

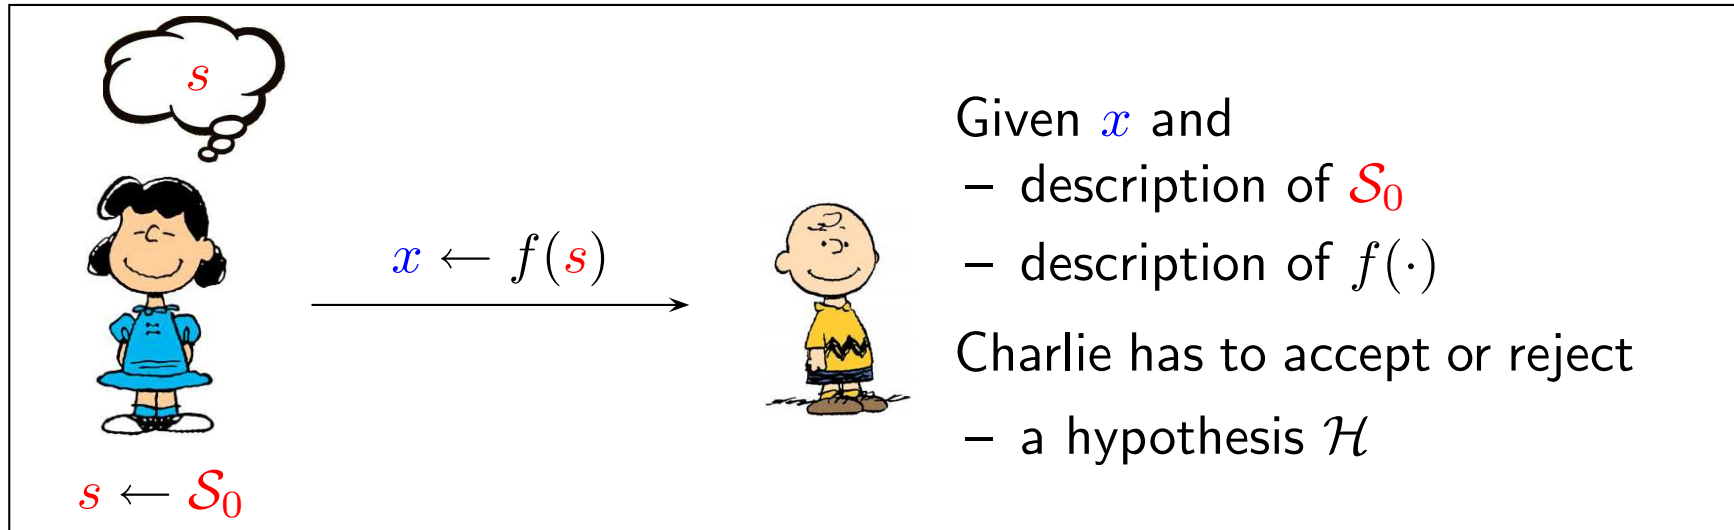
Computational Indistinguishability

Sven Laur
swen@math.ut.ee

University of Tartu

Indistinguishability

A quick recap of hypothesis testing



There are several types of hypotheses:

- ▷ **simple hypotheses** $\mathcal{H} = [s \stackrel{?}{=} s_0]$
- ▷ **complex hypotheses** $\mathcal{H} = [s \stackrel{?}{=} s_0 \vee s \stackrel{?}{=} s_1 \vee \dots \vee s \stackrel{?}{=} s_k]$
- ▷ **trivial hypotheses** that always hold or never hold.

Computational distance

To choose between hypotheses $\mathcal{H}_0 = [s \stackrel{?}{=} s_0]$ and $\mathcal{H}_1 = [s \stackrel{?}{=} s_1]$, we have to distinguish two output distributions $\mathcal{X}_0 = f(s_0)$ and $\mathcal{X}_1 = f(s_1)$.

These distributions are (t, ε) -indistinguishable if for all t -time algorithms \mathcal{A} :

$$\text{Adv}_{\mathcal{X}_0, \mathcal{X}_1}^{\text{ind}}(\mathcal{A}) = |\Pr [x \leftarrow \mathcal{X}_0 : \mathcal{A}(x) = 0] - \Pr [x \leftarrow \mathcal{X}_1 : \mathcal{A}(x) = 0]| \leq \varepsilon$$

In other terms, the distributions \mathcal{X}_0 and \mathcal{X}_1 are (t, ε) -indistinguishable if

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) \leq \varepsilon .$$

Basic properties of computational distance

- ▷ **Triangle inequality.** For all simple hypotheses $\mathcal{H}_0, \mathcal{H}_1, \mathcal{H}_2$:

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_2) \leq \text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) + \text{cd}_x^t(\mathcal{H}_1, \mathcal{H}_2) .$$

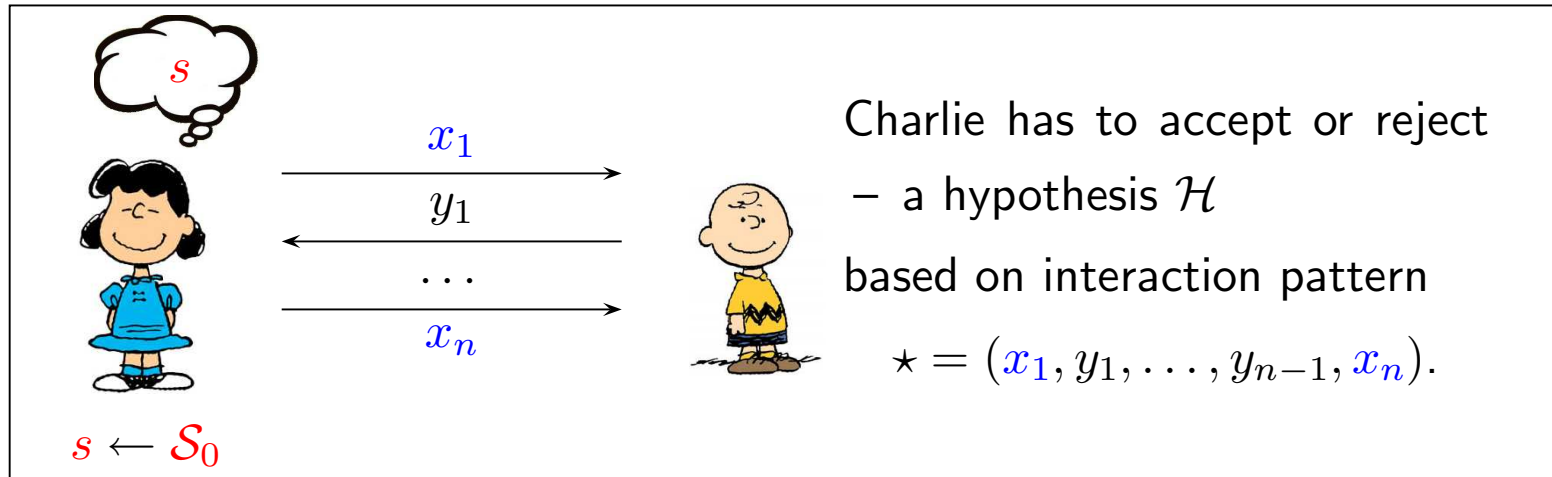
- ▷ **Symmetry.** For any two simple hypothesis \mathcal{H}_0 and \mathcal{H}_1 :

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) = \text{cd}_x^t(\mathcal{H}_1, \mathcal{H}_0) .$$

- ▷ **Positively definiteness.** For any reasonably large time bound t :

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) = 0 \quad \Leftrightarrow \quad \text{sd}_x(\mathcal{H}_0, \mathcal{H}_1) = 0 \quad \Leftrightarrow \quad \mathcal{H}_0 \equiv \mathcal{H}_1 .$$

Interactive hypothesis testing



We use analogous notation for computational and statistical distance:

$$\text{cd}_{\star}^t(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A} \text{ is } t\text{-time}} |\Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_0] - \Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_1]| \quad ,$$

$$\text{sd}_{\star}(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A}} |\Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_0] - \Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_1]| \quad .$$

These measures also satisfy triangle inequality and other distance axioms.

Examples

Pseudorandom functions

Let \mathcal{F}_{all} denote the set of all functions $f : \mathcal{M} \rightarrow \mathcal{C}$ and let $\mathcal{F} \subseteq \mathcal{F}_{\text{all}}$ be a function family. Then we can consider the following interactive hypothesis testing scenario. A t -time adversary \mathcal{A} that makes at most q calls to the oracle $\mathcal{O}(\cdot)$ in order to distinguish two worlds (hypotheses):

- ▷ \mathcal{H}_0 : Oracle chooses $f \xleftarrow{u} \mathcal{F}_{\text{all}}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.
- ▷ \mathcal{H}_1 : Oracle chooses $f \xleftarrow{u} \mathcal{F}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.

We say that \mathcal{F} is (t, q, ε) -**pseudorandom function family** if for any t -time adversary \mathcal{A} that makes at most q queries the corresponding advantage

$$\text{Adv}^{\text{ind}}(\mathcal{A}) = |\Pr [f \xleftarrow{u} \mathcal{F}_{\text{all}} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0] - \Pr [f \xleftarrow{u} \mathcal{F} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0]| \leq \varepsilon .$$

Pseudorandom permutations

Let \mathcal{F}_{prm} denote the set of all permutations $f : \mathcal{M} \rightarrow \mathcal{M}$ and let $\mathcal{F} \subseteq \mathcal{F}_{\text{prm}}$ be a permutation family. Then we can consider the following interactive hypothesis testing scenario. A t -time adversary \mathcal{A} that makes at most q calls to the oracle $\mathcal{O}(\cdot)$ in order to distinguish two worlds (hypotheses):

- ▷ \mathcal{H}_0 : Oracle chooses $f \xleftarrow{u} \mathcal{F}_{\text{prm}}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.
- ▷ \mathcal{H}_1 : Oracle chooses $f \xleftarrow{u} \mathcal{F}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.

We say that \mathcal{F} is (t, q, ε) -**pseudorandom permutation family** if for any t -time adversary \mathcal{A} that makes at most q queries the corresponding advantage

$$\text{Adv}^{\text{ind}}(\mathcal{A}) = |\Pr [f \xleftarrow{u} \mathcal{F}_{\text{prm}} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0] - \Pr [f \xleftarrow{u} \mathcal{F} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0]| \leq \varepsilon .$$

Pseudorandom generators

Let f be a function that stretches m -bit seed s to n -bit string. Then we can consider the following classical hypothesis testing scenario. A t -time adversary \mathcal{A} gets x and must distinguish two worlds (hypotheses):

- ▷ \mathcal{H}_0 : The string x is uniformly chosen over $\{0, 1\}^n$.
- ▷ \mathcal{H}_1 : The string $x \leftarrow f(s)$ for uniformly chosen $s \leftarrow_u \{0, 1\}^m$.

We say that f is (t, ε) -**pseudorandom generator** if for any t -time adversary \mathcal{A} the corresponding advantage is bounded

$$\text{Adv}^{\text{ind}}(\mathcal{A}) = |\Pr [x \leftarrow_u \{0, 1\}^n : \mathcal{A}(x) = 0] - \Pr [s \leftarrow_u \{0, 1\}^m : \mathcal{A}(f(s)) = 0]| \leq \varepsilon .$$

Practical implementations

- ▷ **Pseudorandom functions.** Constructing good pseudorandom functions has never been an explicit design goal. Cryptographic hash functions $h : \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{T}$ with implicit or explicit keys are often treated as pseudorandom functions. However, they are also known to contain much more weaknesses than good block ciphers.
- ▷ **Pseudorandom permutations.** Block ciphers are specifically designed to be pseudorandom permutations. This is the most thoroughly studied branch of practical primitive design and we have many good candidates.
- ▷ **Pseudorandom generators.** Stream ciphers are designed to be fast pseudorandom generators. However, we know much more about block ciphers than about stream ciphers. In fact, there is no widely adopted stream cipher standard. There are also more secure constructions based on number theoretical constructions but they are much slower.

Guessing Games

Simplest guessing game

Consider the simplest attack scenario:

1. \mathcal{S}_0 is a uniform distribution over two states s_0 and s_1 .
2. \mathcal{H}_0 and \mathcal{H}_1 denote simple hypotheses $[s \stackrel{?}{=} s_0]$ and $[s \stackrel{?}{=} s_1]$.
3. Given $x \leftarrow f(s)$, Charlie must choose between hypotheses \mathcal{H}_0 and \mathcal{H}_1 .

The probability of an incorrect guess

$$\begin{aligned} \Pr[\text{Failure}] &= \Pr[\mathcal{H}_0] \cdot \Pr[\mathcal{A}(x) = 1|\mathcal{H}_0] + \Pr[\mathcal{H}_1] \cdot \Pr[\mathcal{A}(x) = 0|\mathcal{H}_1] \\ &= \frac{1}{2} \cdot \left(\underbrace{\Pr[\mathcal{A}(x) = 1|\mathcal{H}_0]}_{\text{False negatives}} + \underbrace{\Pr[\mathcal{A}(x) = 0|\mathcal{H}_1]}_{\text{False positives}} \right) \\ &= \frac{1}{2} + \frac{1}{2} \cdot \underbrace{\left(\Pr[\mathcal{A}(x) = 0|\mathcal{H}_1] - \Pr[\mathcal{A}(x) = 0|\mathcal{H}_0] \right)}_{\pm \text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1)} . \end{aligned}$$

Guessing game with a biased coin

Let \mathcal{D} be a distribution over $\{0, 1\}$ such that $\Pr [i \leftarrow \mathcal{D} : i = 0] \leq \frac{1}{2}$ and consider a guessing game \mathcal{G} between a challenger and an adversary \mathcal{A} :

$$\mathcal{G}^{\mathcal{A}} \left[\begin{array}{l} i \leftarrow \mathcal{D} \\ b \leftarrow \mathcal{A}(f(s_i)) \\ \text{return } [b \stackrel{?}{=} i] \end{array} \right.$$

For this game, the adversary succeeds with probability

$$\begin{aligned} \Pr [\text{Success}] &= \Pr [\mathcal{H}_0] \cdot \Pr [\mathcal{A} = 0 | \mathcal{H}_0] + \Pr [\mathcal{H}_1] \cdot \Pr [\mathcal{A} = 1 | \mathcal{H}_1] \\ &\leq \Pr [\mathcal{H}_1] \cdot (1 + \Pr [\mathcal{A} = 0 | \mathcal{H}_0] - \Pr [\mathcal{A} = 0 | \mathcal{H}_1]) \\ &\leq \Pr [\mathcal{H}_1] + \text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) . \end{aligned}$$

Choosing between many values

Now consider a game

$$\mathcal{G}^A \left[\begin{array}{l} s \leftarrow \mathcal{S}_0 \\ s' \leftarrow \mathcal{A}(f(s)) \\ \text{return } [s \stackrel{?}{=} s'] \end{array} \right.$$

If for all possible states $s_i, s_j \in \text{supp}(\mathcal{S}_0)$ distributions $f(s_i)$ and $f(s_j)$ are (t, ε) -indistinguishable, then for all t -time algorithms

$$\min_s \Pr [s] - \varepsilon \leq \Pr [\text{Success}] \leq \max_s \Pr [s] + \varepsilon .$$

The corresponding proof

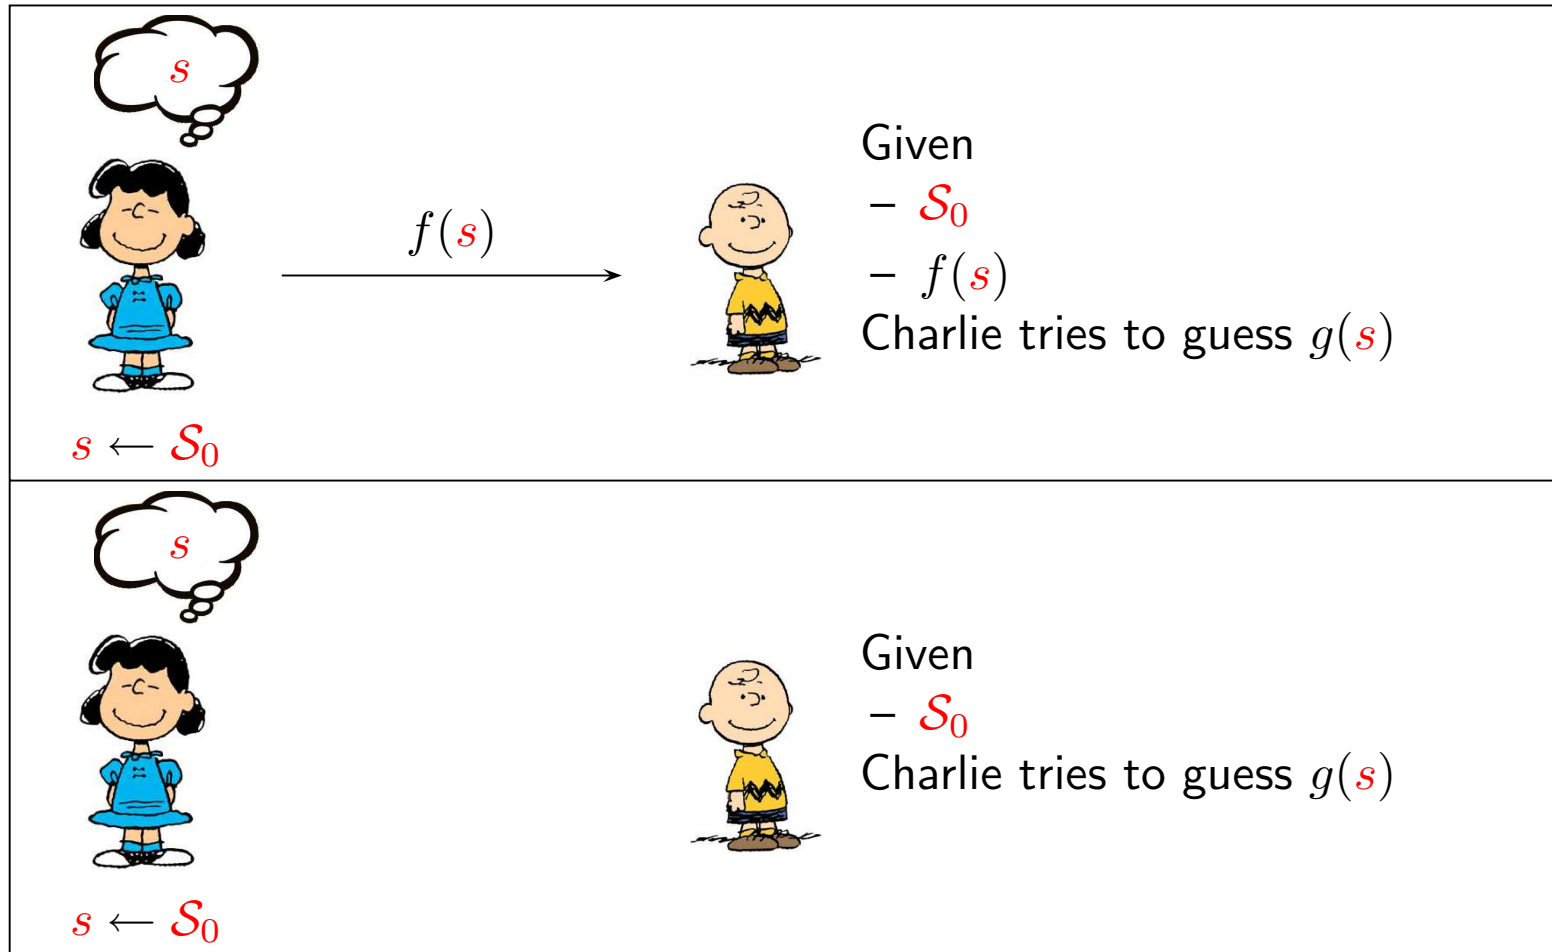
Let s_* the element with the maximal probability over \mathcal{S}_0 . Then

$$\begin{aligned}\Pr[\text{Success}] &= \sum_{s \neq s_*} \Pr[s] \cdot \Pr[\mathcal{A}(f(s)) = s] \\ &\quad + \Pr[s_*] - \sum_{s \neq s_*} \Pr[s_*] \cdot \Pr[\mathcal{A}(f(s_*) = s)] \\ &\leq \Pr[s_*] + \sum_{s \neq s_*} \Pr[s] \cdot \underbrace{|\Pr[\mathcal{A}(f(s)) = s] - \Pr[\mathcal{A}(f(s_*)) = s]|}_{\leq \varepsilon} \\ &\leq \Pr[s_*] + \varepsilon .\end{aligned}$$

The proof of the lower bound is analogous.

Semantic Security

Semantic security



Formal definition

Consider the following games:

$$\begin{array}{cc} \mathcal{G}_0^{\mathcal{A}} & \mathcal{G}_1^{\mathcal{A}} \\ \left[\begin{array}{l} s \leftarrow \mathcal{S}_0 \\ g' \leftarrow \mathcal{A}(f(s)) \\ \text{return } [g' \stackrel{?}{=} g(s)] \end{array} \right. & \left[\begin{array}{l} s \leftarrow \mathcal{S}_0 \\ g' \leftarrow \operatorname{argmax}_{g'} \Pr [g(s) = g'] \\ \text{return } [g' \stackrel{?}{=} g(s)] \end{array} \right. \end{array}$$

Then we can define a true guessing advantage

$$\begin{aligned} \operatorname{Adv}_{f,g}^{\operatorname{sem}}(\mathcal{A}) &= \Pr [\mathcal{G}_0^{\mathcal{A}} = 1] - \Pr [\mathcal{G}_1^{\mathcal{A}} = 1] \\ &= \Pr [s \leftarrow \mathcal{S}_0 : \mathcal{A}(f(s)) = g(s)] - \max_{g'} \Pr [g(s) = g'] \ . \end{aligned}$$

IND \implies SEM

Theorem. If for all $s_i, s_j \in \text{supp}(\mathcal{S}_0)$ distributions $f(s_i)$ and $f(s_j)$ are (t, ε) -indistinguishable, then for all t -time adversaries \mathcal{A} :

$$\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) \leq \varepsilon .$$

Note that

- ▷ function g might be randomised,
- ▷ function $g : \mathcal{S}_0 \rightarrow \{0, 1\}^*$ may be extremely difficult to compute,
- ▷ it might be even infeasible to get samples from the distribution \mathcal{S}_0 .

The corresponding proof

Assume that there exists g and \mathcal{A} such that violate the claim. Then there exists also a deterministic function g such that $\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) > \varepsilon$.

Let $\mathcal{H}_i = [g(\mathbf{s}) = i]$. Then we can prove

$$\text{cd}_x^t(\mathcal{H}_i, \mathcal{H}_j) \leq \varepsilon .$$

Now observe

$$\Pr [\text{Success}] = \sum_{i=0}^t \Pr [\mathcal{H}_i] \cdot \Pr [\mathcal{A}(f(\mathbf{s})) = i | \mathcal{H}_i] \leq \max \Pr [\mathcal{H}_i] + \varepsilon .$$

$$\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) = \Pr [\text{Success}] - \max \Pr [\mathcal{H}_i] \leq \varepsilon .$$

Hence, we have a desired contradiction.

Postmortem

We have now formally shown that if f is a (t, ε) -pseudorandom generator then it is difficult to approximate any predicate $g(\mathbf{s})$ given only $f(\mathbf{s})$.

Similarly, if $f : \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{C}$ is a (t, ε) -pseudorandom function family then

- ▷ it is difficult to approximate a predicate $g(k)$ by adaptively querying $f_k(\cdot)$,
- ▷ it is difficult to approximate a predicate $g(x)$ given only the value $f(x, k)$ and black-box access to the function $f_k(\cdot)$.

However, this general semantic security guarantee has also limitations:

- ▷ The proof is non-constructive.
- ▷ The theorem does not hold if \mathcal{S}_0 is specified by the adversary.
For example, if adversary can influence which messages are enciphered.

Switching lemma

Motivation

Block ciphers are designed to be pseudorandom permutations. However, it is much more easier to work with pseudorandom functions. Therefore, all classical security proofs have the following structure:

1. Replace pseudorandom permutation family \mathcal{F} with the family \mathcal{F}_{prm} .
2. Use the PRP/PRF switching lemma to substitute \mathcal{F}_{prm} with \mathcal{F}_{all} .
3. Solve the resulting combinatorial problem to bound the advantage:
 - ▷ All output values $f(x)$ have uniform distribution.
 - ▷ Each output $f(x)$ is independent of other outputs.

More formally, let \mathcal{G}_0 the original security game and \mathcal{G}_1 and \mathcal{G}_2 be the games obtained after replacement steps. Then

$$\text{Adv}_{\mathcal{G}_0}^{\text{win}}(\mathcal{A}) = \Pr[\mathcal{G}_0^{\mathcal{A}} = 1] \leq \text{cd}_{\star}^t(\mathcal{G}_0, \mathcal{G}_1) + \text{sd}_{\star}(\mathcal{G}_1, \mathcal{G}_2) + \Pr[\mathcal{G}_2^{\mathcal{A}} = 1] .$$

PRP/PRF switching lemma

Theorem. Let \mathcal{M} be the input and output domain for \mathcal{F}_{all} . Then the permutation family \mathcal{F}_{prm} is (q, ε) -pseudorandom function family where

$$\varepsilon \leq \frac{q(q-1)}{2|\mathcal{M}|} .$$

Theorem. Let \mathcal{M} be the input and output domain for \mathcal{F}_{all} . Then for any $q \leq \sqrt{|\mathcal{M}|}$ there exists a $O(q \log q)$ distinguisher \mathcal{A} that achieves

$$\text{Adv}_{\mathcal{F}_{\text{all}}, \mathcal{F}_{\text{prm}}}^{\text{ind}}(\mathcal{A}) \geq 0.316 \cdot \frac{q(q-1)}{|\mathcal{M}|} .$$

Birthday paradox

Obviously $f \notin \mathcal{F}_{\text{prm}}$ if we find a collision $f(x_i) = f(x_j)$ for $i \neq j$.


For the proof note that:

- ▷ If x_1, \dots, x_q are different then the outputs $f(x_1), \dots, f(x_q)$ have uniform distribution over $\mathcal{M} \times \dots \times \mathcal{M}$ when $f \xleftarrow{u} \mathcal{F}_{\text{all}}$.
- ▷ Hence the corresponding adversary \mathcal{A} that outputs 0 only in case of collision obtains

$$\begin{aligned} \text{Adv}_{\mathcal{F}_{\text{all}}, \mathcal{F}_{\text{prm}}}^{\text{ind}}(\mathcal{A}) &= \Pr[\text{Collision} | \mathcal{F}_{\text{all}}] - \Pr[\text{Collision} | \mathcal{F}_{\text{prm}}] \\ &= \Pr[\text{Collision} | \mathcal{F}_{\text{all}}] \geq 0.316 \cdot \frac{q(q-1)}{|\mathcal{M}|} . \end{aligned}$$

Distinguishing strategy as decision tree

Let \mathcal{A} be a deterministic distinguisher that makes **up to** q oracle calls.



Then $\Pr[\text{Vertex } u | \mathcal{F}_{\text{prm}}]$ and $\Pr[\text{Vertex } u | \mathcal{F}_{\text{all}} \wedge \neg \text{Collision}]$ might differ. However, if \mathcal{A} makes **exactly** q queries then all vertices on decision border are sampled with uniform probability and thus

$$\Pr[\mathcal{A} = 0 | \mathcal{F}_{\text{prm}}] = \Pr[\mathcal{A} = 0 | \mathcal{F}_{\text{all}} \wedge \neg \text{Collision}] .$$

The corresponding proof

Obviously, the best distinguisher \mathcal{A} is deterministic and makes exactly q oracle calls. Consequently,

$$\begin{aligned}\Pr[\mathcal{A} = 0 | \mathcal{F}_{\text{all}}] &= \Pr[\text{Collision} | \mathcal{F}_{\text{all}}] \cdot \Pr[\mathcal{A} = 0 | \mathcal{F}_{\text{all}} \wedge \text{Collision}] \\ &\quad + \Pr[\neg \text{Collision} | \mathcal{F}_{\text{all}}] \cdot \Pr[\mathcal{A} = 0 | \mathcal{F}_{\text{all}} \wedge \neg \text{Collision}] \\ &\leq \Pr[\text{Collision} | \mathcal{F}_{\text{all}}] + \Pr[\mathcal{A} = 0 | \mathcal{F}_{\text{prm}}]\end{aligned}$$

and thus also

$$\text{Adv}_{\mathcal{F}_{\text{all}}, \mathcal{F}_{\text{prm}}}^{\text{ind}}(\mathcal{A}) \leq \Pr[\text{Collision} | \mathcal{F}_{\text{all}}] \ .$$

Now observe

$$\Pr\left[\bigvee_{i \neq j} f(x_i) = f(x_j)\right] \leq \sum_{i \neq j} \Pr[f(x_i) = f(x_j)] = \frac{q(q-1)}{2} \cdot \frac{1}{|\mathcal{M}|} \ .$$