Tutorial 2

Gidon Rosalki

2025-11-05

Notice: If you find any mistakes, please open an issue at `https://github.com/robomarvin1501/notes_intro_to_crypto`

1 Recap

So far, we have discussed *perfect secrecy*, and demonstrated its requirements, along with methods such as the one time pad. It came with the significant drawbacks of each key having to be single use, and it needing exceedingly large keys (at least the length of the plaintext). We went on to discuss a slight reduction of the security provided by this with *indistinguishable encryption* (IND), which makes use of pseudo random generators, that take a smaller key, and output a much larger pseudo random output. This can be used with a one time pad, thus allowing shorter keys, but each key can still only be used once.
We also defined *semantic security*, a powerful mathematical definition that should be reviewed in last weeks notes. We also defined *one way functions*, which we will not use extensively, but may well return to come the end of the semester.

2 Pseudorandom Generators (PRGs)

Definition 2.1 (PRG)**.** *A PRG is a polytime computable function* $G : \{0,1\}^n \to \{0,1\}^{l(n)}$ *for* $l(n) > n$.

PRGs enable that for every PPT D, there exists $\mathcal{V}$, a negligible function $\forall n$ such that

$$\left| \Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] \right| < \mathcal{V}(n)$$

Exercise 1. *Given a PRG* G*, such that* $\forall n \in \mathbb{N},\ s \in \{0,1\}^n$ *it holds that* $G(s) \in \{0,1\}^{l(n)}$ *for some function* $l : l(n) > n$*. Show that there exists a negligible function* $\mathcal{V}(\cdot)$ *such that*

$$\frac{1}{8} - \mathcal{V}(n) \leq \Pr_{s \leftarrow \{0,1\}^n} [G(s) \text{ starts with } 000] \leq \frac{1}{8} + \mathcal{V}(n)$$

Solution. We prove by contradiction. Assume there exists a polynomial $p(\cdot)$ such that for infinitely many n,

$$\Pr_{s \leftarrow \{0,1\}^n} [G(s) \text{ starts with } 000] > \frac{1}{8} + \frac{1}{p(n)}$$

or

$$\Pr_{s \leftarrow \{0,1\}^n} [G(s) \text{ starts with } 000] < \frac{1}{8} - \frac{1}{p(n)}$$

Without loss of generality, we will assume that the first inequality is true, for $p(\cdot)$, and infinite ns. We will construct a distinguisher D, such that

$$D\left(y \in \{0,1\}^*\right) = \begin{cases} 1, & \text{if } y \text{ starts with } 000 \\ 0, & \text{else} \end{cases}$$

We will note that D is a PPT. Since in a string of length $l(n)$, each bit is sampled randomly from $\{0,1\}$, we know that

$$\Pr_{r \leftarrow \{0,1\}^{l(n)}} [r \text{ starts with } 000] = \frac{1}{2^3} = \frac{1}{8}$$

Now, D's distinguishing advantage between random strings, and the output of G may be calculated as follows: By the definition of D, and according to the contradicting assumption, for enough values of $n \in \mathbb{N}$ it holds that:

$$\begin{aligned} \left| \Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] \right| &> \left| \frac{1}{8} + \frac{1}{p(n)} - \frac{1}{8} \right| \\ &\geq \frac{1}{p(n)} \end{aligned}$$

which contradicts the assumption that G is a PRG. □

We will note that in this proof we used 2 facts that are related to the start of the string beginning with 000:

1. This may be checked efficiently
2. This property has a probability of $\frac{1}{8}$ of occurring, for a string that is chosen from a uniform distribution

Conclusion: Every feature on binary strings that can be efficiently tested exists for the output of a PRG with approximately the same probability as it exists in a string sampled from the uniform distribution (the difference between the probabilities is bounded b a negligible function).

Exercise 2. *Let there be G_1, G_2 PRGs. We will define*

$$G(s) = G_1(s) \,||\, G_2(s)$$

Where the double line is string joining. Prove or disprove that G is a PRG.

Solution. This is not the case. Consider the case where $G_1 = G_2$. In this case, the output of G will not be pseudo-random, since we can construct a distinguisher that checks if the first half of the string is equal to the second half: More formally, let there be $H = G_1 = G_2$. Therefore, $G(s) = H(s) \,||\, H(s)$. We can theorise that for all H, G is not a PRG. Let us construct $D(z)$:

1. Check if the first and second half are the same.
2. If yes, return 1
3. Else, return 0

For G,

$$\Pr[D(G(s)) = 1] = 1$$

However,

$$\Pr[D(s) = 1] = \frac{1}{2^{\frac{n}{2}}}$$

As a result, G is **not** a PRG. □

3 Indistinguishable proofs

Let there be $\Pi = (KeyGen, Enc, Dec)$, an encryption system. For every algorithm A, and $\forall n \in \mathbb{N}$, we will define the experiment $IND_{\Pi,A}(n)$, as follows:

1. $k \leftarrow KeyGen(1^n)$
2. A receives 1^n as input, and outputs a pair of messages (m_0, m_1)
3. $b \leftarrow \{0,1\}$, and we compute $c^* \leftarrow Enc_k(m_b)$, and pass c^* to A
4. A returns $b' \in \{0,1\}$

$IND_{\Pi,A}(n) = 1$ if $b' = b$. Otherwise, $IND_{\Pi,A}(n) = 0$. We will say that Π is an indistinguishable encryption system if for every PPT function A, there exists a negligible function $v(\cdot)$, such that

$$\Pr[IND_{\Pi,A}(n) = 1] \leq \frac{1}{2} + v(n)$$

for every $n \in \mathbb{N}$.

Exercise 3. *Let there be $l(n) : \forall n \in \mathbb{N}\ l(n) > n$. Let there be a determinstic PT G, such that for all $n \in \mathbb{N}$, and for all $s \in \{0,1\}^n$, it holds that $G(s) \in \{0,1\}^{l(n)}$. We will consider the system $\Pi = (KeyGen, Enc, Dec)$, defined as follows:*

1. $k \leftarrow KeyGen(1^n)$
2. $Enc_k(m) \implies c = G(k) \oplus m$
3. $Dec_k(c) \implies m = c \oplus G(k)$

It is given that Π is IND-secure. Is G necessarily a PRG?

Solution. G is necessarily a PRG. Let us assume the contradiction that it is not: There exists a PPT D, and there exists the polynomial $p(\cdot)$, such that

$$\left| \Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] \right| > \frac{1}{p(n)}$$

In this case, it will hold that Π is not IND-secure. Let us assume that there are infinite values of n, that enable the above inequality, but without the absolute value:

$$\Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] > \frac{1}{p(n)}$$

Let us consider the algorithm A that partakes in the experiment $IND_{\Pi,A}(n)$, defined as follows:

1. A generates $m_0 \leftarrow \{0,1\}^{l(n)}$, and $m_1 \leftarrow 0^{l(n)}$.
2. A receives c^*, for m_b, and runs $D(c^*)$, and returns 1 **if and only if** D returns 1

When A returns 0, then it is also true that $c^* = G(k) \oplus m_0$ is distributed evenly over $\{0,1\}^{l(n)}$. Therefore:

$$\Pr[IND_{\Pi,A}(n) = 1 | b = 0] = 1 - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1]$$

Therefore, when A returns 1, it holds that $c^* = G(k) \oplus 0^{l(n)} = G(k)$, where $k \leftarrow \{0,1\}^n$. Therefore

$$\Pr[IND_{\Pi,A}(n) = 1 | b = 1] = \Pr_{s \leftarrow \{0,1\}^n} [D(r) = 1]$$

Overall, we get for infinite values of $n \in \mathbb{N}$, it holds that

$$\begin{aligned}
\Pr[IND_{\Pi,A}(n) = 1] &= \Pr[IND_{\Pi,A}(n) = 1 | b = 0] \cdot \Pr[b = 0] + \Pr[IND_{\Pi,A}(n) = 1 | b = 1] \cdot \Pr[b = 1] \\
&= \Pr[D(G(\text{PRG})) = 1] \cdot \frac{1}{2} + (1 - \Pr[D(\text{random}) = 0]) \cdot \frac{1}{2} \\
&= \frac{1}{2} \left(\Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] + 1 - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] \right) \\
&\geq \frac{1}{2} + \frac{1}{2p(n)} \\
&> \frac{1}{2} + \frac{1}{p'(n)}
\end{aligned}$$

The final inequalities follows from our initial assumption:

$$\Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] - \Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] > \frac{1}{p(n)}$$

This is a contradiction to the given that Π is IND-secure. Therefore, G is a PRG. If we consider the other direction for our assumption (because of the absolute value), as in:

$$\Pr_{r \leftarrow \{0,1\}^{l(n)}} [D(r) = 1] - \Pr_{s \leftarrow \{0,1\}^n} [D(G(s)) = 1] > \frac{1}{p(n)}$$

Then we may use the same proof, but swapping A to return the opposite of D, i.e., D returns 0, A returns 1. □

Exercise 4. *Let there be $X = \{X_n\}_{n \in \mathbb{N}}$ and $Y = \{Y_n\}_{n \in \mathbb{N}}$, efficiently computable distributions that are computationally indistinguishable.*

1. *Prove that for every PT function f, the distributions $f(Y) = \{f(Y_n)\}_{n \in \mathbb{N}}$ and $f(X) = \{f(X_n)\}_{n \in \mathbb{N}}$ are computationally indistinguishable.*
2. *Is the theorem in part 1 still true if f is not PT?*

Solution. 1. We will assume the contradiction that there exists a PPT function f, for which the distributions $f(X)$ and $f(Y)$ are distinguishable. Then, there exists a PPT distinguisher A such that for infinitely many n:

$$\left| \Pr_{x \leftarrow X_n} [A(1^n, f(x)) = 1] - \Pr_{y \leftarrow Y_n} [A(1^n, f(y)) = 1] \right| > \frac{1}{p(n)}.$$

Let us define a new distinguisher $D(1^n, z) = A(1^n, f(z))$. We will note that since A is a PPT, then so too is D. Therefore, for every $n \in \mathbb{N}$ it holds that:

$$\Pr_{x \leftarrow X_n} [D(1^n, x)] = \Pr_{x \leftarrow X_n} [A(1^n, f(x)) = 1]$$

and also

$$\Pr_{y \leftarrow Y_n}\left[D\left(1^n, y\right)\right] = \Pr_{y \leftarrow Y_n}\left[A\left(1^n, f\left(y\right)\right) = 1\right]$$

Therefore, for infinite values of $n \in \mathbb{N}$

$$\left|\Pr_{x \leftarrow X_n}\left[D\left(1^n, x\right) = 1\right] - \Pr_{y \leftarrow Y}\left[D\left(1^n, y\right) = 1\right]\right| > \frac{1}{p\left(n\right)}$$

So D is a PPT algorithm, that distinguishes between X and Y, with non negligible probability for infinite values of $n \in \mathbb{N}$, which is a contradiction to the assumption that these two distributions are indistinguishable.

2. It is not. Let there be G, a PRG, such that for every $n \in \mathbb{N}$, $G : \{0,1\}^n \to \{0,1\}^{n+1}$. We will look at the following example:

   - $\forall n \in \mathbb{N}$ the distribution X_n takes a random $s \leftarrow \{0,1\}^n$, and returns $G\left(s\right)$
   - $\forall n \in \mathbb{N}$ the distribution X_n takes a random $r \leftarrow \{0,1\}^{n+1}$, and returns r
   - $\forall n \in \mathbb{N}$, given the input $y \in \{0,1\}^{n+1}$, the function f checks if y is in the image of G (as in, if there exists $s \in \{0,1\}^n : y = G\left(s\right)$). If so, $f\left(y\right) = 1$, otherwise $f\left(y\right) = 0$

   Firstly, we will note that since G is a PRG, it holds that the distributions $X = \{X_n\}_{n \in \mathbb{N}}$, and $Y = \{Y\}_{n \in \mathbb{N}}$ are indistinguishable. Now, we will show that $f\left(X\right)$ and $f\left(Y\right)$ are distinguishable. Let there be D, an algorithm that given input $(1^n, b)$, where $b \in \{0,1\}$, returns the bit b. For every $n \in \mathbb{N}$ it holds that:

$$\begin{aligned}
&\left|\Pr_{x \leftarrow X_n}\left[D\left(1^n, f\left(x\right)\right) = 1\right] - \Pr_{y \leftarrow Y_n}\left[D\left(1^n, f\left(y\right) = 1\right)\right]\right| \\
&= \left|\Pr_{s \leftarrow \{0,1\}^n}\left[D\left(1^n, f\left(G\left(s\right)\right)\right) = 1\right] - \Pr_{r \leftarrow \{0,1\}^{n+1}}\left[D\left(1^n, f\left(r\right)\right) = 1\right]\right| \\
&= \Bigg|\Pr_{s \leftarrow \{0,1\}^{n+1}}[D(1^n, 1) = 1] - \Pr_{r \leftarrow \{0,1\}^{n+1}}[r \in \text{Image}(G)] \cdot \Pr_{r \leftarrow \{0,1\}^{n+1}}[D(1^n, f(r)) = 1 \mid r \in \text{Image}(G)] \\
&\quad - \Pr_{r \leftarrow \{0,1\}^{n+1}}[r \notin \text{Image}(G)] \cdot \Pr_{r \leftarrow \{0,1\}^{n+1}}[D(1^n, f(r)) = 1 \mid r \notin \text{Image}(G)]\Bigg| \\
&= \left|1 - \Pr_{r \leftarrow \{0,1\}^{n+1}}\left[r \in Image\left(G\right)\right] \cdot 1 - \Pr_{r \leftarrow \{0,1\}^{n+1}}\left[r \notin Image\left(G\right)\right] \cdot 0\right| \\
&= \left|1 - \frac{\left|Image\left(G\right)\right|}{2^{n+1}}\right| \\
&\geq \frac{1}{2}
\end{aligned}$$

□