

Question 1 (2 Points). Give a proof that the Vernam cipher (Lecture 1, slide 23) is perfectly secret (Definition 2.3, Lecture 2, slide 16).

The proof is analogous to the one for one time pad from the lecture.

We know that the keys are random with $P[K = k] = \frac{1}{26^l}$ for every k . Where l is the length of the message.

Calculate the chance that a certain cyphertext is produced for a known plaintext.

For arbitrary k, m, c :

$$P[C = c | M = m] =$$

Substitute the encryption formula for C . mod is defined as vector modulo operator.

$$P[(M + K) \bmod 26 = c | M = m] =$$

As we already know $M = m$ we can substitute.

$$P[(m + K) \bmod 26 = c | M = m] \stackrel{K, M \text{ indep.}}{=}$$

We know that key-generation and message are independent so we can skip the condition.

$$P[(m + K) \bmod 26 = c] =$$

We refactor to isolate K

$$P[K = (c - m + 26) \bmod 26] =$$

We know the probability for any key is the same so we can solve here.

$$\frac{1}{26^l}$$

Calculate the probability of getting a certain cyphertext c

$$P[C = c] \stackrel{\text{tot. prob.}}{=} \sum_{m'} P[C = c | M = m'] * P[M = m'] =$$

Substitue $P[C = c | M = m']$

$$\frac{1}{26^l} * \sum_{m'} P[M = m'] =$$

The sum of the probability of all possible messages is 1

$$\frac{1}{26^l}$$

Use Bayes theorem to state what is needed for Def. 2.3

$$P[M = m | C = c] \stackrel{\text{Bayes}}{=} \frac{P[C=c|M=m]*P[M=m]}{P[C=c]} = \frac{\frac{1}{26^l} * P[M=m]}{\frac{1}{26^l}} = P[M = m]$$

q.e.d

1.2 Question 2

- (a) The encryption scheme is **not perfectly secret**. I came up with two different proofs, which I am going to discuss both. Without loss of generality, I assumed a uniform distribution over \mathcal{M} .

– **Arithmetic proof**

The message space contains only three possible messages and the key space contains only four possible keys. This means there are only 12 combinations of a message and a key. The modulo operation maps these combinations to only three possible ciphertext values ($\{0, 1, 2\}$), of which all three are equally likely to occur.

This means that every ciphertext fits four combinations of message and key. Since there are only three possible messages, one of the three messages is twice as likely to be the cleartext of a given ciphertext.

To illustrate this, these are the possible combinations for the ciphertext 0:

m	k
0	0
0	3
1	2
2	1

Here, it is clearly visible: Given the ciphertext 0, the probability for the message being 0 is 0.5, while the probabilities for the message being 1 or 2 are 0.25 each. This contradicts the requirement for perfect secrecy that $Pr[M = m|C = c] = Pr[M = m]$.

– **Number-theoretic proof**

Enc_k contains a modulo operation (mod 3) mapping each message to one of three residue classes, but there are four keys ($\{0, 1, 2, 3\}$), of which two (0 and 3) are in the same residue class, $\bar{0}$. For these two keys, the ciphertext is the same. In fact, it is also equal to the cleartext message. This means that for any given ciphertext, the probability for the message being equal to the ciphertext is 0.5, while the probability for the other two cases is 0.25 each, which again contradicts the requirement for perfect secrecy that $Pr[M = m|C = c] = Pr[M = m]$.

- (b) The encryption scheme is **not perfectly secret**. At a first glance, the encryption scheme seems very similar to the One-Time-Pad, except that it might use unnecessarily long keys. In fact, no information about the bits in the cleartext message can be derived from the ciphertext. However, the message space allows messages of length *up to* l . Since Enc_k returns a ciphertext of the same length as the message and discards the extra bits of the key, the length of the cleartext message can be derived from the ciphertext. This is also a violation of the principle of perfect secrecy.

To illustrate this with an example, let's set $l = 2$. This means that there are six possible messages: $\mathcal{M} = \{0, 1, 00, 01, 10, 11\}$. Without loss of generality, I will assume a uniform distribution over \mathcal{M} , so that for any given $m \in \mathcal{M}$, $Pr[M = m] = 1/6$. Now, let us assume that we received the ciphertext 01. While we cannot infer anything about the bits in the cleartext, we know that the original message is two bits long. Therefore, the probabilities for the messages change: $Pr[M = 0|C = 01] = 0 \neq 1/6 = Pr[M = 0]$.

Another way of proving that the scheme is not perfectly secret is to compare the sizes of the message space and the key space. Shannon's theorem for perfect secrecy states that for a perfectly secret encryption scheme, the number of possible keys is equal to or greater than the number of possible messages: $|\mathcal{K}| \geq |\mathbf{M}|$. For the given encryption scheme, the keys are limited to length l , while the messages can also be shorter. This means that $|\mathcal{M}| = \sum_1^l 2^l > 2^l = |\mathcal{K}|$. There are more possible messages than there are possible keys, so the encryption scheme cannot be perfectly secret. The only exception is the case in which $l = 1$, which means that $|\mathcal{M}| = |\mathcal{K}| = 2$. This is equivalent to a One-Time-Pad (and therefore perfectly secret), but it only allows single bits, which of course is of no practical use.

1.3 Question 3

The statement is wrong. For a perfectly secret encryption scheme with message space \mathcal{M} and ciphertext space \mathcal{C} , and for every $m, m' \in \mathcal{M}$ and every $c \in \mathcal{C}$, of course it holds that

$$\begin{aligned}Pr[M = m|C = c] &= Pr[M = m], \text{ and} \\Pr[M = m'|C = c] &= Pr[M = m']\end{aligned}$$

Perfect secrecy, however, does not imply that the messages themselves are equally likely to occur. This depends on the distribution on \mathcal{M} . Since the statement is made for *every* distribution on \mathcal{M} , it cannot hold true.

A simple counter-example would be a distribution over $\mathcal{M} = \{\mathbf{A}, \mathbf{B}\}$ with $Pr[M = \mathbf{A}] = 0.8$ and $Pr[M = \mathbf{B}] = 0.2$. Perfect secrecy means that

$$\begin{aligned}Pr[M = \mathbf{A}|C = c] &= Pr[M = \mathbf{A}] = 0.8 \text{ and} \\Pr[M = \mathbf{B}|C = c] &= Pr[M = \mathbf{B}] = 0.2\end{aligned}$$

Obviously, $Pr[M = \mathbf{A}|C = c] \neq Pr[M = \mathbf{B}|C = c]$.