Fundamentals of Cryptography — Handout 2.

Yu Yu

Perfect, statistical security, leftover hash lemma and privacy amplification.

1 One-time pad and perfectly secret (secure) encryptions

In addition to the perfect secret definition (Definition 2.1) introduced in chapter 2 of the KL book, a more general security definition for any encryption scheme Π=(Gen,Enc,Dec) can based on the following experiment involving an adversary A (decoupled into a pair of algorithms A_1 and D) and a challenger C:

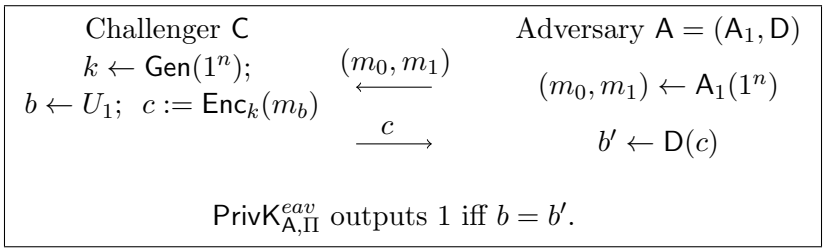

Figure 1: The adversarial indistinguishability experiment $\mathsf{PrivK}^{eav}_{\mathsf{A},\Pi}$ between A and C.

Definition 1 (statistical security)**.** An encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ over message space $\mathcal{M}$ is ε-secure (in presence of eavesdroppers) if for every adversary A it holds that

$$\Pr[\mathsf{PrivK}^{eav}_{\mathsf{A},\Pi} = 1] \leq \frac{1}{2} + \frac{\varepsilon}{2} \ .$$

Note: the above implies that for every A it holds that $\frac{1}{2} - \frac{\varepsilon}{2} \leq \Pr[\mathsf{PrivK}^{eav}_{\mathsf{A},\Pi} = 1] \leq \frac{1}{2} + \frac{\varepsilon}{2}$ (why?). In case that $\varepsilon = 0$ we say that Π is perfectly secret (which is essentially Definition 2.4 from the KL book, and see Proposition 2.5 for its equivalence to the Definition 2.1 from the same book).

Definition 2 (statistical security − alternative definition)**.** An encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ over message space $\mathcal{M}$ and ciphertext space $\mathcal{C}$ is ε-secure if for every distinguisher $\mathsf{D} : \mathcal{C} \to \{0,1\}$ and every pair of $m_0 \neq m_1 \in \mathcal{M}$ it holds that

$$\left| \Pr_{k \leftarrow \mathsf{Gen}(1^n)}[\ \mathsf{D}(\mathsf{Enc}_k(m_0)) = 1\ ] \ - \ \Pr_{k \leftarrow \mathsf{Gen}(1^n)}[\ \mathsf{D}(\mathsf{Enc}_k(m_1)) = 1\ ] \right| \ \leq \ \varepsilon \quad .$$

Lemma 1. *Definition 1 and Definition 2 are equivalent.*

Proof. It suffices to prove that

$$\Pr[b = b'] \ = \ \frac{1}{2} + \frac{1}{2} \cdot \Big(\Pr[\mathsf{D}(\mathsf{Enc}_k(m_1)) = 1] \ - \ \Pr[\mathsf{D}(\mathsf{Enc}_k(m_0)) = 1] \Big) \ .$$

□

The famous one-time pad (aka. Vernam's Cipher) is an encryption scheme with perfect secrecy.

Definition 3 (One-time pad)**.** The one-time pad encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is defined as follows:

1. For any security parameter n, let $\mathcal{M} = \mathcal{C} = \mathcal{K} = \{0,1\}^n$.
2. On input 1^n, Gen outputs a string $k \in \mathcal{K}$ uniformly at random, i.e., $k \leftarrow U_n$
3. Encryption Enc works by bitwise XORing (exclusive-or) every message bit with corresponding key bit, i.e.,
$$\mathsf{Enc}_k(m) \;=\; m \oplus k \;\;.$$
4. Decryption works by bitwise XORing the ciphertext with the key, i.e.,
$$\mathsf{Dec}_k(c) \;=\; c \oplus k \;\;.$$

Theorem 2 (Theorem 2.6 from the KL book)**.** *The one-time pad is a perfectly-secret encryption.*

2 Statistically Secure Encryption

The *statistical distance* between X and Y, denoted by $\mathsf{SD}(X,Y)$, is defined by

$$\mathsf{SD}(X,Y) \stackrel{\mathsf{def}}{=} \frac{1}{2}\sum_x |\Pr[X=x] - \Pr[Y=x]|$$

For joint r.v.s (X,Z) and (Y,Z) that has a common Z, we use shorthand

$$\mathsf{SD}(\; X,Y \mid Z \;) \;\stackrel{\mathsf{def}}{=}\; \mathsf{SD}(\; (X,Z)\;,\;(Y,Z)\;) \;\;.$$

For $\mathsf{SD}(X,Y) \le \varepsilon$ we say that X is ε-close to Y, and for $\mathsf{SD}(X,Y|Z) \le \varepsilon$ we say that X is ε-close to Y conditioned on Z.

Lemma 3. *For random variables X and Y (defined over set $\mathcal{S}$), and for any distinguisher $\mathsf{D} : \mathcal{S} \to \{0,1\}$, it holds that*

$$\mid \Pr[\mathsf{D}(X)=1] \;-\; \Pr[\mathsf{D}(Y)=1] \mid \;\le\; \mathsf{SD}(X,Y) \tag{1}$$

Proof. Denote $p_s \stackrel{\mathsf{def}}{=} \Pr[X=s]$ and $q_s \stackrel{\mathsf{def}}{=} \Pr[Y=s]$. Then, we have

$$\begin{aligned}\mathsf{SD}(X,Y) \;&=\; \frac{1}{2}\Big(\sum_{s\in\mathcal{S}:p_s\ge q_s}(p_s-q_s) \;+\; \sum_{s\in\mathcal{S}:p_s<q_s}(q_s-p_s)\Big)\\ &=\; \sum_{s\in\mathcal{S}:p_s\ge q_s}(p_s-q_s)\end{aligned}$$

which is due to

$$\sum_{s\in\mathcal{S}:p_s\ge q_s}(p_s-q_s) \;-\; \sum_{s\in\mathcal{S}:p_s<q_s}(q_s-p_s) \;=\; \sum_{s\in\mathcal{S}}p_s \;-\; \sum_{s\in\mathcal{S}}q_s \;=\; 0\;\;.$$

Assume without loss of generality that $\Pr[\mathsf{D}(X)=1] \geq \Pr[\mathsf{D}(Y)=1]$, it follows that

$$\begin{aligned} & |\Pr[\mathsf{D}(X)=1] - \Pr[\mathsf{D}(Y)=1]| \\ = & \Pr[\mathsf{D}(X)=1] - \Pr[\mathsf{D}(Y)=1] \\ \leq & \sum_{s\in\mathcal{S}:p_s\geq q_s} (p_s - q_s) = \mathsf{SD}(X,Y) \ . \end{aligned}$$

□

In fact, for every X and Y there exists D (how to define such D?) for the equality of Equation (1) to hold, and thus an alternative definition of statistical distance is

$$\mathsf{SD}(X,Y) \stackrel{\text{def}}{=} \max_{\mathsf{D}} |\Pr[\mathsf{D}(X)=1] - \Pr[\mathsf{D}(Y)=1]|$$

STATISTICAL DISTANCE IS A METRIC. That is, it has the following properties.

- (non-negativity):$\mathsf{SD}(X,Y) \geq 0$.
- (identity of indiscernibles): $\mathsf{SD}(X,Y)=0$ iff X is identically distributed to Y.
- (symmetry): $\mathsf{SD}(X,Y) = \mathsf{SD}(Y,X)$.
- (triangle inequality): $\mathsf{SD}(X,Z) \leq \mathsf{SD}(X,Y) + \mathsf{SD}(Y,Z)$.

Further, it satisfies the following:

- (no greater than 1): $\mathsf{SD}(X,Y) \leq 1$ with equality holds iff X and Y have disjoint supports.
- (replacement): for every function f, it holds that $\mathsf{SD}(f(X),f(Y)) \leq \mathsf{SD}(X,Y)$.

Theorem 4 (One-time pad using statistical random keys)**.** *Let encryption scheme $\tilde{\Pi} = (\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ with $\mathcal{K}=\mathcal{M}=\mathcal{C}=\{0,1\}^n$ be the same as the one-time pad except that Gen samples k from $\tilde{K}$ which is ε-close to U_n (i.e., $\mathsf{SD}(\tilde{K},U_n)\leq\varepsilon$). Then, $\tilde{\Pi}$ is 2ε-secure (with respect to Definition 1 and Definition 2).*
We say that $\tilde{\Pi}$ is statistically secure if ε is negligible in n.

Proof. For any m_0, $m_1 \in \mathcal{M}$, and any distinguisher D

$$\begin{aligned} & |\Pr[\mathsf{D}(m_0\oplus\tilde{K})=1] - \Pr[\mathsf{D}(m_1\oplus\tilde{K})=1]| \leq \mathsf{SD}(m_0\oplus\tilde{K}, m_1\oplus\tilde{K}) \\ & \leq \mathsf{SD}(m_0\oplus\tilde{K}, m_0\oplus U_n) + \mathsf{SD}(m_0\oplus U_n, m_1\oplus\tilde{K}) \\ & = \mathsf{SD}(m_0\oplus\tilde{K}, m_0\oplus U_n) + \mathsf{SD}(m_1\oplus U_n, m_1\oplus\tilde{K}) \\ & \leq 2\varepsilon \end{aligned}$$

where the first inequality is by Lemma 3 and the second inequality is the triangle. □

3 Unpredictability, Min-entropy and The Leftover Hash Lemma

In practice, we may only have some secret (or a uniformly random key with some information leaked to the adversary) that is highly unpredictable but not indistinguishable from uniform.

Definition 4 (min-entropy and unpredictability)**.** A random variable X is ε-unpredictable if for every algorithm A it holds that

$$\Pr[\mathsf{A}(1^n) = X] \ \leq \ \varepsilon \ .$$

The min-entropy of X, denoted by $\mathbf{H}_\infty(X)$, is defined by

$$\mathbf{H}_\infty(X) \stackrel{\text{def}}{=} -\log \ (\ \max_{x \in \mathcal{X}} \ \Pr[X = x] \ )$$

where log refers to the logarithm to the base 2.

Fact 3.1. *A random variable X is ε-unpredictable iff $\mathbf{H}_\infty(X) \geq \log(1/\varepsilon)$.*

The proof is easy to see as the best guessing strategy for A is to output the x with maximal probability.

Why negative logarithm? This is for the convenience of human beings. When talking about security, people feel more comfortable to say that "this crypto-system has 128 bits of security" than to argue that "the system is 2^{-128} (or $0.0000\cdots01$) secure against any attacks."

Definition 5 (average min-entropy and conditional unpredictability)**.** For joint random variable (X,Z), we say that $\mathcal{X}$ is ε-unpredictable given Z if for every algorithm A it holds that

$$\Pr[\mathsf{A}(1^n, Z) = X] \ \leq \ \varepsilon \ .$$

The average min-entropy of X conditioned on Z, denoted by $\mathbf{H}_\infty(X \mid Z)$, is defined by

$$\mathbf{H}_\infty(X) \stackrel{\text{def}}{=} -\log \ \mathbb{E}_{z \leftarrow Z} \left(\ \max_{x \in \mathcal{X}} \ \Pr[X = x | Z = z] \right) \ .$$

Note: one may think of X as a secret key and Z as the leakage about X such that it remains ε-hard to predict the secret key even if Z is given to the adversary.

Similar to the unconditional case, it is easy to see the following fact.

Fact 3.2. *For joint random variable (X,Z), X is ε-unpredictable given Z iff $\mathbf{H}_\infty(X|Z) \geq \log(1/\varepsilon)$.*

In general, encryptions are not secure by using keys that have high min-entropy entropy. Consider the OTP pad as an example, if the key $k \in \{0,1\}^n$ is sampled from distribution K that is constant 0 at the first bit, and uniformly random at the rest positions (i.e., $\mathbf{H}_\infty(K) = n - 1$). The adversary can choose two message m_0 and m_1 whose first bits are 0 and 1 respectively, and always wins the game (since the first bit is not encrypted).

Randomness Extraction Now we introduce functions that extract statistical random bits from random variables with some non-trivial amount of min-entropy, and we call such functions *randomness extractors*. We define universal hash functions below and show that they are randomness extractors (formalized by the leftover hash lemma).

Definition 6 (universal hash functions)**.** $\mathcal{H} \subseteq \ \{h : \{0,1\}^l \to \{0,1\}^t\}$ is a family of universal hash functions if for any distinct $x_1, x_2 \in \{0,1\}^l$, it holds that

$$\Pr_{h \stackrel{\$}{\leftarrow} \mathcal{H}} \ [h(x_1) = h(x_2)] \leq \frac{1}{2^t} \ .$$

$\mathcal{H}$ is not necessarily the full set of functions from l bits to t (as otherwise $|\mathcal{H}| = 2^{t \cdot 2^l}$ and not every $h \in \mathcal{H}$ is efficiently computable). For example, $\mathcal{H} = \{h_a : h_a(x) \stackrel{\text{def}}{=} (a \cdot x)_{[t]} \mid a, x \in \{0,1\}^l\}$ is a family of universal hash functions, where a and x are interpreted as elements over $GF(2^l)$, "$\cdot$" denotes multiplication over $GF(2^l)$, and $(a \cdot x)_{[t]}$ denotes the first t bits of $(a \cdot x)$.
That is, each value of a (i.e., the description of h_a) corresponds to a function $h_a \in \mathcal{H}$ and thus $\mathcal{H}$ is a set of efficient functions with $|\mathcal{H}| = 2^l$.

FINITE FIELD ARITHMETIC The finite field with p^n elements is denoted $\mathrm{GF}(p^n)$ and is also called the Galois Field. In cryptography, we often use $\mathrm{GF}(2^n)$ and $\mathrm{GF}(p)$, where p is a prime. Here we introduce the arithmetic operations over $\mathrm{GF}(2^n)$. Every $GF(2^n)$ is defined by an irreducible[1] polynomial in the form of

$$P(x) = c_n x^n + c_{n-1} x^{n-1} + \ \ldots \ + c_0$$

For every $a \in \{0,1\}^n$, denote its binary representation by $a = a_{n-1} a_{n-2} \ldots a_0$, and thus is interpreted as polynomial:

$$a_{n-1} x^{n-1} \ + \ a_{n-2} x^{n-2} \ + \ \ldots \ + \ a_0$$

and so is it with $b = b_{n-1} b_{n-2} \ldots b_0$:

$$b_{n-1} x^{n-1} \ + \ b_{n-2} x^{n-2} \ + \ \ldots \ + \ b_0$$

Thus, denote the multiplication over $\mathrm{GF}(2^n)$ by '+', then $a + b$ yields

$$(a_{n-1} \oplus b_{n-1}) x^{n-1} \ + \ (a_{n-2} \oplus b_{n-2}) x^{n-2} \ + \ \ldots \ + \ (a_0 \oplus b_0) x^0$$

which is $(a_{n-1} \oplus b_{n-1})(a_{n-2} \oplus b_{n-2}) \ldots (a_0 \oplus b_0)$ in binary, where '$\oplus$' is addition modulo 2 (i.e., XOR). Substraction is the same as addition. For multiplication between a and b, denoted by '$a \times b$', we get

$$(a_{n-1} x^{n-1} \ + \ a_{n-2} x^{n-2} \ + \ \ldots \ + \ a_0 x^0) \times (b_{n-1} x^{n-1} \ + \ b_{n-2} x^{n-2} \ + \ \ldots \ + \ b_0 x^0)$$

which is a polynomial of degree $2(n-1)$, followed by division using $P(x)$ as the divisor, yielding a remainder polynomial (which is of degree $n-1$) as the result of $a \times b$.

A CONCRETE EXAMPLE. The AES (advanced encryption standard) uses $GF(2^8)$ with reducing polynomial being

$$P(x) = x^8 + x^4 + x^3 + x + 1$$

(01010011) $\times$ (11001010) yields (example from Wikipedia)

$$(x^6 + x^4 + x + 1)(x^7 + x^6 + x^3 + x) \ = \ x^{13} + x^{12} + x^{11} + x^{10} + x^9 + x^8 + x^6 + x^5 + x^4 + x^3 + x^2 + x$$

modulo $P(x) = x^8 + x^4 + x^3 + x + 1$ we get 1.

Lemma 5 (The Leftover Hash Lemma)**.** *For any integers $d \leq k \leq l$, let $\mathcal{H} \subseteq \{ h : \{0,1\}^l \to \{0,1\}^{k-d}\}$ be a family of universal hash functions. Then, for any random variable X defined over $\{0,1\}^l$ with min-entropy[2] $\mathbf{H}_\infty(X) \geq k$ it holds that*

$$\mathsf{SD}(\ H(X),\ U_{k-d} \mid H) \ \leq \ 2^{-\frac{d}{2}-1}$$

[1] A polynomial is irreducible if it cannot be factored into nontrivial polynomials over the same field.

[2] In fact, the lemma only requires that X has collision entropy (a weaker form of entropy than min-entropy) at least k. That is, define the collision probability $\mathsf{CP}(X) = \sum_x \Pr[X = x]^2$, and collision entropy $\mathbf{H}_2(X) = -\log(\mathsf{CP}(X))$. For any X, $\mathbf{H}_\infty(X) \geq k$ implies $\mathbf{H}_2(X) \geq k$

where H is the random variable that is uniformly distributed over all members of $\mathcal{H}$.
We refer to d as the entropy loss (the difference between the amount of entropy and the number of bits extracted), and we call H as the random seed.

Proof. We recall the Cauchy-Schwartz inequality $|\sum_i a_i b_i| \le \sqrt{(\sum_i a_i^2)\cdot(\sum_i b_i^2)}$, and we denote $\mathcal{S} \stackrel{\text{def}}{=} \{0,1\}^{k-d}$ and $p_{s|h} = \Pr[H(X) = s | H = h]$

$$\begin{aligned}
&\mathsf{SD}(\ H(X),\ U_{k-d} \mid H) \\
=\ & \sum_{h\in\mathcal{H}} \frac{1}{|\mathcal{H}|}\cdot\frac{1}{2}\sum_{s\in\mathcal{S}} |p_{s|h} - \frac{1}{|\mathcal{S}|}| \\
=\ & \frac{1}{2}\cdot \sum_{h\in\mathcal{H},s\in\mathcal{S}} \sqrt{\frac{1}{|\mathcal{S}|\cdot|\mathcal{H}|}}\cdot\left(\sqrt{\frac{|\mathcal{S}|}{|\mathcal{H}|}}\cdot\left|p_{s|h} - \frac{1}{|\mathcal{S}|}\right|\right) \\
\le\ & \frac{1}{2}\cdot \sqrt{\sum_{h\in\mathcal{H},s\in\mathcal{S}} (\frac{1}{|\mathcal{S}|\cdot|\mathcal{H}|}) \sum_{h\in\mathcal{H},s\in\mathcal{S}} \frac{|\mathcal{S}|}{|\mathcal{H}|}(p_{s|h} - \frac{1}{|\mathcal{S}|})^2} \\
=\ & \frac{1}{2}\sqrt{(|S|\sum_{h\in\mathcal{H}}\frac{1}{|\mathcal{H}|}\sum_{s\in\mathcal{S}} p_{s|h}^2) - 1}
\end{aligned}$$

We denote with $\mathsf{CP}(H(X)|H) \stackrel{\text{def}}{=} \sum_{h\in\mathcal{H}}\frac{1}{|\mathcal{H}|}\sum_{s\in\mathcal{S}} p_{s|h}^2$, which can be considered as the collision probability of $H(X_1)$ and $H(X_2)$ averaged over H (uniform over $\mathcal{H}$), where X_1 and X_2 are i.i.d. to X. We have by the universal hash functions that

$$\mathsf{CP}(H(X)|H) \ \le\ \Pr[X_1 = X_2] + \Pr_{x_1\neq x_2}[H(x_1) = H(x_2)] \ \le\ 2^{-k} + 2^{-k+d}$$

It follows that

$$\frac{1}{2}\sqrt{|S|\cdot\mathsf{CP}(H(X)|H) - 1} \ \le\ \frac{1}{2}\sqrt{2^{k-d}(2^{-k} + 2^{-k+d}) - 1} \ \le\ 2^{-\frac{d}{2}-1}$$

which completes the proof. □

Corollary 3.1. *For integers $d \le k \le l$ and $\mathcal{H} : \{0,1\}^l \to \{0,1\}^{k-d}$ be the same as assumed in Lemma 5. For any random variable (X,Z) where X is over $\{0,1\}^l$ with average min-entropy $\mathbf{H}_\infty(X|Z) \ge k$ it holds that*

$$\mathsf{SD}(\ H(X),\ U_{k-d} \mid H, Z) \ \le\ 2^{-\frac{d}{2}-1}$$

where H is the random variable that is uniformly distributed over all members of $\mathcal{H}$.

The proof of the above corollary is left as an exercise.

Applications of randomness extraction. Say we have a random physical source that produces random numbers with high min-entropy but not perfectly uniform. Then we just sample a random $h \stackrel{\$}{\leftarrow} \mathcal{H}$ and apply it to the source, and the output will be statistically close to uniform even if h is made public. Another application is privacy amplification below, Alice and Bob initially share some secret randomness W, but somehow the adversary Eve manages to learn some information (e.g. individual bits) about W, denoted by Z. Alice and Bob don't know which part of W is leaked but the only thing that can be guaranteed is that W remains of some amount (say

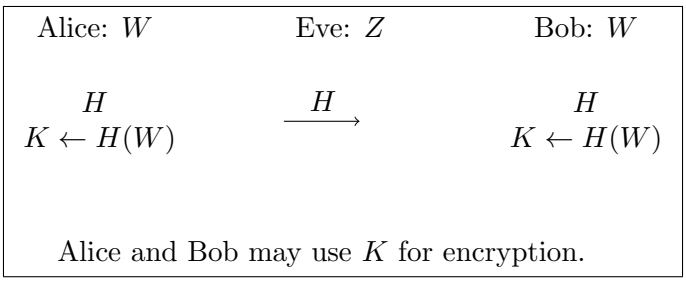

Figure 2: The problem of privacy amplification, where Alice and Bob shares a secret W, and adversary Eve learns some partial information Z correlated with W.

k) of min-entropy given Z, i.e., $\mathbf{H}_\infty(W|Z) \geq k$. They engage in the following protocol: Alice sends a random universal hash H (that output $k-d$ bits), and both parties compute $K = H(W)$. By the leftover hash lemma (Corollary 3.1), we have that K looks $2^{-d/2-1}$-close to uniform to Eve (who sees H and Z), namely,

$$\mathsf{SD}(\ K = H(W)\ ,\ U_{k-d}\ |H, Z\ )\ \leq\ 2^{-\frac{d}{2}-1}$$

Thus, by executing the protocol both parties agree on a statistically random key K and they may use it in the subsequent communications. For example, use K as a secret key in one-time pad encryption (see Theorem 4).

4 Homework 2

Exercise 1. Prove Corollary 3.1.

Exercise 2. The leftover hash lemma tells us that for any X with some amount of min-entropy, applying a random universal hash function $h \xleftarrow{\$} \mathcal{H}$ to X yields statistically random bits. Explain why it is not possible to use a single deterministic function h (instead of sampling it from a set of functions), namely, there exists no deterministic h such that $h(X)$ is statistically random for all X's with some non-trivial amount of min-entropy.
Hint: Try to come up with a counterexample for which h fails.

Exercise 3 (One-Time Message Authentication Codes)**.** Consider the following scenario where Alice and Bob share a secret $W \in \{0,1\}^{2n}$ which is partitioned into equal-length strings W_1 and W_2, and an active adversary Eve has some knowledge Z correlated with W satisfying $\mathbf{H}_\infty(W|Z) \geq n + t$ and she may modify the message transmitted over the channel. To detect Eve's tampering, for any message $m \in \{0,1\}^n$ Alice computes tag $\sigma := W_1 + m \cdot W_2$ and send m along with the tag σ to Bob, where '$+$' and '$\cdot$' denote addition and multiplication over $\mathrm{GF}(2^n)$ respectively. Upon receiving (possibly modified) message and tag (m', σ'), Bob verifies whether $\sigma' = W_1 + m' \cdot W_2$ or not, and output '$\perp$' if negative. Show that for any $m \in \{0,1\}^n$ and any adversary Eve it holds that

$$\Pr_{((w_1,w_2),z)\leftarrow(W,Z)}\ [(m', \sigma') \leftarrow Eve(\sigma, z, m) : m \neq m'\ \wedge\ \sigma' = w_1 + m' \cdot w_2]\ \leq\ 2^{-t}$$

Namely, Bob detects any tampering with probability at least $(1 - 2^{-t})$.

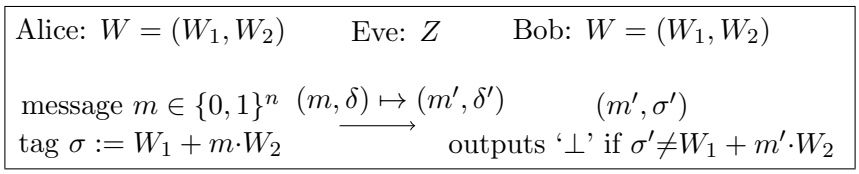

Exercise 4. For random variable X, define the collision probability $\mathsf{CP}(X) \stackrel{\text{def}}{=} \sum_x \Pr[X = x]^2$, and collision entropy $\mathbf{H}_2(X) = -\log(\mathsf{CP}(X))$. Show that for any X with $\mathbf{H}_2(X) \geq k$ and any $0 < \delta < 1$ there exists some Y with $\mathbf{H}_\infty(Y) \geq k - \log(1/\delta)$ such that $\mathsf{SD}(X, Y) \leq \delta$.