**Cryptography, winter term 16/17:**
**Sample solution to assignment 1**

Cornelius Brand, Marc Roth

---

**Exercise 1.1 (Shift Cipher)** The following is an encryption of English text using a shift-cipher. Find the key and decrypt the ciphertext.

O QFMDHCGMGHSA GVCIZR PS GSQIFS SJSB WT SJSFMHVWBU OPCIH HVS GMGHSA SLQSDH HVS YSM WG DIPZWQ YBCKZSRUS

**Solution 1.1 (Shift Cipher)** The key is 14 and the decrypted text states Kerckhoff's Principle:
*A cryptosystem should be secure even if everything about the system, except the key, is public knowledge.*

**Exercise 1.2 (Bonus: Vigenère Cipher)** For some reasons, the TAs of the Cryptography lecture have to communicate over an unsafe channel while they create Problem 1.5. Lurking around in front of their office, you overheard them talking about it. Apparently, they decided to use the Vigenère cipher to encrypt their messages. The next day, you were able to eavesdrop a part of their communication and to obtain an encrypted chat protocol. Can you decrypt it? Which principle did they violate? Maybe there is some information hidden that can help you to solve Problem 1.5. You can find the encrypted protocol here:

`http://www-cc.cs.uni-saarland.de/crypto1617/encrypted_protocol/`

**Solution 1.2 (Bonus: Vigenère Cipher)** For any fixed key length, say $t$, one can perform a frequency analysis on $m_i, m_{i+t}, m_{i+2t}, \ldots$ for $i \in \{1, \ldots, t\}$ and check whether the highest match is the plaintext. To crack the cipher one just starts with key length 1 and iteratively performs the frequency analysis with increasing key length until the key is found. It is advisable to perform at least parts of this method with the help of a computer. In this exercise the key is GOAT and the decrypted text is:

**C:** Are you sure that using the vigenere cipher is a good idea? I mean, we could just meet at Icoffee.

**M:** Don't worry, when it was invented it resisted decryption for three centuries.

**C:** I know, but they did not have computers five hundred years ago.

**M:** Yeah, but they don't even know that we use the vigenere cipher, so we should be fine.

**C:** Well, ok... did you read my mail about the proposal for the remaining exercise?

**M:** Yes, I did. The Monty Hall Problem is a great example showing that intuition can go totally wrong when dealing with probabilities. But we should not include the name of the problem

**C:** Yeah, you are right. If we do, they can search for it online.

**M:** Then we have to come up with another setting and title...

**C:** Halloween is approaching. Lets create some bad pun!

**Exercise 1.3 (False positives)** In the sprit of interdisciplinary education, picture yourself as a student of criminal law, presented with the following case:

A murder has happened, and witnesses are nowhere to be found. Luckily, the culprit blundered and left some of his DNA at the crime scene. The police then conduct a massive screening of 10 Million people, and perform DNA tests on each of them. The test they use, on average, identifies *wrongly* one out of every million people as having the same DNA, i.e. the test reports a match with the sample DNA although they are actually distinct one millionth of the time. At the same time, if the samples actually *are* identical, the test *always* correctly reports this. You, in your future role of a judge, are now to decide the fate of a person whose DNA produces a positive test result. Would you count this as evidence against this person? Why? (In other words, calculate the probability that this person actually is the sought murderer, given that the murderer is among the sampled people.)

**Solution 1.3 (False positives)** For a randomly chosen person, let $M$ be the event that this person is the murderer, and let $T$ be the event that the test reports this person's and there murderers DNA to be indentical.

Directly from the statement of the exercise, we find

$$\begin{aligned}\Pr[M] &= 1/10^7,\\ \Pr[T|M] &= 1,\\ \Pr[T|\neg M] &= 1/10^6\end{aligned}$$

and we want to compute $\Pr[M|T]$. By Bayes' theorem, $\Pr[M|T] = \Pr[T|M] \cdot \frac{\Pr[M]}{\Pr[T]}$, and

$$\Pr[T] = \Pr[T|M] \cdot \Pr[M] + \Pr[T|\neg M] \cdot \Pr[\neg M] = 1 \cdot 1/10^7 + 1/10^6 \cdot (1 - 1/10^7)\,.$$

Plugging this in gives

$$\Pr[M|T] = \Pr[T|M] \cdot \frac{\Pr[M]}{\Pr[T]} = 1 \cdot \frac{1/10^7}{1/10^7 + 1/10^6 \cdot (1 - 1/10^7)} = \frac{1}{11 - 10^{-6}}\,.$$

**Exercise 1.4 (Coin tossing)** In this exercise, you will show how to simulate fair coins with biased coins and vice versa.

(a) You are given a biased coin, i.e. $\Pr[H] = p$ for some $0 < p < 1$, and $\Pr[T] = 1 - p$. Demonstrate an algorithm that, using only this biased coin, produces the outputs $H$ and $T$ with equal probability and prove that it is correct and halts with probability 1.

**Hint / Bonus:** Consider the experiment where the biased coin is flipped two times and let $C_1$ and $C_2$ denote the outcomes of the flips. Compute the probability $\Pr[C_1 = H | C_1 \neq C_2]$. For bonus points, calculate the expected runtime (i.e., number of iterations) of your procedure.

(b) You are given a fair coin, i.e. $\Pr[H] = \Pr[T] = 1/2$, and some $0 < t < 1$.[1] Let $(a_i)_i$ be the binary expansion of $t$, i.e. $t = \sum_{i=1}^{\infty} \frac{a_i}{2^i}$. Consider the following strategy, where initially, $n = 1$.

(i) Throw the coin. Let the result be $r$.

(ii) If $r = H$, compute the $a_n$. If $a_n = 1$, output $H$, otherwise, output $T$.

(iii) If $r \neq H$, start over and increase $n$ by one.

Prove that this strategy terminates after two iterations in expectation, and that it in fact produces $H$ with probability $t$.

**Bonus:** Take a moment and ogle in awe at the mindboggling simplicity of this procedure and the supreme power it entails: You can now imitate any (computably) biased coin with—in expectation—just two coin tosses.

**Solution 1.4 (Coin tossing)** (a) We first acknowledge the given hint and calculate

$$\Pr[C_1 = H | C_1 \neq C_2] = \frac{\Pr[C_1 = H \wedge C_1 \neq C_2]}{\Pr[C_1 \neq C_2]} = \frac{p \cdot (1-p)}{p \cdot (1-p) + p \cdot (1-p)} = \frac{1}{2}\,.$$

This means that, once we know that two independent tosses of a coin have distinct outcomes, the result of the first coin toss is distributed uniformly, i.e. like an unbiased coin.

The obvious strategy is now the following: We perform two consecutive, independent tosses of the coin, and repeat this until they are distinct (i.e., with the above terminology, until $C_1 \neq C_2$). Then we report the outcome of $C_1$ as our result. As already argued, this producedure outputs $H$ or $T$ with equal probability, if it terminates, and the probability that it terminates is 1.

To wit, the expected value of the number of coin tosses before termination is given by $\frac{1}{2p(1-p)}$.

$$\sum_{k=1}^{\infty} k \cdot (p^2 + (1-p)^2)^{k-1} \cdot 2p(1-p) = 2p(1-p) \cdot \sum_{k=1}^{\infty} k \cdot (p^2 + (1-p)^2)^{k-1} = \frac{1}{2p(1-p)}$$

[1] For those acquainted with the term: $t$ needs to be computable.

(b) The event of $H$ coming up has probability $\frac{1}{2}$, so the expected number of coin tosses until this happens is 2 (as can be seen analogous to the expectation in (a)). As for the probability, per definition of the procedure, the probability that $H$ is output by the procedure

$$\begin{aligned}
&\Pr[H] = \\
&\Pr\left[\bigvee_{i=1}^{\infty}\left(\bigwedge_{j=1}^{i-1}(r_j \neq H) \wedge r_i = H \wedge a_i = 1\right)\right] = \\
&\sum_{i=1}^{\infty}\Pr\left[\bigwedge_{j=1}^{i-1}(r_j \neq H) \wedge r_i = H \wedge a_i = 1\right] = \\
&\sum_{i=1}^{\infty}\prod_{j=1}^{i-1}\Pr[r_j \neq H] \cdot \Pr[r_i = H] \cdot \Pr[a_i = 1)] \ .
\end{aligned}$$

Since the $a_i$ are constant, by definition $\Pr[a_i = 1] = a_i$ for all $i$, and $\Pr[r_j = H] = \Pr[r_j \neq H] = \frac{1}{2}$. Therefore,

$$\Pr[H] = \sum_{i=1}^{\infty}\frac{a_i}{2^i}$$

and by the definition of binary expansion, this is nothing else than

$$\Pr[H] = t$$

**Exercise 1.5 (The Haunty Mall Problem)** It is the evening of the 31st of October, and after a long day's studies, you decide to make a recreational trip to the local mall. But — who could have seen it coming! — your visit turns into a living nightmare when you are captured by dreadful creatures (demons, ghosts and suchlike) and dragged down into the eternal fires of hell. You then find yourself trapped in a gruesome game: The soulless dwellers of hell thrive on seeing students suffer, and hence they confront you with a riddle.

There are 666 closed gates in front of you. Of these gates, 665 harbour only everlasting, hellish pain and sorrow, but one of them, so you are promised, will lead you back to your earthly home. You have to decide for one of these 666 gates, and after you do so, one of the demons will open, 664 of the remaining 665 portals you did not choose, such that behind *all* of them there is, you know, everlasting, hellish pain and sorrow. You can now decide: Will you stick with your original choice, or will you change your mind and pick the only remaining closed gate?

Prove to the demons that your years at university have made you immune against soulless creatures thriving on seeing students suffer, by showing how to maximize your chance of returning home. Good luck!

**Solution 1.5 (The Haunty Mall Problem)** It is always better to change. The changing candidate loses if and only if his original first choice was already the good gate.

This happens with probability $1/n$, where $n = 666$ is the number of gates. Conversely, this means that the candidate wins with probability $(n-1)/n$ if he changes.

Despite this proof's simplicity, make sure you *really* understand what's happening here. This problem, in the case of only three gates, has fooled some of the greatest minds of the past century, including Paul Erdős.