ETH Zürich, D-INFK
Spring 2018

Prof. Ueli Maurer
Christian Badertscher
Fabio Banfi

Cryptography Foundations

Solution Exercise 1

1.1 Variants of the CPA Game for Symmetric Encryption Schemes

a) i. The idea is to construct a distinguisher D' for the bit-guessing problem (S_t^{ind}, B). D' internally runs the assumed distinguisher D and emulates a view towards D that is identical to the interaction that D would have in the bit-guessing problem (S_t^{rch}, B). This way, the decision of D will help D' to guess the bit.

For this, D' first forwards every message queried by D during phase 2 to the challenger of S_t^{ind} and returns the obtained encryptions back to D.

Once D reaches phase 3 and provides a challenge message m, D' samples a uniformly random message $\widetilde{m}$ of length $|m|$ and provides as challenge the two messages m and $\widetilde{m}$ to S_t^{ind}. Upon receiving the encryption c (of either m or $\widetilde{m}$), D' forwards this to D. Should D query any other message during phase 5, those will be again forwarded by D', and the encryptions thereof given back to D. Finally, D' issues as its guess, which we denote by Z', the bit that D issues in the emulated interaction, which we denote by Z.

This emulation by D' perfectly mimics (S_t^{rch}, B) towards D, therefore we have

$$\Lambda^{D'}\big((S_t^{\mathsf{ind}}, B)\big) = 2{\cdot}\Pr^{D'(S_t^{\mathsf{ind}},B)}[Z' = B]{-}1 = 2{\cdot}\Pr^{D(S_t^{\mathsf{rch}},B)}[Z = B]{-}1 = \Lambda^{D}\big((S_t^{\mathsf{rch}}, B)\big).$$

ii. The idea is to construct a distinguisher D' for (S_t^{rch}, B) which internally runs the distinguisher D. We again emulate a view towards D such that its guess helps D' to guess the bit B.

For this, D' first forwards every message queried by D during phase 2 to S_t^{rch}, and returns the encryptions back to D. Once D reaches phase 3 and provides a pair of challenge messages m_0 and m_1, D' samples a uniformly random bit $\widetilde{B}$ and provides as challenge the message $m_{\widetilde{B}}$ to S_t^{rch}. Upon receiving the encryption c (of either m_0 or m_1), D' forwards this to D. Should D query any other message during phase 5, those will be again forwarded by D' and the encryptions thereof given back to D.

Finally, D' issues as its guess the bit defined by $Z' := Z \oplus \widetilde{B}$, where Z is the bit guessed by D. This means that $Z' = 0$ if $Z = \widetilde{B}$ and $Z' = 1$ otherwise..

Two important observations have to be made here:

1) It is clear that if the challenger's choice in S_t^{rch} is $B = 0$ (the challenge message is really encrypted), then D' perfectly emulates the bit-guessing problem $(S_t^{\mathsf{ind}}, \widetilde{B})$ towards D.
2) On the other hand, if $B = 1$ (the challenge message is ignored and a uniformly random message is encrypted instead), then D guesses $\widetilde{B}$ with probability only $\frac{1}{2}$. To see this, note that any value given to D in this interaction is independent of the value $\widetilde{B}$.

Therefore, we have

$$\begin{aligned}
\Lambda^{D'}\big((S_t^{\mathsf{rch}}, B)\big) &= 2 \cdot \Pr^{D'(S_t^{\mathsf{rch}},B)}[Z' = B] - 1 \\
&= 2 \cdot \Pr^{D'(S_t^{\mathsf{rch}},B)}[Z \oplus \widetilde{B} = B] - 1 \\
&= 2 \cdot \Big(\Pr^{D'(S_t^{\mathsf{rch}},B)}[Z \oplus \widetilde{B} = B \,|\, B = 0] \cdot \frac{1}{2} \\
&\qquad + \underbrace{\Pr^{D'(S_t^{\mathsf{rch}},B)}[Z \oplus \widetilde{B} = B \,|\, B = 1]}_{=\frac{1}{2}} \cdot \frac{1}{2} \Big) - 1 \\
&= 2 \cdot \Pr^{D'(S_t^{\mathsf{rch}},B)}[Z = \widetilde{B} \,|\, B = 0] \cdot \frac{1}{2} + 2 \cdot \frac{1}{2} \cdot \frac{1}{2} - 1 \\
&= \Pr^{D(S_t^{\mathsf{ind}},B)}[Z = B] - \frac{1}{2} \\
&= \frac{1}{2} \Lambda^{D}\big((S_t^{\mathsf{ind}}, B)\big).
\end{aligned}$$

b) i. The solution to this case follows exactly the same idea as in **a) ii.**
D' forwards every message queried by D during phase 2 or phase 5 the challenger of S_{t+1}^{ror}, and returns the encryptions back to D.
Once D reaches phase 3 and provides a challenge message m, D' samples a uniformly random bit $\widetilde{B}$ and if $\widetilde{B} = 0$ it provides as challenge the message m to S_{t+1}^{ror}, otherwise it provides a uniformly sampled message $\widetilde{m}$ of length $|m|$. Upon receiving the encryption c (of either m or some uniformly random message), D' forwards this to D. Finally, D' returns $Z' := Z \oplus \widetilde{B}$ where Z is the guess issued by D. It is again clear that if the bit B chosen in S_{t+1}^{ror} is $B = 0$ (the challenge message is really encrypted), then D' perfectly emulates the guessing problem $(S_t^{\mathsf{rch}}, \widetilde{B})$ towards D, whereas if $B = 1$ the success probability of D in guessing $\widetilde{B}$ is at most $\frac{1}{2}$. Therefore, we again conduct the same steps as in **a) ii.** to obtain

$$\begin{aligned}
\Lambda^{D'}\big((S_{t+1}^{\mathsf{ror}}, B)\big) &= 2 \cdot \Pr^{D'(S_{t+1}^{\mathsf{ror}},B)}[Z' = B] - 1 \\
&= 2 \cdot \Pr^{D'(S_{t+1}^{\mathsf{ror}},B)}[Z \oplus \widetilde{B} = B] - 1 \\
&= \Pr^{D(S_t^{\mathsf{rch}},B)}[Z = B] - \frac{1}{2} \\
&= \frac{1}{2} \Lambda^{D}\big((S_t^{\mathsf{rch}}, B)\big).
\end{aligned}$$

ii. The idea is to construct a distinguisher D' for $(S_{t-1}^{\mathsf{rch}}, B)$ which internally runs the distinguisher D so that the view of the latter is in part the same as if it was interacting with (S_t^{ror}, B). For this, D' first samples T uniformly at random from $\{1, \ldots, t\}$. Then for each of the first $T-1$ queries, it ignores the message m provided by D, and provides instead a uniformly random message $\widetilde{m}$ of length $|m|$ as query to $(S_{t-1}^{\mathsf{rch}}, B)$, and returns the encryptions back to D. It then forwards the T-th query message m_T provided by D as challenge message for $(S_{t-1}^{\mathsf{rch}}, B)$, and upon receiving the encryption c (of either m_T or some other uniformly random message), D' forwards this to D. D' proceeds by normally forwarding all remaining queries, giving back the encryptions thereof to D. Finally, D' returns the same bit Z as D.
For the formal analysis, let $S_{t-1}^{\mathsf{rch}\text{-}b}$ be the same as S_{t-1}^{rch}, but where the bit B is fixed to the value b, and analogously for $S^{\mathsf{ror}\text{-}b}$. Also define hybrid systems H_τ which consists of the three steps between the challenger and the adversary:

1. The challenger chooses a key k according to the key distribution.
2. The adversary can choose up to t messages; for the i-th message m, the challenger acts as follows:
   - If $i < \tau$, the challenger chooses a uniformly random message $\widetilde{m}$ of length $|m|$ and computes the encryption of $\widetilde{m}$, i.e., $\widetilde{c} = e(\widetilde{m}, k, \widetilde{r})$ for fresh and independent randomness value $\widetilde{r} \in \mathcal{R}$, and returns $\widetilde{c}$ to the adversary.
   - If instead $i \geq \tau$, it computes the encryption of m, i.e., $c = e(m, k, r)$ for fresh and independent randomness value $r \in \mathcal{R}$, and returns c to the adversary.
3. The adversary issues a guess Z.

We observe the following:

- The system emulated towards D by D' when interacting with $(S_{t-1}^{\mathsf{rch}\text{-}0}, B)$ and with $T = \tau$ is the same as the system H_τ, therefore
$$\Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}0}}[Z = z \,|\, T = \tau] = \Pr^{D H_\tau}[Z = z].$$
- The system emulated towards D by D' when interacting with $(S_{t-1}^{\mathsf{rch}\text{-}1}, B)$ and with $T = \tau$ is the same as the system $H_{\tau+1}$, therefore
$$\Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}1}}[Z = z \,|\, T = \tau] = \Pr^{D H_{\tau+1}}[Z = z].$$
- The system H_1 is the same as the system $S_t^{\mathsf{ror}\text{-}0}$, therefore
$$\Pr^{D H_1}[Z = z] = \Pr^{D S_t^{\mathsf{ror}\text{-}0}}[Z = z].$$
- The system H_{t+1} is the same as the system $S_t^{\mathsf{ror}\text{-}1}$, therefore
$$\Pr^{D H_{t+1}}[Z = z] = \Pr^{D S_t^{\mathsf{ror}\text{-}1}}[Z = z].$$

Therefore we have

$$\begin{aligned}
\Lambda^{D'}\big((S_{t-1}^{\mathsf{rch}}, B)\big) &\overset{(1)}{=} \Delta^{D'}(S_{t-1}^{\mathsf{rch}\text{-}0}, S_{t-1}^{\mathsf{rch}\text{-}1}) \\
&= \Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}1}}[Z = 1] - \Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}0}}[Z = 1] \\
&= \sum_{\tau=1}^{t} \Big(\Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}1}}[Z = 1 \,|\, T = \tau] \cdot \Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}1}}[T = \tau] \\
&\qquad - \Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}0}}[Z = 1 \,|\, T = \tau] \cdot \Pr^{D' S_{t-1}^{\mathsf{rch}\text{-}0}}[T = \tau] \Big) \\
&= \frac{1}{t} \sum_{\tau=1}^{t} \big(\Pr^{D H_{\tau+1}}[Z = 1] - \Pr^{D H_\tau}[Z = 1] \big) \\
&= \frac{1}{t} \sum_{\tau=1}^{t} \Delta^{D}(H_\tau, H_{\tau+1}) \\
&\overset{(2)}{=} \frac{1}{t} \Delta^{D}(H_1, H_{t+1}) \\
&= \frac{1}{t} \big(\Pr^{D H_{t+1}}[Z = 1] - \Pr^{D H_1}[Z = 1] \big) \\
&= \frac{1}{t} \Big(\Pr^{D S_t^{\mathsf{ror}\text{-}1}}[Z = 1] - \Pr^{D S_t^{\mathsf{ror}\text{-}0}}[Z = 1] \Big) \\
&= \frac{1}{t} \Delta^{D}(S_t^{\mathsf{ror}\text{-}0}, S_t^{\mathsf{ror}\text{-}1}) \overset{(1)}{=} \frac{1}{t} \Lambda^{D}\big((S_{t-1}^{\mathsf{ror}}, B)\big),
\end{aligned}$$

where for (1) we used Lemma 2.3 and for (2) we used Lemma 2.2.
Note that what we do here is known in the literature as *hybrid argument*.

c) We would like to phrase the answer very informally to convey the taste of statements proven above.

Recall that in most parts of cryptography, we do not have absolute security, i.e., we cannot conclude that something is secure without any assumptions. This is why many statements are actually implication statements of the form: "If A is true, then B is true", that relate two problems A and B.

Consider sub-task **a) i.**. The statement can be understood as follows: "If a scheme is IND-CPA secure, then it is also RCH-CPA secure". This implication is then proven using the technique called *indirect proof of an implication*, i.e., by assuming there is an arbitrary attacker for RCH-CPA and translating it into an attacker for IND-CPA, such that the performance of the new attacker is equal to the performance of the attacker for RCH-CPA. This illustrates that "solving the RCH-CPA problem" is at least as hard as "solving the IND-CPA problem". Hence, if IND-CPA is hard to solve (for a given scheme) then so is RCH-CPA (for that scheme).

Consider sub-task **a) ii.**. This can be understood as follows: "If a scheme is RCH-CPA secure, then it is also IND-CPA secure". Again, this is proven by assuming there is an attacker against the IND-CPA game and translating it into an attacker against RCH-CPA, such that the performance of the new attacker is *roughly* equal to the performance of the attacker against IND-CPA. This also says that solving the IND-CPA problem is at least as hard as solving the RCH-CPA problem.

Putting both conclusions together, this means that both problems (for an encryption scheme) are roughly *equally hard* to solve for an adversary D. Note that this statement is not absolute, we have only put two problems in relation to each other.

The analogous interpretations hold for **1b)** as well.

1.2 On the Security of the One-Time Pad

First note that for any $g, h \in \mathbb{G}$ we have

$$\begin{aligned}\Pr[U + X = g, X = h] &= \Pr[U = g + (-h), X = h] \\ &= \Pr[U = g + (-h)] \cdot \Pr[X = h] \\ &= \frac{1}{|\mathbb{G}|} \Pr[X = h], \end{aligned} \tag{1}$$

where in the second step we used that U and X are independent, and in the last step that U is uniformly distributed. Then by the law of total probability, for any $g \in \mathbb{G}$ we have

$$\Pr[U + X = g] = \sum_{h \in \mathbb{G}} \Pr[U + X = g, X = h] \overset{(1)}{=} \frac{1}{|\mathbb{G}|} \sum_{h \in \mathbb{G}} \Pr[X = h] = \frac{1}{|\mathbb{G}|}. \tag{2}$$

Hence, for any $g, h \in \mathbb{G}$,

$$\Pr[U + X = g, X = h] \overset{(1,2)}{=} \Pr[U + X = g] \cdot \Pr[X = h],$$

which is exactly the definition of $U + X$ and X being independent.

Note that the proof of security for the one-time pad over bitstrings of length n as introduced in the lecture notes is simply the instantiation of the above where U is the key, X is the message, $\mathbb{G}$ is the set $\{0,1\}^n$, the operation $+$ is the bit-wise XOR, the inverse operation is the identity function on $\{0,1\}^n$, and the neutral element is 0^n (the bitstring consisting of n zeros).

1.3 Properties of the Distinguishing Advantage

a) The core of this statement is to argue about the optimal distinguisher for random variables X and Y i.e., their associated distributions P_X and P_Y. Once we have found the optimal distinguisher, we can compute its advantage and conclude the statement. Not surprisingly, classifying two distributions in an optimal way follows the well-known *maximum likelihood* strategy. Intuitively, a distinguisher decides that a sample x comes from distribution P_X if and only if x is more likely under P_X than under P_Y. We give a detailed derivation below to make this intuition formal.

First, note that the distinguisher always has to output 0 or 1 based on an observed sample x. We write $d : \mathcal{X} \times \mathcal{R} \to \{0,1\}$, where $\mathcal{X}$ is the range of X and Y, to denote the (possibly probabilistic) strategy of the distinguisher (where the probabilistic aspect is captured as usual as an independent random variable R distributed over the randomness space $\mathcal{R}$). Let us introduce the following bit-guessing problem (S_U, U): U is a uniformly random bit and $S_0 := X$ and $S_1 := Y$, and Z denotes the usual output of the distinguisher. We first compute the success probability of an arbitrary distinguisher that decides based on the strategy d.

$$\begin{aligned}
\Pr^{D(S_U,U)}[Z = U] &= \frac{1}{2}\Pr[d(X,R) = 0] + \frac{1}{2}\Pr[d(Y,R) = 1] \\
&= \frac{1}{2}\sum_{x\in\mathcal{X}} \mathsf{P}_X(x) \cdot \Pr[d(x,R) = 0] + \frac{1}{2}\sum_{x\in\mathcal{X}} \mathsf{P}_Y(x) \cdot \Pr[d(x,R) = 1] \qquad (3)\\
&= \frac{1}{2}\sum_{x\in\mathcal{X}} (\mathsf{P}_X(x) \cdot \Pr[d(x,R) = 0] + \mathsf{P}_Y(x) \cdot \Pr[d(x,R) = 1]).
\end{aligned}$$

Since the distinguisher has to decide on either 0 or 1, we have that $\forall x : \Pr[d(x,R) = 0] + \Pr[d(x,R) = 1] = 1$ over the randomness of the strategy. We can therefore conclude that $\frac{1}{2}\sum_{x\in\mathcal{X}} \mathsf{P}_X(x) \cdot \Pr[d(x,R) = 0] + \mathsf{P}_Y(x) \cdot \Pr[d(x,R) = 1] \leq \frac{1}{2}\sum_{x\in\mathcal{X}} \max\{\mathsf{P}_X(x), \mathsf{P}_Y(x)\}$. This inequality is very crucial. It tells us that the success of the best strategy is upper bounded. We further see that the following *deterministic* strategy $d(x,R) := d_{\text{det}}(x)$ exactly achieves this upper bound and hence is optimal: define $\mathcal{X}^\star := \{x \in \mathcal{X} \mid \mathsf{P}_Y(x) \geq \mathsf{P}_X(x)\}$ and define $d_{\text{det}}(x)$ to be 1 if and only if $x \in \mathcal{X}^\star$.

In the following, let D be the distinguisher that on input a value $x \in \mathcal{X}$ outputs $d_{\text{det}}(x)$.

$$\begin{aligned}
&\Delta^D(X,Y) = \Delta^D(S_0,S_1) \overset{\text{(Lemma 2.3)}}{=} \Lambda^D\big((S_U,U)\big) = 2 \cdot \Pr^{D(S_U,U)}[Z = U] - 1 \\
&\overset{(1)}{=} \left(\sum_{x\in\mathcal{X}\setminus\mathcal{X}^\star} \mathsf{P}_X(x) + \sum_{x\in\mathcal{X}^\star} \mathsf{P}_Y(x)\right) - 1 = \Pr[Y \in \mathcal{X}^\star] - \Pr[X \in \mathcal{X}^\star].
\end{aligned}$$

Where (1) follows by plugging our concrete strategy $d_{\text{det}}(\cdot)$ into Equation 3, and the last step follows by the complement relation $\sum_{x\in\mathcal{X}\setminus\mathcal{X}^\star} \mathsf{P}_X(x) = 1 - \sum_{x\in\mathcal{X}^\star} \mathsf{P}_X(x)$.

As a final step, we can conduct a straightforward calculation that this advantage is identical to the statistical distance which will conclude the proof. We thereby use the facts

that for any $\alpha \in \mathbb{R}$ we have $|\alpha| = \alpha$ if $\alpha \geq 0$, whereas $|\alpha| = -\alpha$ whenever $\alpha < 0$.

$$\begin{aligned}
\delta(X,Y) &= \frac{1}{2}\sum_{x\in\mathcal{X}} |\mathsf{P}_X(x) - \mathsf{P}_Y(x)| \\
&= \frac{1}{2}\sum_{x\in\mathcal{X}^\star} (\mathsf{P}_Y(x) - \mathsf{P}_X(x)) + \frac{1}{2}\sum_{x\in\mathcal{X}\setminus\mathcal{X}^\star} (\mathsf{P}_X(x) - \mathsf{P}_Y(x)) \\
&= \frac{1}{2}\left(\Pr[Y \in \mathcal{X}^\star] - \Pr[X \in \mathcal{X}^\star]\right) + \frac{1}{2}\left(\Pr[X \notin \mathcal{X}^\star] - \Pr[Y \notin \mathcal{X}^\star]\right) \\
&= \frac{1}{2}\left(\Pr[Y \in \mathcal{X}^\star] - \Pr[X \in \mathcal{X}^\star]\right) + \frac{1}{2}\left(1 - \Pr[X \in \mathcal{X}^\star] - (1 - \Pr[Y \in \mathcal{X}^\star])\right) \\
&= \Pr[Y \in \mathcal{X}^\star] - \Pr[X \in \mathcal{X}^\star].
\end{aligned}$$

b) First we define the new distinguisher D' which internally uses the distinguisher D: it simply provides D with access to S and gives D a uniformly sampled bit U. Once D outputs its decision bit Z, D' outputs its guess bit $Z' := Z \oplus U$, that is, D' outputs 0 if and only if $Z = U$ (see Figure 1). Then, using the law of total probability on the partition defined by the (disjoint) events $U = B$ and $U \neq B$ (each clearly occurring with probability $\frac{1}{2}$), we have

$$\begin{aligned}
\Lambda^{D'}\big((S,B)\big) &= 2 \cdot \Pr^{D'S}[Z' = B] - 1 \\
&= 2 \cdot \Pr^{D'S}[Z \oplus U = B] - 1 \\
&= 2 \cdot \Pr^{D'S}[Z = U \oplus B \,|\, U = B] \cdot \frac{1}{2} \\
&\qquad + 2 \cdot \Pr^{D'S}[Z = U \oplus B \,|\, U \neq B] \cdot \frac{1}{2} - 1 \\
&= \Pr^{D(S,B)}[Z = 0] + \Pr^{D(S,\overline{B})}[Z = 1] - 1 \\
&= \Pr^{D(S,\overline{B})}[Z = 1] - \Pr^{D(S,B)}[Z = 1] \\
&= \Delta^{D}\big((S,B),(S,\overline{B})\big).
\end{aligned}$$

Recall from the lecture that we consider here the problem of distinguishing pairs (S,B) and (S,U) where (S,U) stands for a pair that is sampled like (S,B), but where the distinguisher does not see the bit B (correlated with S), but an independent and uniformly distributed bit U. We again apply the law of total probability on the event whether B and U are identical and conclude:

$$\begin{aligned}
\Delta^{D}\big((S,B),(S,U)\big) &= \Pr^{D(S,U)}[Z = 1] - \Pr^{D(S,B)}[Z = 1] \\
&= \Pr^{D(S,U)}[Z = 1 \,|\, U = B] \cdot \frac{1}{2} \\
&\qquad + \Pr^{D(S,U)}[Z = 1 \,|\, U \neq B] \cdot \frac{1}{2} - \Pr^{D(S,B)}[Z = 1] \\
&= \frac{1}{2}\Pr^{D(S,B)}[Z = 1] \\
&\qquad + \frac{1}{2}\Pr^{D(S,\overline{B})}[Z = 1] - \Pr^{D(S,B)}[Z = 1] \\
&= \frac{1}{2}\Pr^{D(S,\overline{B})}[Z = 1] - \frac{1}{2}\Pr^{D(S,B)}[Z = 1] \\
&= \frac{1}{2}\Delta^{D}\big((S,B),(S,\overline{B})\big).
\end{aligned}$$

Finally, putting the two equalities together, we have

$$\Delta^{D}\big((S,B),(S,U)\big) = \frac{1}{2}\Lambda^{D'}\big((S,B)\big).$$

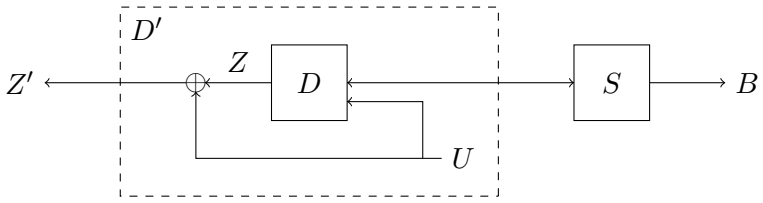

Figure 1: The mapping $D \mapsto D'$.