Homework 1

Please submit your homeworks through CS Dropbox:
https://dropbox.cs.princeton.edu/COS433_S2018/HW1

1 Problem 1 (20 points)

(a) Show that a simple substitution cipher does *not* have perfectly secrecy, even if the permutation is chosen uniformly at random. That is, provide two messages of the same length such that the distributions of encryptions of those two messages are different. Your messages should consist of grammatically correct English sentences.

(b) Show that a Homophonic substitution cipher that encrypts a single character at a time can *never* have perfect secrecy, no matter how large the output alphabet is. (again, using grammatically correct English messages)

(c) Prove that the Vigenère cipher has perfect secrecy if the keyword is a sequence of random characters, and the message length is no longer than the key length

(d) Show that the Vigenère cipher does *not* have perfect secrecy if the message length is even one character longer than the key length. (again, using grammatically correct English messages)

(e) Show that the Vigenère cipher does *not* have perfect secrecy if the keyword is a random sequence of English words, even if the message is shorter than the key. (again, using grammatically correct English messages)

(f) Show that a transposition cipher cannot have perfect secrecy, even if the permutation is chosen uniformly at random.

(g) Show that the following cipher has perfect secrecy. Messages are ℓ bit strings. The key is a random permutation P on 2ℓ items. To encrypt a message m, write down m, followed by $\overline{m}$, the bitwise complement of m. Then permute the bits of the resulting 2ℓ-bit string $m||\overline{m}$ according to the permutation described by the key.

2 Problem 2 (10 points)

(a) Devise an encryption scheme such that (1) given an encryption of any message, an adversary can figure out 90% of the secret key, but (2) the scheme is still perfectly secure, despite 90% of the key being revealed. Do not forget to prove that the scheme is secure and that it is correct.

(b) Devise an encryption scheme such that (1) given an encryption of any message, an adversary learns *nothing* about the secret key, but (2) the scheme is completely broken (as in, given the ciphertext, an adversary can completely recover the plaintext).

3 Problem 3 (30 points)

Suppose we say that two messages are *adjacent* if they differ by at most a single bit.

Definition 1. *An encryption scheme* $(\mathsf{Enc}, \mathsf{Dec})$ *for* ℓ*-bit messages has* adjacent-message perfectly secrecy *if, for any two* ℓ*-bit* adjacent *messages* m_0, m_1*, the distributions* $\mathsf{Enc}(k, m_0)$ *and* $\mathsf{Enc}(k, m_1)$ *are identical.*

This is the same definition as perfect secrecy seen in class, except for the restriction that it only applies to m_0, m_1 that are adjacent. Therefore, it is a seemingly weaker definition.

(a) Prove that any encryption scheme that has *adjacent-message perfect* secrecy must in fact have *perfect* secrecy.
Hint: suppose m_0, m_1 *differed in just 2 bits. How would you prove that the distributions of their encryptions are identical? Generalize this to arbitrary messages.*

(b) Suppose instead we used a different notion of adjacent messages. Interpret each message as an integer from 0 to $2^\ell - 1$, and say that two messages m_0, m_1 are adjacent if $m_0 = m_1 \pm 1$. Suppose we changed Definition 1 to apply only to adjacent messages in this sense. Does an encryption scheme satisfying this notion of adjacent-message perfect secrecy necessarily satisfy normal perfect secrecy? If yes, prove it. If not, provide an example of a scheme with adjacent-message perfect secrecy that is does not have perfect secrecy.

(c) Suppose instead we used yet a different notion of adjacent messages. Here, two messages are adjacent if m_0 can be obtained by cyclically rotating the ℓ bits of m_1. So `"10110010"` is adjacent to `"11001010"` (move the first two bits of the first message to the end to obtain the second message), but `"10110010"` is

not adjacent to `"10110001"`. Suppose we changed Definition 1 to apply only to adjacent messages in this sense. Does an encryption scheme satisfying this notion of adjacent-message perfect secrecy necessarily satisfy normal perfect secrecy? If yes, prove it. If not, provide an example of a scheme with adjacent-message perfect secrecy that is does not have perfect secrecy.

(d) Provide a simple, tight characterization of the kinds of notions of "adjacent" messages for which adjacent-message perfect secrecy implies perfect secrecy. Prove that, for any notion of adjacency meeting your criteria and any encryption scheme satisfying adjacent-message perfect secrecy for that notion of adjacency, that the scheme must actually have perfect secrecy. For any notion of adjacency that does *not* meet your criteria, show how to build an encryption scheme that has adjacent message perfect secrecy, but not perfect secrecy. Your scheme does not need to be efficient.