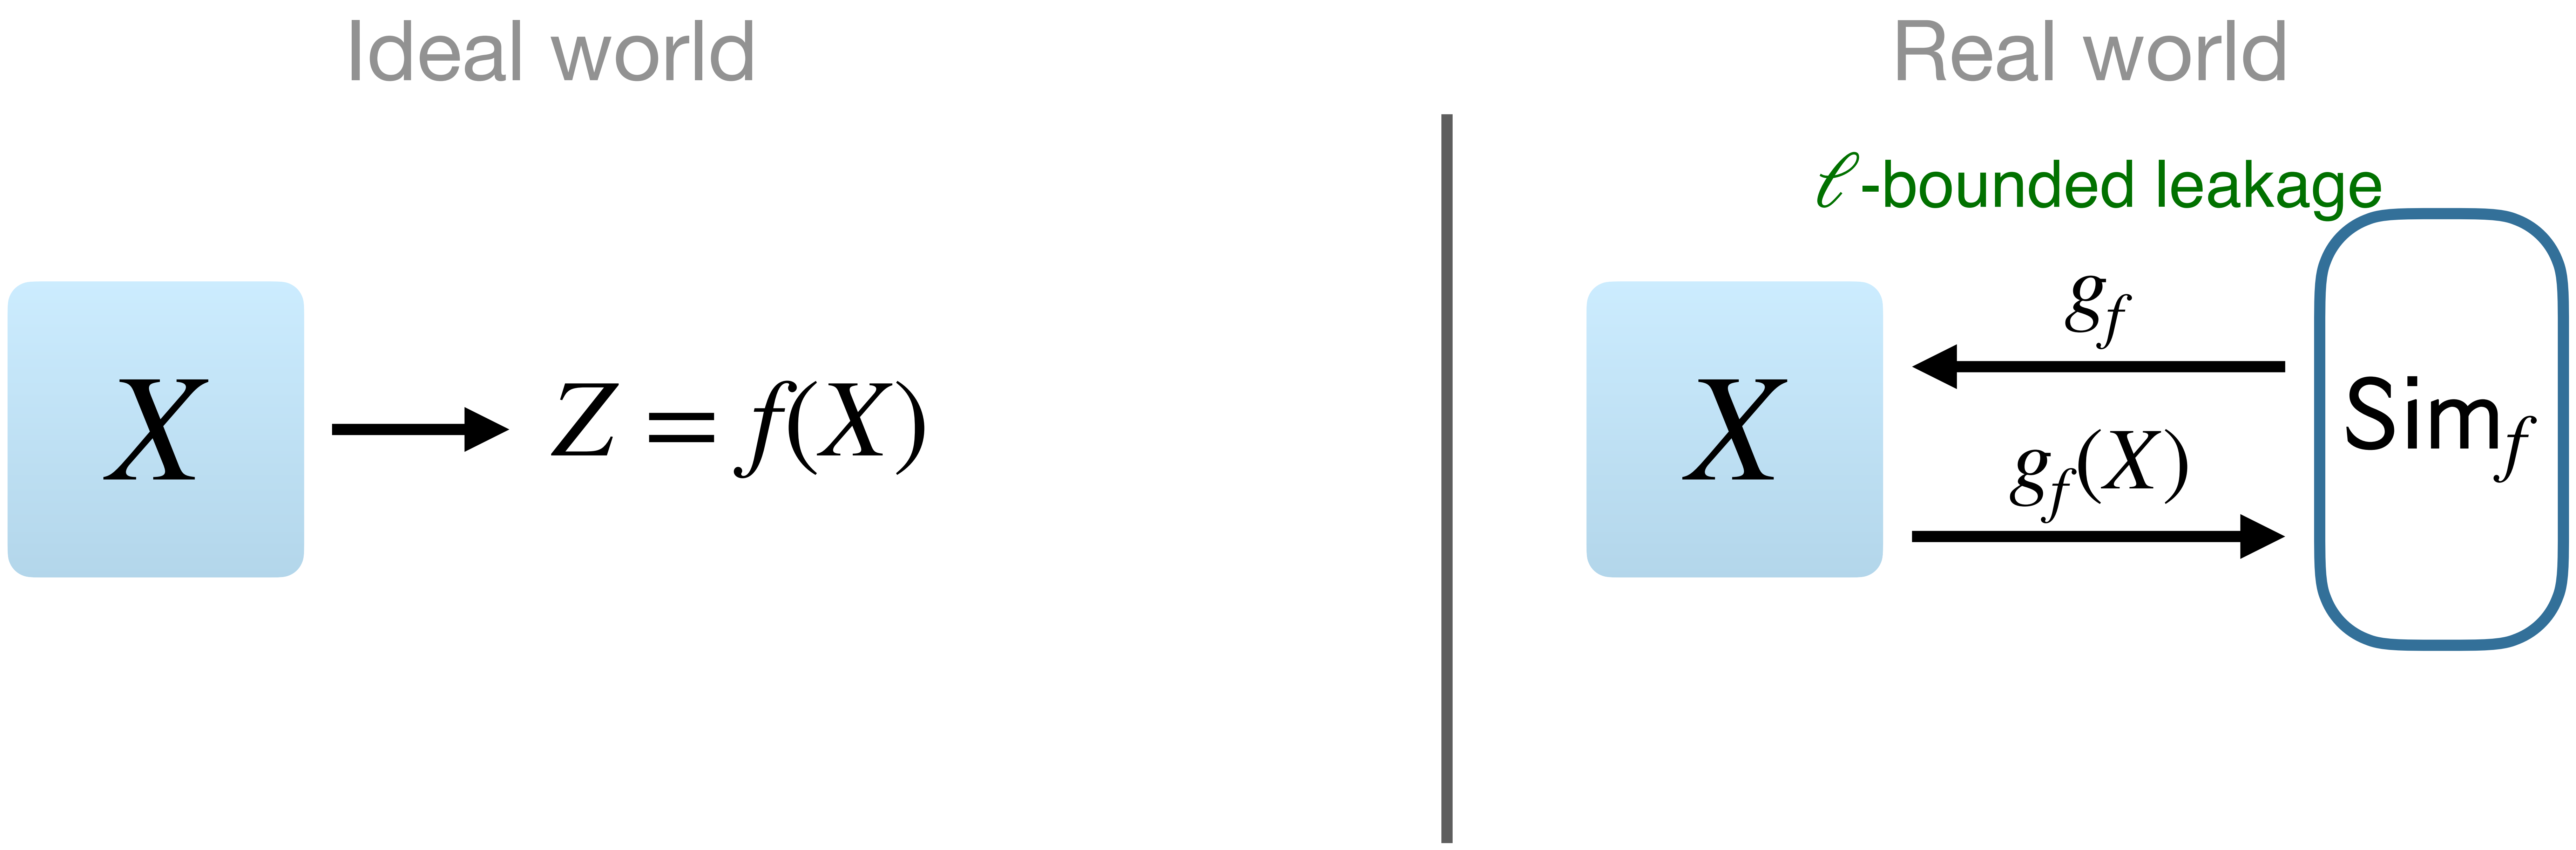

The leakage simulation paradigm

Secret X, randomized leakage $Z = f(X)$

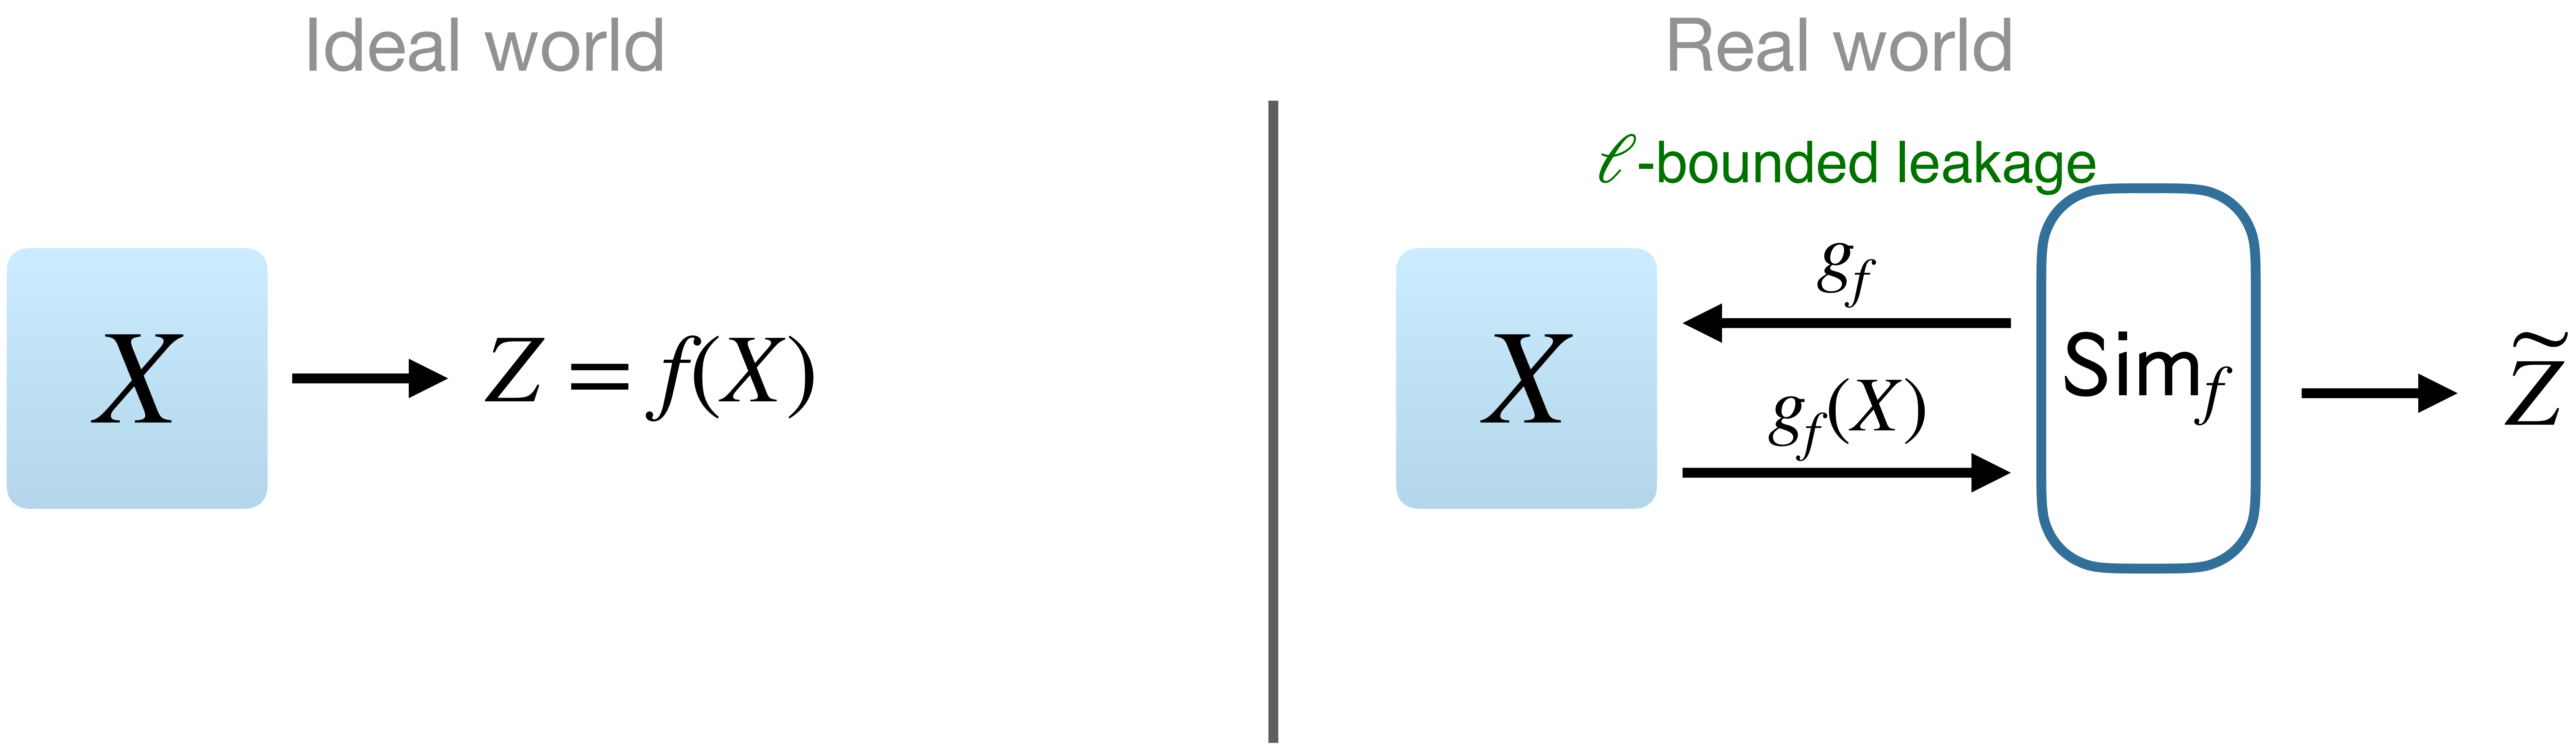

The leakage simulation paradigm

Secret X, randomized leakage $Z = f(X)$

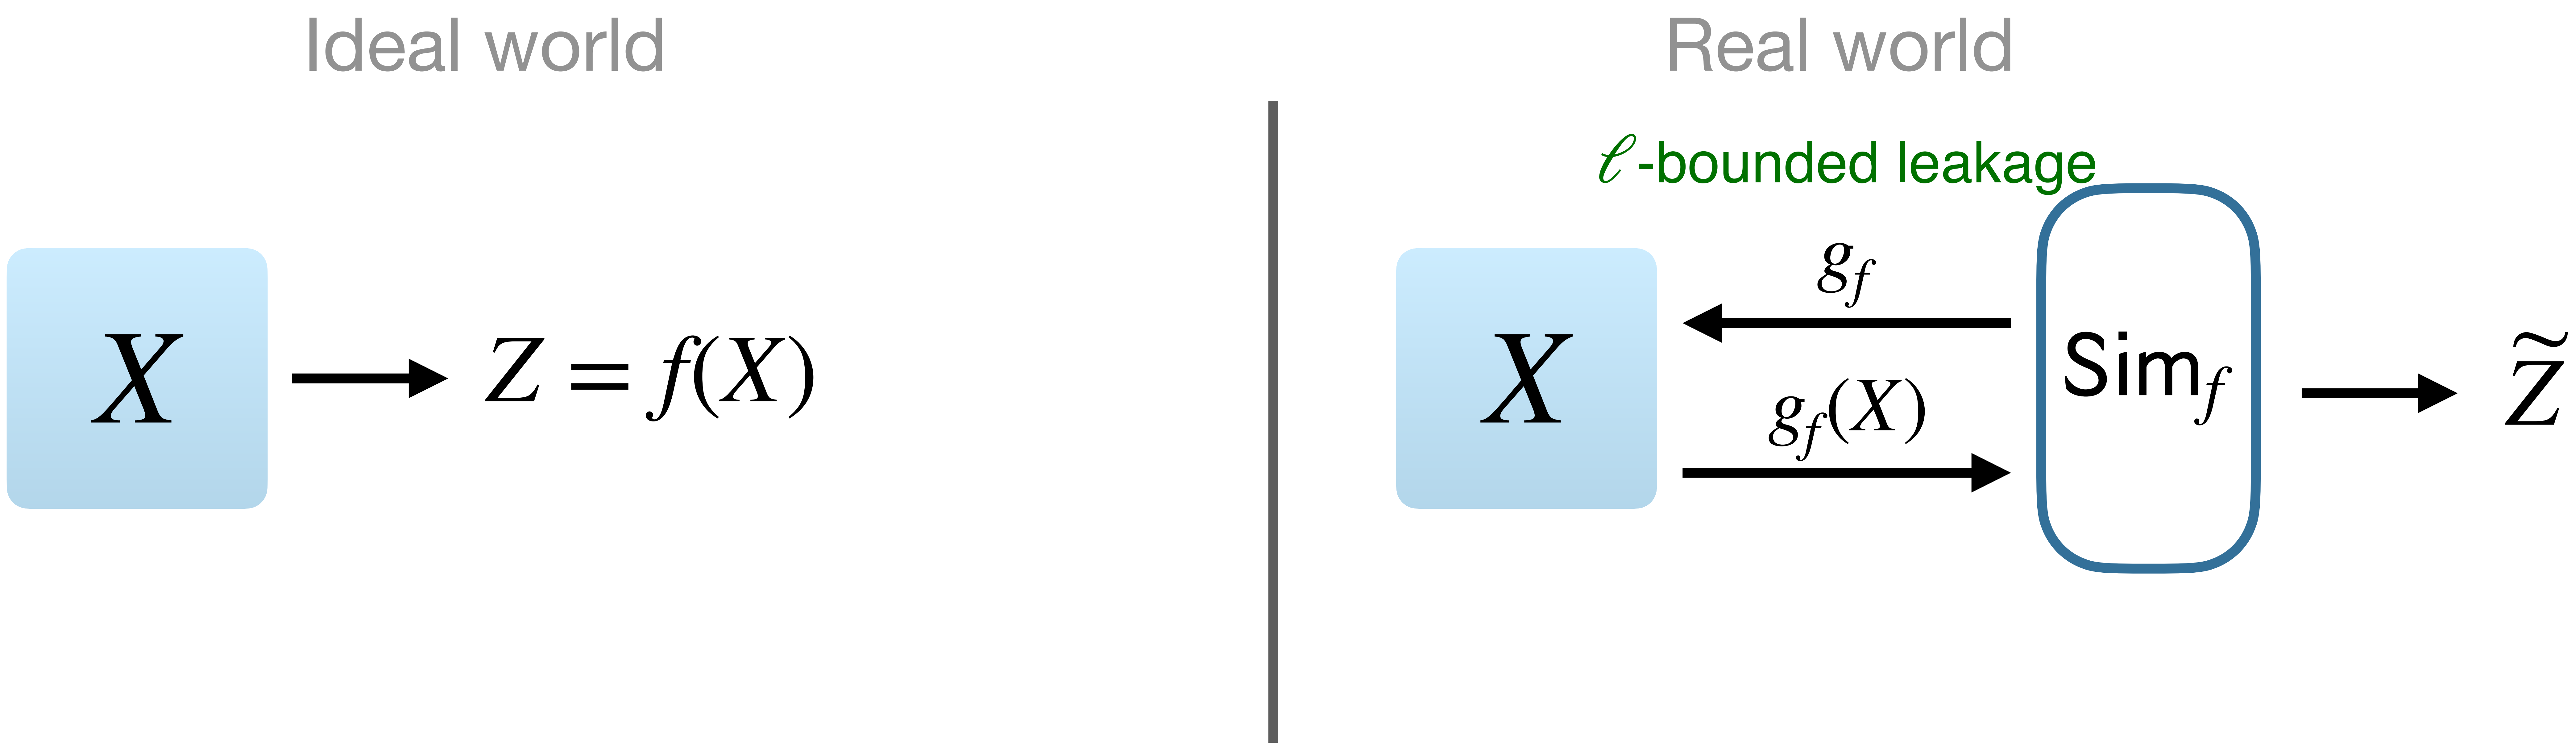

ε-simulation of Z by ℓ-bounded leakage: $\mathrm{SD}(P_{XZ}\,;P_{X\widetilde{Z}}) \leq \varepsilon$

statistical distance

Mutual information vs. Bounded leakage

Q: Can we simulate all leakages Z with low $I(X;Z)$ using a small amount of bounded leakage?

Mutual information vs. Bounded leakage

Q: Can we simulate all leakages Z with low $I(X;Z)$ using a small amount of bounded leakage?

A: Nope.

Mutual information vs. Bounded leakage

Q: Can we simulate all leakages Z with low $I(X;Z)$ using a small amount of bounded leakage?

A: Nope.

$$X \text{ uniform over } \{0,1\}^n \text{ and } Z = \begin{cases} X, & \text{with prob. } \delta, \\ \perp, & \text{with prob. } 1-\delta. \end{cases}$$

Mutual information vs. Bounded leakage

Q: Can we simulate all leakages Z with low $I(X;Z)$ using a small amount of bounded leakage?

A: Nope.

$$X \text{ uniform over } \{0,1\}^n \text{ and } Z = \begin{cases} X, & \text{with prob. } \delta, \\ \perp, & \text{with prob. } 1-\delta. \end{cases}$$

- $I(X;Z) = \delta n$.

Mutual information vs. Bounded leakage

Q: Can we simulate all leakages Z with low $I(X;Z)$ using a small amount of bounded leakage?

A: Nope.

$$X \text{ uniform over } \{0,1\}^n \text{ and } Z = \begin{cases} X, & \text{with prob. } \delta, \\ \perp, & \text{with prob. } 1-\delta. \end{cases}$$

- $I(X;Z) = \delta n$.
- Can't simulate with error below $\delta/2$, even with $n-1$ bits of bounded leakage.

Mutual information vs. Bounded leakage

Q: Can we simulate all leakages Z with low $I(X;Z)$ using a small amount of bounded leakage?

A: Nope.

$$X \text{ uniform over } \{0,1\}^n \text{ and } Z = \begin{cases} X, & \text{with prob. } \delta, \\ \perp, & \text{with prob. } 1-\delta. \end{cases}$$

- $I(X;Z) = \delta n$.
- Can't simulate with error below $\delta/2$, even with $n-1$ bits of bounded leakage.

"Low mutual information" is too loose, need to come up with a different measure.

Coming up with another noise measure

(RECAP) We would like to find a noisy leakage model that:

1. Can be simulated using a small amount of bounded leakage.
2. Captures real-world side-channel attacks with good parameters.

Coming up with another noise measure

(RECAP) We would like to find a noisy leakage model that:

1. Can be simulated using a small amount of bounded leakage.
2. Captures real-world side-channel attacks with good parameters.

We'll do it backwards…

(1) come up with a nice simulator, (2) reverse-engineer the noise measure.

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;
- **for** $i = 1, \ldots, L$**:**

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;
- **for** $i = 1, \ldots, L$**:**
  - With probability $\dfrac{P(z_i)}{T \cdot Q(z_i)}$, return z_i;

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;
- **for** $i = 1, \ldots, L$**:**
  - With probability $\frac{P(z_i)}{T \cdot Q(z_i)}$, return z_i;
- return $\perp$;

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;
- **for** $i = 1,\ldots,L$**:**
  - With probability $\frac{P(z_i)}{T \cdot Q(z_i)}$, return z_i;
- return $\perp$;

- Conditioned on output $\neq \perp$, output is distributed according to P

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;
- **for** $i = 1, \ldots, L$**:**
  - With probability $\frac{P(z_i)}{T \cdot Q(z_i)}$, return z_i;
- return $\perp$;

- Conditioned on output $\neq \perp$, output is distributed according to P
- $\Pr[\text{output} = \perp] = (1 - 1/T)^L \leq e^{-L/T}$

Rejection Sampling 101

Setting: Want to sample from P, but only have access to i.i.d. samples from Q and Bernoulli random variables.

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from Q;
- **for** $i = 1, \ldots, L$**:**
  - With probability $\frac{P(z_i)}{T \cdot Q(z_i)}$, return z_i;
- return $\perp$;

- Conditioned on output $\neq \perp$, output is distributed according to P
- $\Pr[\text{output} = \perp] = (1 - 1/T)^L \leq e^{-L/T}$
- **Need $P(z) \leq T \cdot Q(z)$ for all z**

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1,\ldots,L$**:**

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1,\ldots,L$**:**
    - With probability $\frac{P_{Z|X=x}(z_i)}{T \cdot P_Z(z_i)}$, return i;

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1,\ldots, L$**:**
    - With probability $\dfrac{P_{Z|X=x}(z_i)}{T \cdot P_Z(z_i)}$, return i;
  - return L;

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1, \ldots, L$**:**
    - With probability $\dfrac{P_{Z|X=x}(z_i)}{T \cdot P_Z(z_i)}$, return i;
  - return L;
- Output $z_{g_{\vec{z}}(X)}$ as the simulated leakage.

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1, \ldots, L$**:**
    - With probability $\dfrac{P_{Z|X=x}(z_i)}{T \cdot P_Z(z_i)}$, return i;
  - return L;
- Output $z_{g_{\vec{z}}(X)}$ as the simulated leakage.

- $g_{\vec{z}}$ has output length $\log L$

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1, \ldots, L$**:**
    - With probability $\dfrac{P_{Z|X=x}(z_i)}{T \cdot P_Z(z_i)}$, return i;
  - return L;
- Output $z_{g_{\vec{z}}(X)}$ as the simulated leakage.

- $g_{\vec{z}}$ has output length $\log L$
- simulation error = rej. samp. fails $\approx e^{-L/T}$

Simulation from bounded leakage via rejection sampling

Idea: Simulate true leakage distribution $P_{Z|X=x}$ by rejection sampling from marginal P_Z

- Sample $z_1, z_2, \ldots, z_L$ i.i.d. from P_Z
- "Rej. samp." leakage function $g_{\vec{z}}(x)$:
  - **for** $i = 1, \ldots, L$**:**
    - With probability $\dfrac{P_{Z|X=x}(z_i)}{T \cdot P_Z(z_i)}$, return i;
  - return L;
- Output $z_{g_{\vec{z}}(X)}$ as the simulated leakage.

- $g_{\vec{z}}$ has output length $\log L$
- simulation error = rej. samp. fails $\approx e^{-L/T}$
- Need $P_{Z|X=x}(z) \leq T \cdot P_Z(z)$ for **most** z

Which noisy leakages are good for rejection sampling?

Hockey-Stick Divergences (generalize statistical distance):

$$\mathrm{SD}_t(P; Q) \leq \delta \text{ if and only if } P(S) \leq 2^t \cdot Q(S) + \delta \text{ for all sets } S.$$

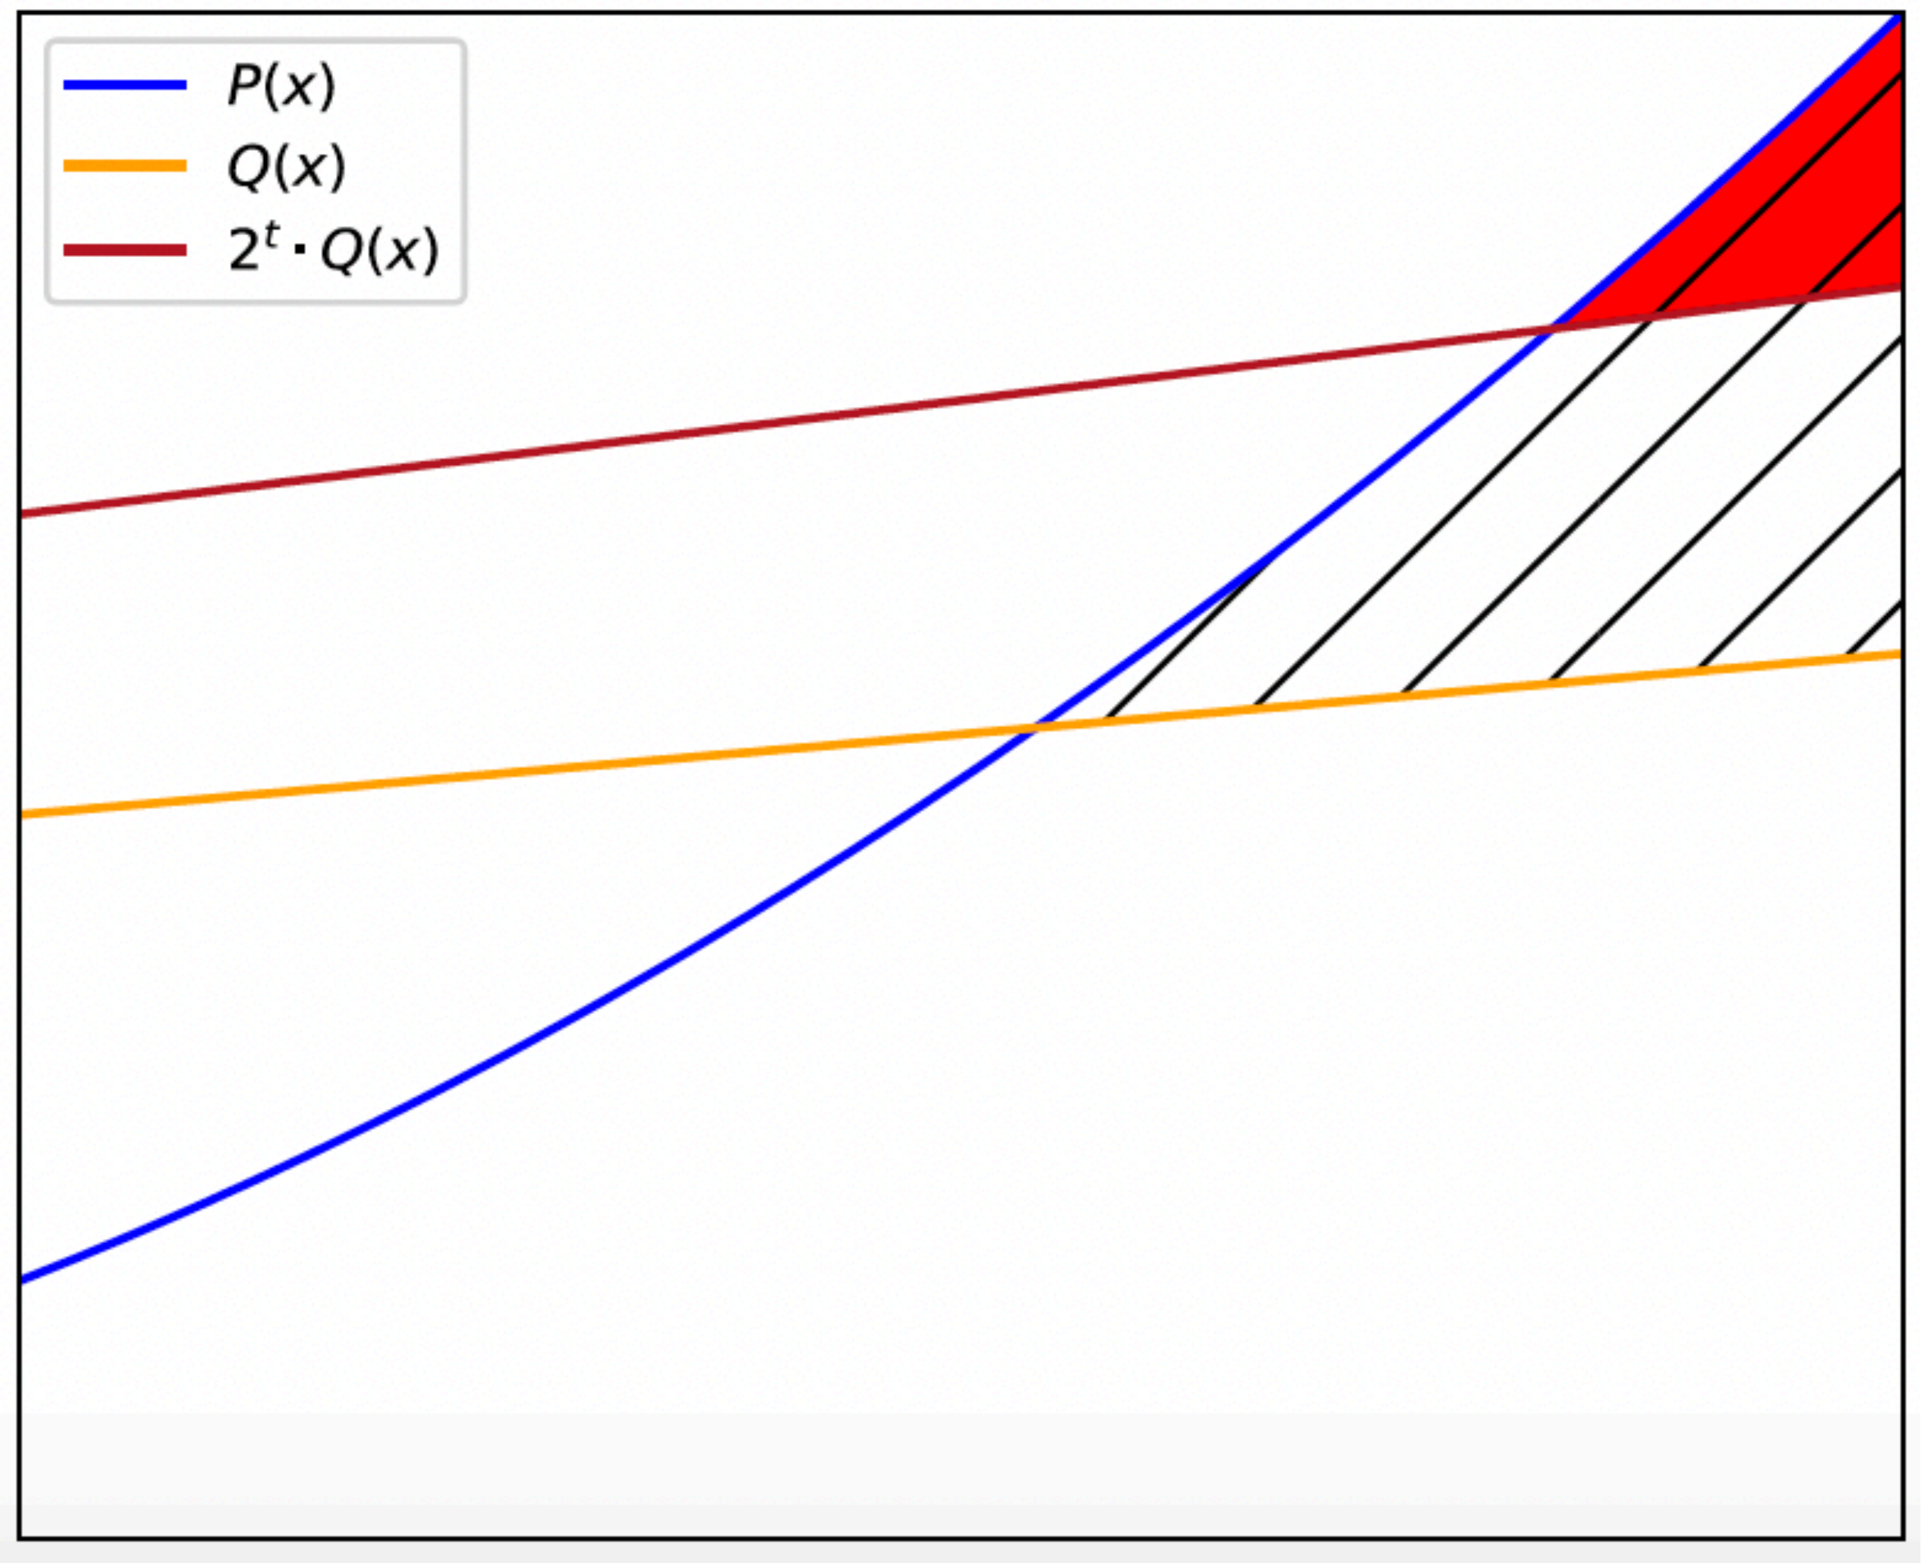

The (t, δ)-SD-noisy model

$Z = f(X)$ is (t, δ)-SD-noisy leakage from X if $\mathrm{SD}_t(P_{XZ}; P_X \otimes P_Z) \leq \delta$

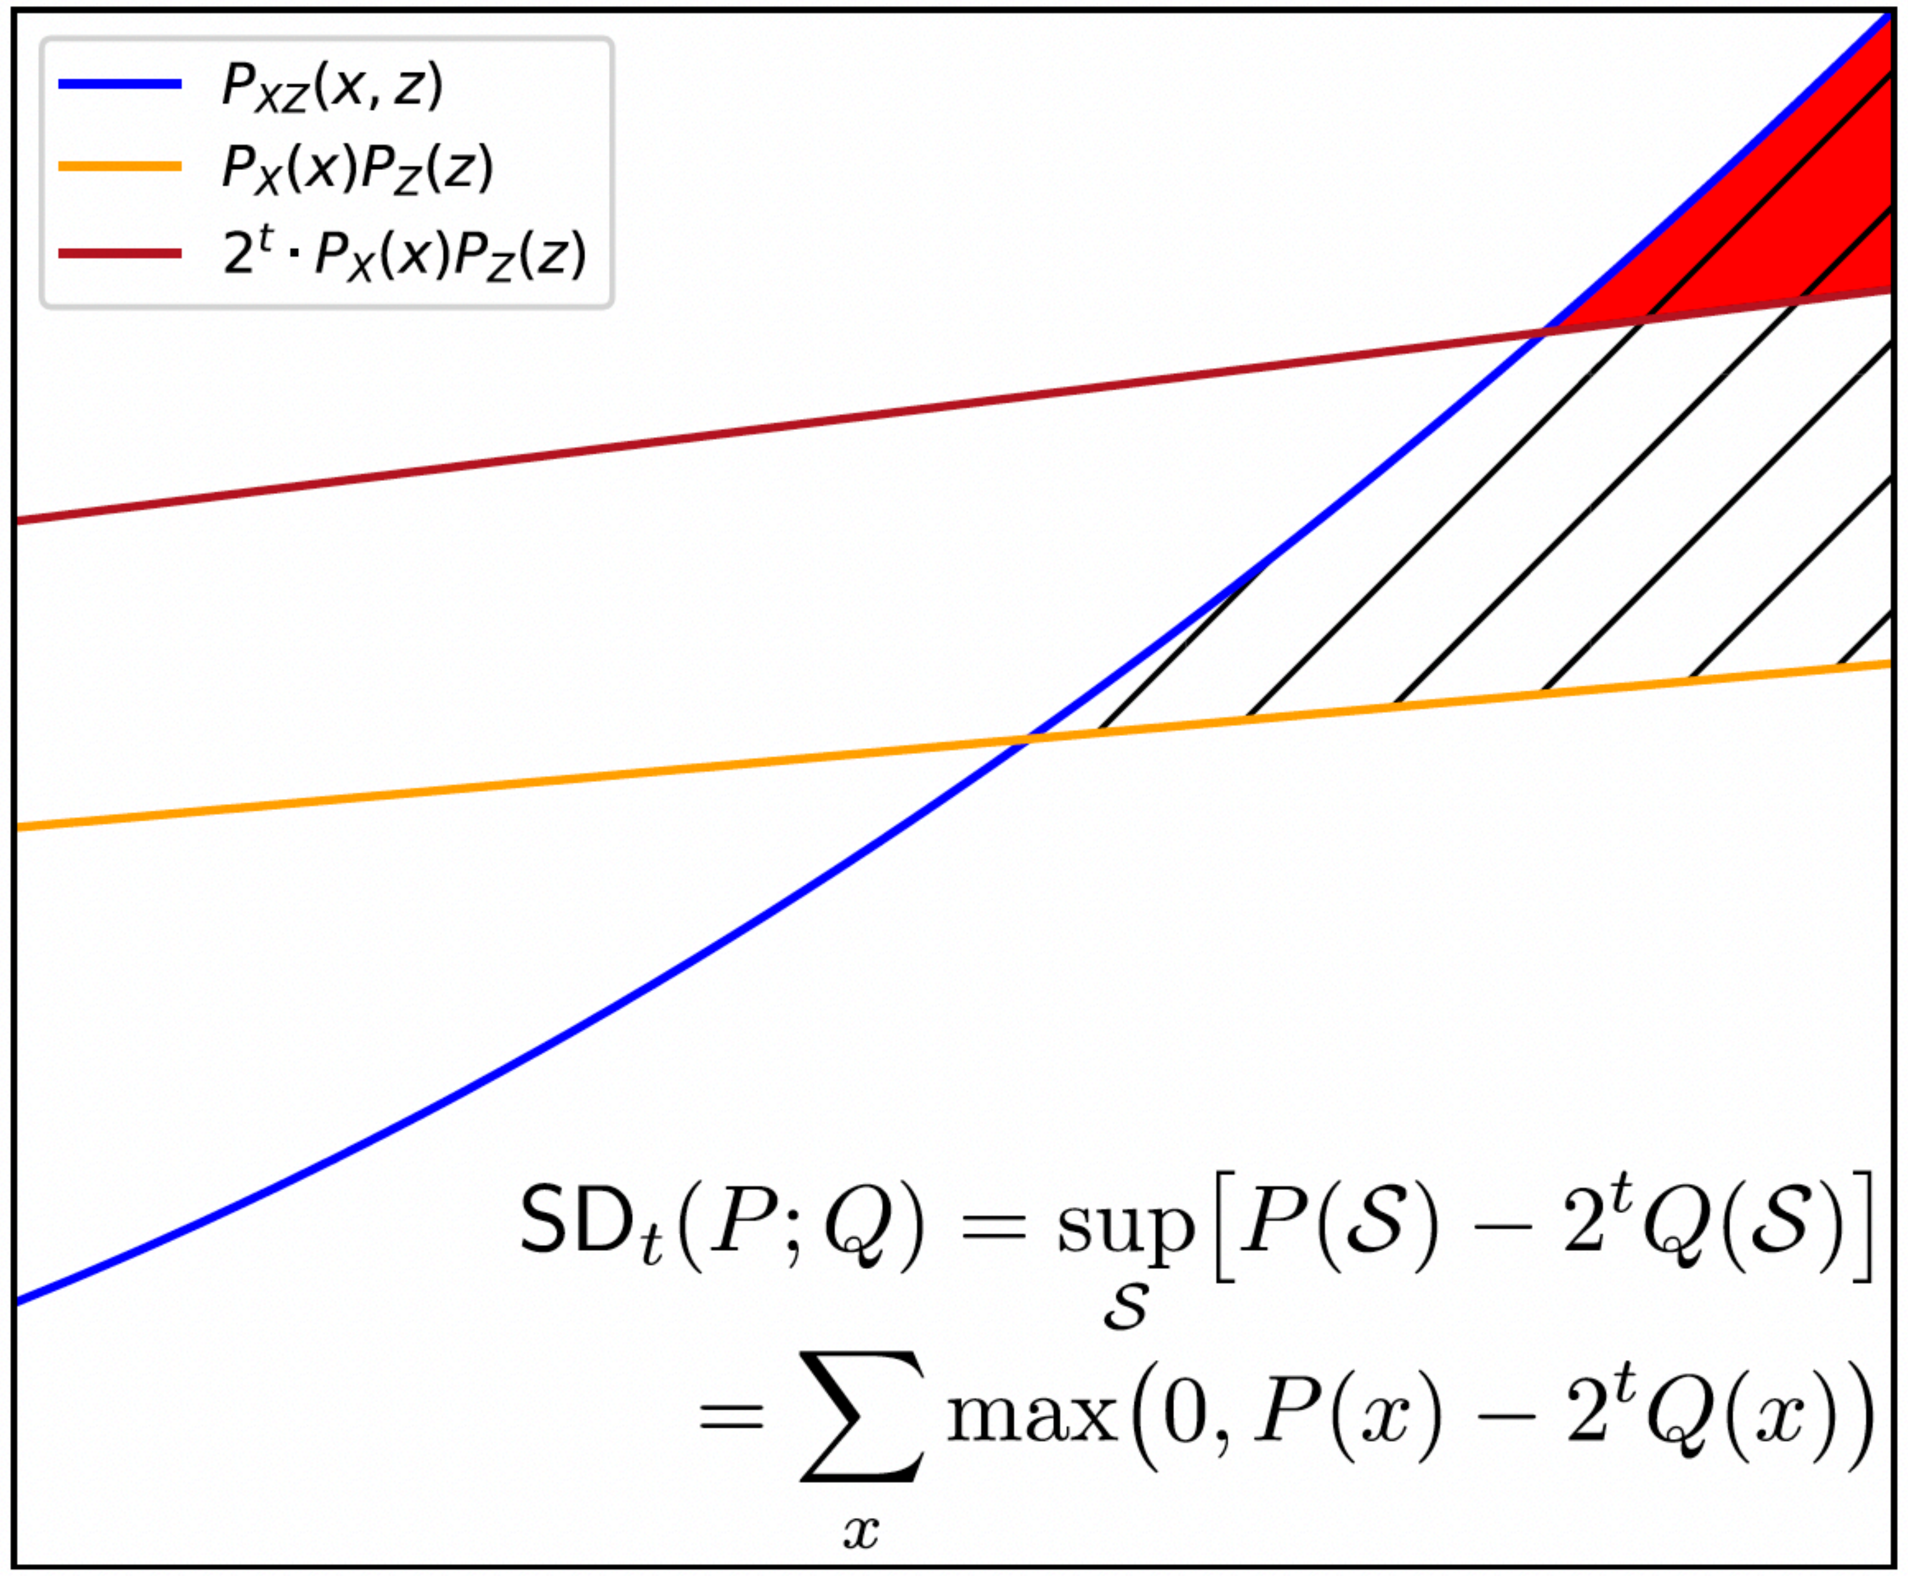

Simulation by bounded leakage

For any $\alpha > 0$, (t, δ)-SD-noisy leakage is $(\varepsilon = \delta + \alpha)$-simulatable from $t + \log \ln(1/\alpha)$ bits of bounded leakage.

Essentially,

$t \approx$ amount of bounded leakage,

$\delta \approx$ simulation error.

Also in our work

Also in our work

- (t, δ)-SD-noisy leakage existing noisy leakage models used by practitioners with good parameters — **as we increase t, the error δ falls very quickly**;

Also in our work

- (t, δ)-SD-noisy leakage existing noisy leakage models used by practitioners with good parameters — **as we increase t, the error δ falls very quickly**;

- (t, δ)-SD-noisy leakages compose nicely — **uses connection to differential privacy**

Wrapping up

Wrapping up

- Theory of leakage-resilience focuses on the bounded leakage model;

Wrapping up

- Theory of leakage-resilience focuses on the bounded leakage model;
- Real-world side-channel attacks produce noisy unbounded leakage;

Wrapping up

- Theory of leakage-resilience focuses on the bounded leakage model;
- Real-world side-channel attacks produce noisy unbounded leakage;
- ε-resilience to t-bounded leakage implies $(\varepsilon + \delta)$-resilience to (t, δ)-SD-noisy leakage;

Wrapping up

- Theory of leakage-resilience focuses on the bounded leakage model;
- Real-world side-channel attacks produce noisy unbounded leakage;
- ε-resilience to t-bounded leakage implies $(\varepsilon + \delta)$-resilience to (t, δ)-SD-noisy leakage;
- (t, δ)-SD-noisy leakage captures practical noisy leakage models with good parameters;

Wrapping up

- Theory of leakage-resilience focuses on the bounded leakage model;
- Real-world side-channel attacks produce noisy unbounded leakage;
- ε-resilience to t-bounded leakage implies $(\varepsilon + \delta)$-resilience to (t, δ)-SD-noisy leakage;
- (t, δ)-SD-noisy leakage captures practical noisy leakage models with good parameters;
- (t, δ)-SD-noisy leakages compose nicely.

Wrapping up

- Theory of leakage-resilience focuses on the bounded leakage model;
- Real-world side-channel attacks produce noisy unbounded leakage;
- ε-resilience to t-bounded leakage implies $(\varepsilon + \delta)$-resilience to (t, δ)-SD-noisy leakage;
- (t, δ)-SD-noisy leakage captures practical noisy leakage models with good parameters;
- (t, δ)-SD-noisy leakages compose nicely.

Thanks!