Defining security of cryptographic primitives
The hybrid argument

Formally defining security of cryptoprimitives

- Let us move back to "computational" world:
  - Messages are bit-strings;
  - Encryption, decryption, key generation, signing, etc. are PPT algorithms on bit-strings.
  - Adversary is an(y) interactive PPT algorithm.
- Primitive is secure if adversary's succeeds in breaking it with a low probability.
  - A function $f : \mathbb{N} \to \mathbb{R}$ is negligible if for all polynomials, $\lim_{\eta\to\infty} f(\eta) \cdot p(\eta) = 0$.
  - I.e. the inverse of f is superpolynomial.
  - η is the security parameter
    - Where does it come from?

Security parameter

- We need an integer parameter for speaking about asymptotic security.
- η is something that
  - the work of honest participants is polynomial in η;
  - the work of the adversary is hopefully superpolynomial in η.
- It could be
  - the key / plaintext length in asymmetric encryption and signing;
  - the length of the challenge in identification protocols.
- But also
  - key / block length in block ciphers / symmetric encryption;
  - key / tag length in MACs;
  - output length in hash functions

  although the common definitions for those are usually not parameterized.

Security of symmetric encryption

- We want the ciphertext to hide all partial information.
  - At least information that can be found in polynomial time.
- Let $H : \{0,1\}^* \to \{0,1\}^*$ be a polynomial-time algorithm.
- We pick a plaintext x.
- We give η and $y = \mathcal{E}_k(\eta, x)$ to the adversary.
- The adversary answers with $z \in \{0,1\}^*$.
- The adversary wins if $z = H(x)$.
- We want the adversary's winning probability to be negligible in η.

Exercise. What is wrong with this definition?

Semantic security

- For all polynomial-time algorithms $H : \{0,1\}^* \rightarrow \{0,1\}^*$
- for all polynomial-time constructible families of probability distributions $\{M_\eta\}_{\eta\in\mathbb{N}}$ over bit-strings
- for all PPT adversaries $\mathcal{A}$
- the probability

$$\Pr[h^* = h \,|\, x \leftarrow M_\eta, h = H(x), y \leftarrow \mathcal{E}_k(\eta, x), h^* \leftarrow \mathcal{A}(\eta, y)]$$

is at most negligibly larger than the probability

$$\Pr[h^* = h \,|\, x, x' \leftarrow M_\eta, h = H(x'), y \leftarrow \mathcal{E}_k(\eta, x), h^* \leftarrow \mathcal{A}(\eta, y)]$$

- Then $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ has semantic security against chosen-plaintext attacks.

Simplifying semantic security

- H, M and $\mathcal{A}$ are all polynomial-time algorithms.
- Put them all into $\mathcal{A}$:
  - $\mathcal{A}$ first outputs H and M;
  - then x is picked according to M and $y = \mathcal{E}_k(\eta, x)$ is given to $\mathcal{A}$;
  - then $\mathcal{A}$ tries to find $H(x)$.
- Restrict $\mathcal{A}$:
  - Let H be identity function.
  - Let M_η be a distribution that assigns 50% to some m_0, 50% to some m_1 and nothing to any other bit-string.
    - To specify M_η, $\mathcal{A}$ outputs m_0 and m_1.
    - m_0 and m_1 must have equal length.

Find-then-guess security

- $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ — a symmetric encryption scheme.
- Let k be generated by $\mathcal{K}(\eta)$.
- Let $b \in_R \{0, 1\}$ be uniformly generated.
- The adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ works as follows:
  - $\mathcal{A}_1(\eta)$ returns two messages m_0, m_1 of equal length and some internal state s.
  - Invoke $\mathcal{E}_k(\eta, m_b)$. Let y be the result.
  - $\mathcal{A}_2(s, y)$ outputs a bit b^*.
- Encryption scheme has find-then-guess security against chosen-plaintext attacks if the probability of $b = b^*$ is not larger than $1/2 + f(\eta)$ for some negligible f.

Exercise. Show that find-then-guess security implies semantic security.

Indistinguishability of probability distributions

- For each $\eta \in \mathbb{N}$ let D^0_η and D^1_η be probability distributions over bit-strings.
- The families of probability distributions $D^0 = \{D^0_\eta\}_{\eta\in\mathbb{N}}$ and $D^1 = \{D^1_\eta\}_{\eta\in\mathbb{N}}$ are indistinguishable if
  - for any adversary $\mathcal{A}$
    - The running time of $\mathcal{A}(\eta,\cdot)$ must be polynomial in η
  - the difference of probabilities

$$\Pr[\mathcal{A}(\eta,x)=1 \,|\, x \leftarrow D^0_\eta] - \Pr[\mathcal{A}(\eta,x)=1 \,|\, x \leftarrow D^1_\eta]$$

    is a negligible function of η.
- Denote $D^0 \approx D^1$.

Transitivity

Theorem. If $D^0 \approx D^1$ and $D^1 \approx D^2$, then $D^0 \approx D^2$.

Proof.

- Suppose that $D^0 \not\approx D^2$.
- Let $\mathcal{A}$ be a polynomial-time adversary such that $\mathcal{A}$ can distinguish D^0 and D^2 with non-negligible advantage.
- For $i \in \{0, 1, 2\}$, let

$$p^i_\eta = \Pr[\mathcal{A}(\eta, x) = 1 \,|\, x \leftarrow D^i_\eta]$$

- There is a polynomial q, such that for infinitely many η, $|{p^0}_\eta - {p^2}_\eta| \geq q(\eta)$.
- For any such η, either $|{p^0}_\eta - {p^1}_\eta| \geq q(\eta)/2$ or $|{p^1}_\eta - {p^2}_\eta| \geq q(\eta)/2$.
- Either $|{p^0}_\eta - {p^1}_\eta| \geq q(\eta)/2$ holds for infinitely many η, or $|{p^1}_\eta - {p^2}_\eta| \geq q(\eta)/2$ holds for infinitely many η.
- $\mathcal{A}$ distinguishes either D^0 and D^1, or D^1 and D^2. $\square$

Independent components

- Let D^0, D^1, E be families of probability distributions.
- Define the probability distribution F^i_η by

  1. Let $x \leftarrow D^i_\eta$.
  2. Let $y \leftarrow E_\eta$.
  3. Output (x, y).

- E is polynomial-time constructible if there is a polynomial-time algorithm $\mathcal{E}$, such that the output of $\mathcal{E}(\eta)$ is distributed identically to E_η.
- **Theorem.** If $D^0 \approx D^1$ and E is polynomial-time constructible, then $F^0 \approx F^1$.

Proof

- Suppose that $F^0 \not\approx F^1$.
- Let $\mathcal{A}$ be a polynomial-time adversary such that $\mathcal{A}$ can distinguish F^0 and F^1 with non-negligible advantage.
- Construct $\mathcal{B}$ as follows: on input (η, x), it will
  - call $\mathcal{E}(\eta)$, giving y;
  - call $\mathcal{A}(\eta, (x, y))$, giving b;
  - return b.
- We see that
  - if x is distributed according to ${D^0}_\eta$, then the argument to $\mathcal{A}$ is distributed according to ${F^0}_\eta$;
  - if x is distributed according to ${D^1}_\eta$, then the argument to $\mathcal{A}$ is distributed according to ${F^1}_\eta$;

hence the advantage of $\mathcal{B}$ is equal to the advantage of $\mathcal{A}$. □

Multiple sampling

- Let $D^0 = \{D^0_\eta\}_{\eta\in\mathbb{N}}$ and $D^1 = \{D^1_\eta\}_{\eta\in\mathbb{N}}$ be two families of probability distributions.
- Let p be a positive polynomial.
- Let $\vec{D}^b_\eta$ be a probability distribution over tuples

$$(x_1, x_2, \ldots, x_{p(\eta)}) \in (\{0,1\}^*)^{p(\eta)}$$

  such that

  - each x_i is distributed according to D^b_η;
  - each x_i is is independent of all other x-s.

Multiple sampling

- Let $D^0 = \{D^0_\eta\}_{\eta\in\mathbb{N}}$ and $D^1 = \{D^1_\eta\}_{\eta\in\mathbb{N}}$ be two families of probability distributions.
- Let p be a positive polynomial.
- Let $\vec{D}^b_\eta$ be a probability distribution over tuples

$$(x_1, x_2, \ldots, x_{p(\eta)}) \in (\{0,1\}^*)^{p(\eta)}$$

  such that
  - each x_i is distributed according to D^b_η;
  - each x_i is is independent of all other x-s.
- To sample $\vec{D}^b_\eta$, sample D^b_η $p(\eta)$ times and construct the tuple of sampled values.

$\vec{D}$-s indistinguishable $\Rightarrow$ D-s indistinguishable

Theorem. If $\vec{D}^0 \approx \vec{D}^1$ then $D^0 \approx D^1$.

$\vec{D}$-s indistinguishable $\Rightarrow$ D-s indistinguishable

Theorem. If $\vec{D}^0 \approx \vec{D}^1$ then $D^0 \approx D^1$.
If ●●● $\approx$ ●●● then ● $\approx$ ●.

Contrapositive: if ● $\not\approx$ ● then ●●● $\not\approx$ ●●●

$\vec{D}$-s indistinguishable $\Rightarrow$ D-s indistinguishable

Theorem. If $\vec{D}^0 \approx \vec{D}^1$ then $D^0 \approx D^1$.
If ●●● $\approx$ ●●● then ● $\approx$ ●.

Contrapositive: if ● $\not\approx$ ● then ●●● $\not\approx$ ●●●
If ● $\not\approx$ ● then there exists a PPT distinguisher $\mathcal{A}$:

$$\Pr[\mathcal{A}(\eta, x) = 0 \,|\, x \leftarrow D^0_\eta] - \Pr[\mathcal{A}(\eta, x) = 0 \,|\, x \leftarrow D^1_\eta] \geq 1/q(\eta)$$

for some polynomial q and infinitely many η.

$\vec{D}$-s indistinguishable $\Rightarrow$ D-s indistinguishable

Theorem. If $\vec{D}^0 \approx \vec{D}^1$ then $D^0 \approx D^1$.
If ●●● $\approx$ ●●● then ● $\approx$ ●.

Contrapositive: if ● $\not\approx$ ● then ●●● $\not\approx$ ●●●
If ● $\not\approx$ ● then there exists a PPT distinguisher $\mathcal{A}$:

$$\Pr[\mathcal{A}(\eta, x) = 0 \,|\, x \leftarrow D^0_\eta] - \Pr[\mathcal{A}(\eta, x) = 0 \,|\, x \leftarrow D^1_\eta] \geq 1/q(\eta)$$

for some polynomial q and infinitely many η.

Let $\mathcal{B}(\eta, (x_1, \ldots, x_{p(\eta)})) = \mathcal{A}(\eta, x_1)$.
Then $\mathcal{B}$ distinguishes ●●● and ●●●.

$\vec{D}$-s indistinguishable $\Rightarrow$ D-s indistinguishable

Theorem. If $\vec{D}^0 \approx \vec{D}^1$ then $D^0 \approx D^1$.
If ●●● $\approx$ ●●● then ● $\approx$ ●.

Contrapositive: if ● $\not\approx$ ● then ●●● $\not\approx$ ●●●
If ● $\not\approx$ ● then there exists a PPT distinguisher $\mathcal{A}$:

$$\Pr[\mathcal{A}(\eta, x) = 0 \,|\, x \leftarrow D^0_\eta] - \Pr[\mathcal{A}(\eta, x) = 0 \,|\, x \leftarrow D^1_\eta] \geq 1/q(\eta)$$

for some polynomial q and infinitely many η.

Let $\mathcal{B}(\eta, (x_1, \ldots, x_{p(\eta)})) = \mathcal{A}(\eta, x_1)$.
Then $\mathcal{B}$ distinguishes ●●● and ●●●.

I.e. we can distinguish ●●● from ●●● by just considering the first elements of the tuples.

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if $\bullet \approx \bullet$ then $\bullet\bullet\bullet \approx \bullet\bullet\bullet$. (let $p = 3$)

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if • ≈ • then ••• ≈ •••. (let $p = 3$)

Our lemmas said (• ≈ • ∧ • ≈ •) ⇒ • ≈ • and • ≈ • ⇒ •• ≈ ••.

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if $\bullet \approx \bullet$ then $\bullet\bullet\bullet \approx \bullet\bullet\bullet$. (let $p = 3$)

Our lemmas said $(\bullet \approx \bullet \wedge \bullet \approx \bullet) \Rightarrow \bullet \approx \bullet$ and $\bullet \approx \bullet \Rightarrow \bullet\bullet \approx \bullet\bullet$.

$\bullet\bullet\bullet$

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if • ≈ • then ••• ≈ •••. (let $p = 3$)

Our lemmas said (• ≈ • ∧ • ≈ •) ⇒ • ≈ • and • ≈ • ⇒ •• ≈ ••.

••• ≈ •••

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if $\bullet \approx \bullet$ then $\bullet\bullet\bullet \approx \bullet\bullet\bullet$. (let $p = 3$)

Our lemmas said $(\bullet \approx \bullet \wedge \bullet \approx \bullet) \Rightarrow \bullet \approx \bullet$ and $\bullet \approx \bullet \Rightarrow \bullet\bullet \approx \bullet\bullet$.

$\bullet\bullet\bullet \approx \bullet\bullet\bullet \approx \bullet\bullet\bullet$

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if ${\color{red}\bullet} \approx {\color{blue}\bullet}$ then ${\color{red}\bullet\bullet\bullet} \approx {\color{blue}\bullet\bullet\bullet}$. (let $p = 3$)

Our lemmas said $({\color{red}\bullet} \approx {\color{green}\bullet} \wedge {\color{green}\bullet} \approx {\color{blue}\bullet}) \Rightarrow {\color{red}\bullet} \approx {\color{blue}\bullet}$ and ${\color{red}\bullet} \approx {\color{blue}\bullet} \Rightarrow {\color{red}\bullet}{\color{green}\bullet} \approx {\color{blue}\bullet}{\color{green}\bullet}$.

${\color{red}\bullet\bullet\bullet} \approx {\color{red}\bullet\bullet}{\color{blue}\bullet} \approx {\color{red}\bullet}{\color{blue}\bullet\bullet} \approx {\color{blue}\bullet\bullet\bullet}$.

D-s indistinguishable $\Rightarrow$ $\vec{D}$-s indistinguishable

(Interesting) theorem. If $D^0 \approx D^1$ and there exist polynomial-time algorithms $\mathcal{D}^0$ and $\mathcal{D}^1$, such that the output distribution of $\mathcal{D}^b(\eta)$ is equal to D^b_η, then $\vec{D}^0 \approx \vec{D}^1$.

Assume for now that the polynomial p is a constant. I.e. the length of the vector $\vec{x}$ does not depend on the security parameter η.
Let p be the common value of $p(\eta)$ for all η.

Theorem statement: if • ≈ • then ••• ≈ •••. (let $p = 3$)

Our lemmas said (• ≈ • ∧ • ≈ •) ⇒ • ≈ • and • ≈ • ⇒ •• ≈ ••.

••• ≈ ••• ≈ ••• ≈ •••. By transitivity, ••• ≈ •••.

(Actually, we're done with this case)

Constructing the distinguisher

Contraposition: if $\color{red}{\bullet\bullet\bullet} \not\approx \color{blue}{\bullet\bullet\bullet}$ then $\color{red}{\bullet} \not\approx \color{blue}{\bullet}$.

Constructing the distinguisher

Contrapositive: if ●●● $\not\approx$ ●●● then ● $\not\approx$ ●.
If ●●● $\not\approx$ ●●● then there exists a PPT distinguisher $\mathcal{A}$:

$$\Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow \vec{D}^0_\eta] - \Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow \vec{D}^1_\eta] \geq 1/q(\eta)$$

for some polynomial q and infinitely many η.

Hybrid distributions

If ${\color{red}\bullet\bullet\bullet} \not\approx {\color{blue}\bullet\bullet\bullet}$ then

$$\left({\color{red}\bullet\bullet\bullet} \not\approx {\color{red}\bullet\bullet}{\color{blue}\bullet}\right) \vee \left({\color{red}\bullet\bullet}{\color{blue}\bullet} \not\approx {\color{red}\bullet}{\color{blue}\bullet\bullet}\right) \vee \left({\color{red}\bullet}{\color{blue}\bullet\bullet} \not\approx {\color{blue}\bullet\bullet\bullet}\right)$$

Hybrid distributions

If ${\color{red}\bullet\bullet\bullet} \not\approx {\color{blue}\bullet\bullet\bullet}$ then

$$({\color{red}\bullet\bullet\bullet} \not\approx {\color{red}\bullet\bullet}{\color{blue}\bullet}) \vee ({\color{red}\bullet\bullet}{\color{blue}\bullet} \not\approx {\color{red}\bullet}{\color{blue}\bullet\bullet}) \vee ({\color{red}\bullet}{\color{blue}\bullet\bullet} \not\approx {\color{blue}\bullet\bullet\bullet})$$

Let $\vec{E}_\eta^k$, where $0 \leq k \leq p$, be a probability distribution over tuples $(x_1, \ldots, x_p)$, where

- each x_i is independent of all other x-s;
- $x_1, \ldots, x_k$ are distributed according to ${\color{red}D_\eta^0}$;
- $x_{k+1}, \ldots, x_p$ are distributed according to ${\color{blue}D_\eta^1}$.

Thus $\vec{E}_\eta^0 = {\color{blue}\vec{D}_\eta^1}$ and $\vec{E}_\eta^p = {\color{red}\vec{D}_\eta^0}$. Define $P_\eta^k = \Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow \vec{E}_\eta^k]$. Then for infinitely many η:

$$1/q(\eta) \leq P_\eta^p - P_\eta^0 = \sum_{i=1}^{p} (P_\eta^i - P_\eta^{i-1}) \ .$$

And for some j_η, $P_\eta^{j_\eta} - P_\eta^{j_\eta - 1} \geq 1/(p \cdot q(\eta))$.

$\mathcal{A}$ distinguishes hybrids

There exists j, such that $j = j_\eta$ for infinitely many η. Thus

$$\Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow \vec{E}_\eta^j] - \Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow \vec{E}_\eta^{j-1}] \geq 1/(p \cdot q(\eta))$$

for infinitely many η. We have $\vec{E}^{j-1} \not\approx \vec{E}^j$.

$\mathcal{A}$ distinguishes hybrids

There exists j, such that $j = j_\eta$ for infinitely many η. Thus

$$\Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow {\color{red}\vec{E}^j_\eta}] - \Pr[\mathcal{A}(\eta, \vec{x}) = 0 \,|\, \vec{x} \leftarrow {\color{blue}\vec{E}^{j-1}_\eta}] \geq 1/(p \cdot q(\eta))$$

for infinitely many η. We have $\vec{E}^{j-1} \not\approx \vec{E}^j$.

If we can distinguish

$$\vec{E}^j = \underbrace{{\color{red}\bullet\bullet}\cdots{\color{red}\bullet}}_{j-1}{\color{red}\bullet}\underbrace{{\color{blue}\bullet\bullet}\cdots{\color{blue}\bullet}}_{p-j}$$

from

$$\vec{E}^{j-1} = \underbrace{{\color{red}\bullet\bullet}\cdots{\color{red}\bullet}}_{j-1}{\color{blue}\bullet}\underbrace{{\color{blue}\bullet\bullet}\cdots{\color{blue}\bullet}}_{p-j}$$

using $\mathcal{A}$, then how do we distinguish ${\color{red}\bullet}$ and ${\color{blue}\bullet}$?

Distinguisher for D^0 and D^1

On input (η, x):

1. Let $x_1 := \mathcal{D}^0(\eta), \ldots, x_{j-1} := \mathcal{D}^0(\eta)$.
2. Let $x_j := x$
3. Let $x_{j+1} := \mathcal{D}^1(\eta), \ldots, x_p := \mathcal{D}^1(\eta)$
4. Let $\vec{x} = (x_1, \ldots, x_p)$.
5. Call $b^* := \mathcal{A}(\eta, \vec{x})$ and return b^*.

The advantage of this distinguisher is at least $1/(p \cdot q(\eta))$.

Distinguisher for D^0 and D^1

On input (η, x):

1. Let $x_1 := \mathcal{D}^0(\eta), \ldots, x_{j-1} := \mathcal{D}^0(\eta)$.
2. Let $x_j := x$
3. Let $x_{j+1} := \mathcal{D}^1(\eta), \ldots, x_p := \mathcal{D}^1(\eta)$
4. Let $\vec{x} = (x_1, \ldots, x_p)$.
5. Call $b^* := \mathcal{A}(\eta, \vec{x})$ and return b^*.

The advantage of this distinguisher is at least $1/(p \cdot q(\eta))$.

Unfortunately, the above construction was not constructive.

Being constructive

For infinitely many η we had

$$1/q(\eta) \leq P_\eta^p - P_\eta^0 = \sum_{i=1}^{p} (P_\eta^i - P_\eta^{i-1}) \ .$$

Hence the average value of $P_\eta^j - P_\eta^{j-1}$ is $\geq 1/(p \cdot q(\eta))$.

Being constructive

For infinitely many η we had

$$1/q(\eta) \leq P_\eta^p - P_\eta^0 = \sum_{i=1}^{p} (P_\eta^i - P_\eta^{i-1}) \ .$$

Hence the average value of $P_\eta^j - P_\eta^{j-1}$ is $\geq 1/(p \cdot q(\eta))$.

Consider the following distinguisher $\mathcal{B}(\eta, x)$:

1. Let $j \in_R \{1, \ldots, p\}$.
2. Let $x_1 := \mathcal{D}^0(\eta), \ldots, x_{j-1} := \mathcal{D}^0(\eta)$.
3. Let $x_j := x$
4. Let $x_{j+1} := \mathcal{D}^1(\eta), \ldots, x_p := \mathcal{D}^1(\eta)$
5. Let $\vec{x} = (x_1, \ldots, x_p)$.
6. Call $b^* := \mathcal{A}(\eta, \vec{x})$ and return b^*.

What $\mathcal{B}$ does

If (for example) $p = 5$, then $\mathcal{B}$ tries to distinguish

$\color{blue}{\bullet\bullet\bullet\bullet\bullet}$ and $\color{red}{\bullet}\color{blue}{\bullet\bullet\bullet\bullet}$ with probability $1/5$
$\color{red}{\bullet}\color{blue}{\bullet\bullet\bullet\bullet}$ and $\color{red}{\bullet\bullet}\color{blue}{\bullet\bullet\bullet}$ with probability $1/5$
$\color{red}{\bullet\bullet}\color{blue}{\bullet\bullet\bullet}$ and $\color{red}{\bullet\bullet\bullet}\color{blue}{\bullet\bullet}$ with probability $1/5$
$\color{red}{\bullet\bullet\bullet}\color{blue}{\bullet\bullet}$ and $\color{red}{\bullet\bullet\bullet\bullet}\color{blue}{\bullet}$ with probability $1/5$
$\color{red}{\bullet\bullet\bullet\bullet}\color{blue}{\bullet}$ and $\color{red}{\bullet\bullet\bullet\bullet\bullet}$ with probability $1/5$

The advantage of $\mathcal{B}$ is $1/p$ times the sum of $\mathcal{A}$'s advantages of distinguishing these pairs of distributions.

The advantage of $\mathcal{B}$ is

$$\frac{1}{p}\sum_{j=1}^{p} P_\eta^j - P_\eta^{j-1} = \frac{1}{p}(P_\eta^p - P_\eta^0) \geq \frac{1}{p \cdot q(\eta)} \ .$$

If p depends on η

$\mathcal{B}(\eta, x)$ is:

1. Let $j \in_R \{1, \ldots, p(\eta)\}$.
2. Let $x_1 := \mathcal{D}^0(\eta), \ldots, x_{j-1} := \mathcal{D}^0(\eta)$.
3. Let $x_j := x$
4. Let $x_{j+1} := \mathcal{D}^1(\eta), \ldots, x_{p(\eta)} := \mathcal{D}^1(\eta)$
5. Let $\vec{x} = \big(x_1, \ldots, x_{p(\eta)}\big)$.
6. Call $b^* := \mathcal{A}(\eta, \vec{x})$ and return b^*.

The advantage of $\mathcal{B}$ is at least $1/(p(\eta) \cdot q(\eta))$. $\square$

Left-or-right security

- Consider again symmetric encryption $(\mathcal{K}, \mathcal{E}, \mathcal{D})$.
- Let k be generated by $\mathcal{K}(\eta)$.
- Let $\mathcal{O}_b$ be the following oracle:
  - On input (m_0, m_1) where $|m_0| = |m_1|$, it returns an encryption of m_b with the key k.
- Let $b \in_R \{0, 1\}$ be uniformly generated.
- Let $\mathcal{A}$ have access to the oracle $\mathcal{O}_b$.
  - $\mathcal{A}$ can make as many oracle queries as it wants to.
- Encryption system has left-or-right security against chosen-plaintext attacks if no PPT $\mathcal{A}$ can guess b with probability more that $1/2 + f(\eta)$, where f is negligible.

Exercise. Show that an encryption system has left-or-right security against CPA iff it has find-then-guess security against CPA.

Real-or-constant security

- Let $\mathcal{O}_0$ be the following oracle:
  - On input m, it returns an encryption of m with the key k.
- Let $\mathcal{O}_1$ be the following oracle:
  - On input m, it returns an encryption of $\mathbf{0}^{|m|}$ with the key k.
- Let $b \in_R \{0, 1\}$ be uniformly generated.
- Let $\mathcal{A}$ have access to the oracle $\mathcal{O}_b$.
- Encryption system has real-or-constant security against chosen-plaintext attacks if no PPT $\mathcal{A}$ can guess b with probability more that $1/2 + f(\eta)$, where f is negligible.

Exercise. Show that an encryption system has left-or-right security against CPA iff it has real-or-constant security against CPA.