TD 1 – Introduction

Exercise 1. *One-time pad*

1. Let X, R be two independent random variables over $\{0,1\}$, with $\Pr[X=0]=p$ for some p, and $\Pr[R=0]=\frac{1}{2}$. Compute the following quantities, using the law of total probability and Bayes' formula.
   i. $\Pr[X \oplus R = 0]$
   ii. $\Pr[X \oplus R = 1]$
   iii. $\Pr[X = 0 | X \oplus R = 0]$
   iv. $\Pr[X = 0 | X \oplus R = 1]$
2. We now assume that $\Pr[R=0]=q$ for some arbitrary q. Recompute $\Pr[X = 0 | X \oplus R = 0]$.
3. Let now X, R be independent random variables over $\{0,1\}^n$, and assume R to be uniformly distributed in $\{0,1\}^n$.
   i. For arbitrary $y, z \in \{0,1\}^n$, compute $\Pr[X \oplus R = y]$ and $\Pr[X = z | X \oplus R = y]$.
   ii. Explain why knowing $X \oplus R$ does not reveal any information about X.
   iii. Let Y be another random variable over $\{0,1\}^n$. Explain why knowing $(X \| Y) \oplus (R \| R)$ *does* reveal information about $X \| Y$, where $\|$ denotes string concatenation.

Exercise 2. *One-time pad for variable length messages*

Let us consider the space $\mathcal{M} = \{0,1\}^{\leq \ell}$ of binary string of length $\leq \ell$.

1. We consider the following encryption scheme: the key is uniformly sampled from $\mathcal{K} = \{0,1\}^\ell$ and we define $\mathsf{Enc}_k(m) = k_{[0,|m|[} \oplus m$ where $k_{[0,t[}$ is made of the first t bits of k.
   i. Write the decryption algorithm.
   ii. Prove that this scheme is not perfectly secret. *First give an intuitive explanation, and then a proof using the indistinguishability experiment: describe an adversary whose advantage is nonzero.*
2. Propose a perfectly secret encryption scheme for $\mathcal{M}$. *Provide the encryption and decryption algorithms, and prove that it is perfectly secret (using the result on the one-time-pad).*

Exercise 3. *ε-indistinguishability and key lengths*

1. Consider the one-time pad for length-ℓ messages, but using a key sampled uniformly from a set $\mathcal{K}$ of size $(1-\varepsilon)2^\ell$, for $0 < \varepsilon \leq \frac{1}{2}$. Prove that this scheme is ε-indistinguishable. *Indication. Prove actually the stronger claim that the scheme is $(\varepsilon/2(1-\varepsilon))$-indistinguishable.*

We shall prove that if an encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is ε-indistinguishable, then $|\mathcal{K}| \geq (1-2\varepsilon)|\mathcal{M}|$.

2. By contrapositive, we assume $|\mathcal{K}| < (1-2\varepsilon)|\mathcal{M}|$ and define an adversary A for the experiment $\mathsf{Exp}^{\mathsf{IND}}_{\mathsf{Enc}}$. To produce m_0 and m_1, it draws them independently and uniformly from $\mathcal{M}$. Once it receives c, it checks whether there exists $k \in \mathcal{K}$ such that $\mathsf{Dec}_k(c) = m_0$. It returns 0 if this is the case, and 1 otherwise.
   i. If $b = 0$, what is the probability that A returns 0?
   ii. Assume now that $b = 1$. Bound the probability that there exists k such that $\mathsf{Dec}_k(c) = m_0$. Deduce a bound on the probability that A returns 0 in that case.
   iii. Prove that A has advantage $\geq \varepsilon$.

Exercise 4. *Secrecy and indistinguishability*

Let $(\mathsf{Enc}, \mathsf{Dec})$ be a encryption scheme. Let M, K, C be random variables describing the message, the key and the ciphertext respectively. They satisfy $C = \mathsf{Enc}_K(M)$. *We assume without loss of generality that for every $m \in \mathcal{M}$ and $c \in \mathcal{C}$, $\Pr[M=m] > 0$ and $\Pr[C=c] > 0$, that is $\mathcal{M}$ and $\mathcal{C}$ do not contain any impossible message or ciphertext.*

Recall that the scheme is perfectly secure if for any $m \in \mathcal{M}$ and $c \in \mathcal{C}$, $\Pr[M = m | C = c] = \Pr[M = m]$. This is equivalent to saying that the two random variables M and C are independent.

1. We will prove that perfect secrecy is equivalent to *perfect indistinguishability*: the distribution of $\mathsf{Enc}_K(m)$ (when K is random) does not depend on m.

   i. Prove that for any $m \in \mathcal{M}$ such that $\Pr[M = m] > 0$ and any $c \in \mathcal{C}$, $\Pr[C = c||M = m] = \Pr[\mathsf{Enc}_K(m) = c]$.
   ii. Deduce that the scheme is perfectly secret if and only if for every $m \in \mathcal{M}$ and $c \in \mathcal{C}$, $\Pr[\mathsf{Enc}_K(m) = c] = \Pr[C = c]$.
   iii. Prove that the scheme is perfectly secret if and only if for every $m, m' \in \mathcal{M}$, and $c \in C$, $\Pr[\mathsf{Enc}_K(m) = c] = \Pr[\mathsf{Enc}_K(m') = c]$.

2. We will now prove that perfect secrecy is equivalent to perfect *adversarial indistinguishability*, as defined in the course.

   i. Assume that the scheme is perfectly secret, and consider a *deterministic* adversary A: we can partition $\mathcal{C} = \mathcal{C}_0 \sqcup \mathcal{C}_1$ such that A outputs 0 if $c \in \mathcal{C}_0$ and 1 if $c \in \mathcal{C}_1$. Prove that the advantage of A in $\mathsf{Exp}^{\mathsf{IND}}_{\mathsf{Enc}}$ is exactly 0.
   ii. Prove that the results holds with a randomized adversary. *Change the viewpoint: A randomized adversary is a random choice amongst several possible deterministic adversaries.*
   iii. We want to prove the converse. For, we assume that the scheme is not perfectly secret and construct an adversary that has a nonzero advantage. Let m_0, $m_1 \in \mathcal{M}$ and $c^\star \in \mathcal{C}$ such that $\Pr[c^\star = \mathsf{Enc}_K(m_0)] > \Pr[c^\star = \mathsf{Enc}_K(m_1)]$. Consider the following adversary: It provides m_0 and m_1, and when it receives c, it outputs 0 if $c = c^\star$ and a uniform bit if $c \neq c^\star$. Prove that its advantage is nonzero.

Exercise 5. *Probability reminders*

— A (discrete) *probability space* is a pair (Ω, p) made of a finite or countable *sample space* (*a.k.a. universe*) Ω and a *probability mass function* $p : \Omega \to [0, 1]$ which associates to each *outcome* $\omega \in \Omega$ a *probability* $p(\omega)$, such that $\sum_{\omega \in \Omega} p(\omega) = 1$.
— An *event* is a subset of Ω. The probability of a event E is $\Pr[E] = \sum_{\omega \in \Omega} p(\omega)$. We use $E \wedge F$ to denote the event $E \cap F$, $E \vee F$ to denote $E \cup F$, and $\neg E$ to denote $\Omega \setminus E = \{\omega \in \Omega : \omega \notin E\}$.
— Given two events E, $F \subset \Omega$, the *conditional probability of E given F* is $\Pr[E|F] = \Pr[E \wedge F]/\Pr[F]$ (provided $\Pr[F] \neq 0$). The intuitive meaning is the probability of the event E *within the restricted universe* F: In particular, $\Pr[E] = \Pr[E|\Omega]$ for all E.
— Two events E and F are *independent* if $\Pr[E|F] = \Pr[E]$, or equivalently if $\Pr[F|E] = \Pr[F]$, or equivalently if $\Pr[E \wedge F] = \Pr[E]\Pr[F]$.
— A (discrete) *random variable* is a function $X : \Omega \to S$. Each $x \in S$ defines and *event* $[X = x] = \{\omega \in \Omega : X(\omega) = x\}$, and similarly for $[X \geq x]$, $[X < x]$, ...
— The (conditional) *expectation* of a random variable $X : \Omega \to S$ is $\mathbb{E}[X|E] = \sum_{x \in S} x \Pr[X = x|E]$. Expectation is linear: $\mathbb{E}[X + Y|E] = \mathbb{E}[X|E] + \mathbb{E}[Y|E]$. The *standard* expectation is $\mathbb{E}[X] = \mathbb{E}[X|\Omega]$.

Prove the following (*almost obvious but very useful!*) results.

1. For two events E and F,

   i. $\Pr[\neg E] = 1 - \Pr[E]$, and
   ii. $\Pr[E \vee F] = \Pr[E] + \Pr[F] - \Pr[E \wedge F] \leq \Pr[E] + \Pr[F]$. (Union bound)

2. For two events E and F,

$$\Pr[E|F]\Pr[F] = \Pr[F|E]\Pr[E] = \Pr[E \wedge F]. \qquad \text{(Bayes' formula)}$$

3. Let F_1, ..., F_n be a partition of Ω, that is $\bigcup_i F_i = \Omega$ and $F_i \cap F_j = \emptyset$ if $i \neq j$. Then,

   i. for any event E, $\Pr[E] = \sum_{i=1}^{n} \Pr[E|F_i]\Pr[F_i] = \sum_{i=1}^{n} \Pr[E \wedge F_i]$, and (Law of total probability)
   ii. for any random variable X, $\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X|F_i]\Pr[F_i]$. (Law of total expectation)

4. Let $X : \Omega \to \mathbb{N}$ be a random variable with nonnegative integer values. Then $\mathbb{E}[X] = \sum_{i \geq 1} \Pr[X \geq i]$.