Introduction to Modern Cryptography 2nd edition Katz Solutions Manual

SKU: 9781466570269-SOLUTIONS

Chapter 2

Perfectly Secret Encryption – Solutions

2.1 Prove that, by redefining the key space, we may assume that the key-generation algorithm Gen chooses a key uniformly at random from the key space, without changing $\Pr[C = c \mid M = m]$ for any m, c.

Solution: If Gen is a randomized algorithm, we may view it as a deterministic algorithm that takes as input a random tape ω of some length; the distribution on the output of Gen is, by definition, the distribution obtained by choosing uniform ω and then running $\mathsf{Gen}(\omega)$. So, rather than letting the key be the output of Gen, we can simply let the key be ω itself (and redefine the key space accordingly).

Formally, given a scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ in which Gen is randomized, construct a new scheme $(\mathsf{Enc}', \mathsf{Dec}')$ where the key is a uniform ω. Then define $\mathsf{Enc}'_\omega(m)$ to compute $k := \mathsf{Gen}(\omega)$ followed by $\mathsf{Enc}_k(m)$, and define decryption analogously.

2.2 Prove that, by redefining the key space, we may assume that Enc is deterministic without changing $\Pr[C = c \mid M = m]$ for any m, c.

Solution: As in the previous exercise, if Enc is a randomized algorithm then we may view it as being a deterministic algorithm that also takes a random tape ω as additional input. The distribution on the output of $\mathsf{Enc}_k(m)$ is then, by definition, the distribution obtained by choosing uniform ω and then computing $\mathsf{Enc}_k(m;\omega)$. We then define key generation to include ω as well as k (and redefine the key space accordingly).

Formally, given a scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ in which Enc is randomized, construct a new scheme $(\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}' = \mathsf{Dec})$ as follows. Gen' computes $k \leftarrow \mathsf{Gen}$ and also chooses uniform ω; the key is (k, ω). Then define $\mathsf{Enc}'_{(k,\omega)}(m)$ to be $\mathsf{Enc}_k(m;\omega)$.

2.3 Prove or refute: An encryption scheme with message space $\mathcal{M}$ is perfectly secret if and only if for every probability distribution over $\mathcal{M}$ and every $c_0, c_1 \in \mathcal{C}$ we have $\Pr[C = c_0] = \Pr[C = c_1]$.

Solution: This is not true. Consider modifying the one-time pad so encryption appends a bit that is 0 with probability 1/4 and 1 with probability 3/4. This scheme will still be perfectly secret, but ciphertexts ending in 1 are more likely than ciphertexts ending in 0.

K16478_SM_Cover.indd 13 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

2.4 Prove the second direction of Lemma 2.4.

Solution: Say $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is perfectly secret. Fix two messages m, m' and a ciphertext c that occurs with nozero probability, and consider the uniform distribution over $\{m, m'\}$. Perfect secrecy implies that $\Pr[M = m \mid C = c] = 1/2 = \Pr[M = m' \mid C = c]$. But

$$\begin{aligned}\frac{1}{2} = \Pr[M = m \mid C = c] &= \frac{\Pr[C = c \mid M = m] \cdot \Pr[M = m]}{\Pr[C = c]} \\ &= \frac{\frac{1}{2}\Pr[C = c \mid M = m]}{\Pr[C = c]},\end{aligned}$$

and so $\Pr[C = c \mid M = m] = \Pr[\mathsf{Enc}_K(m) = c] = \Pr[C = c]$. Since an analogous calculation holds for m' as well, we conclude that $\Pr[\mathsf{Enc}_K(m) = c] = \Pr[\mathsf{Enc}_K(m') = c]$.

2.5 Prove Lemma 2.6.

Solution: We begin by proving that any encryption scheme that is perfectly secret is perfectly indistinguishable. Every adversary $\mathcal{A}$ participating in $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}$ defines a fixed pair of plaintext messages m_0, m_1 that it outputs in the first step of the experiment. (Note that since $\mathcal{A}$ is a deterministic algorithm, it always outputs the same pair of messages.) Fix $\mathcal{A}$ and fix m_0, m_1 output by $\mathcal{A}$. By Lemma 2.4, for every $m_0, m_1 \in \mathcal{M}'$ and every $c \in \mathcal{C}$,

$$\Pr[\mathsf{Enc}_K(m_0) = c] = \Pr[\mathsf{Enc}_K(m_1) = c]. \tag{2.1}$$

In particular, the above holds for m_0, m_1 output by $\mathcal{A}$ and for any c. Let $\mathcal{C}_0$ (resp., $\mathcal{C}_1$) denote the set of ciphertexts for which $\mathcal{A}$ outputs 0 (resp., 1) at the conclusion of experiment $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}$; since $\mathcal{A}$ is deterministic these sets are well-defined. Note that since $\mathcal{A}$ must output either 0 or 1, it follows that $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{C}_1$. We have:

$$\begin{aligned}\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1] &= \frac{1}{2} \cdot \Pr[\mathcal{A} \text{ outputs } 0 \mid b = 0] + \frac{1}{2} \cdot \Pr[\mathcal{A} \text{ outputs } 1 \mid b = 1] \\ &= \frac{1}{2} \cdot \sum_{c \in \mathcal{C}_0} \Pr[c = \mathsf{Enc}_K(m_0)] + \frac{1}{2} \cdot \sum_{c \in \mathcal{C}_1} \Pr[c = \mathsf{Enc}_K(m_1)] \\ &= \frac{1}{2} \cdot \sum_{c \in \mathcal{C}_0} \Pr[c = \mathsf{Enc}_K(m_0)] + \frac{1}{2} \cdot \sum_{c \in \mathcal{C}_1} \Pr[c = \mathsf{Enc}_K(m_0)] \\ &= \frac{1}{2} \cdot \sum_{c \in \mathcal{C}} \Pr[c = \mathsf{Enc}_K(m_0)] \quad = \quad \frac{1}{2} \cdot (1) = \frac{1}{2}.\end{aligned}$$

where the second last equality is due to Equation (2.1).

We now proceed to the other direction, that any perfectly indistinguishable encryption scheme is also perfectly secret. We prove the contrapositive. Assume encryption scheme Π is not perfectly secret with respect

K16478_SM_Cover.indd 14 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

to Definition 2.1. Then by Lemma 2.4, there must exist two messages $m_0, m_1 \in \mathcal{M}$ and a ciphertext $\tilde{c} \in \mathcal{C}$ for which

$$\Pr[\mathsf{Enc}_K(m_0) = \tilde{c}] \neq \Pr[\mathsf{Enc}_K(m_1) = \tilde{c}]. \tag{2.2}$$

Let $\mathcal{A}$ be an adversary who outputs m_0, m_1 in the first step of $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}$. Then if $\mathcal{A}$ receives any ciphertext $c' \neq \tilde{c}$ it outputs a random bit b'; if it receives the ciphertext $\tilde{c}$, it outputs $b' = 0$.

Since each message is chosen with probability 1/2 in the experiment, we have

$$\begin{aligned}&\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1]\\ &= \frac{1}{2} \cdot \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \mid M = m_0] + \frac{1}{2} \cdot \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \mid M = m_1].\end{aligned}$$

Next,

$$\begin{aligned}&\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \mid M = m_0]\\ &= \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \ \wedge \ \tilde{c} = \mathsf{Enc}_K(m_0)]\\ &\quad + \ \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \ \wedge \ \tilde{c} \neq \mathsf{Enc}_K(m_0)]\\ &= \Pr[\tilde{c} = \mathsf{Enc}_K(m_0)] + \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \ \wedge \ \tilde{c} \neq \mathsf{Enc}_K(m_0)]\\ &= \Pr[\tilde{c} = \mathsf{Enc}_K(m_0)] + \frac{1}{2} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_0)]\end{aligned}$$

where the second-to-last equality is because $\mathcal{A}$ always outputs 0 when receiving $\tilde{c}$ (and when $M = m_0$ this means that $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1$), and the last equality is because $\mathcal{A}$ outputs a random bit when given a ciphertext $c' \neq \tilde{c}$. A similar analysis for the case that $M = m_1$ (the only difference is that now $\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \ \wedge \ \tilde{c} \neq \mathsf{Enc}_K(m_1)] = 0$) gives

$$\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1 \mid M = m_1] = \frac{1}{2} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_1)].$$

Putting this all together we have:

$$\begin{aligned}\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1] &= \frac{1}{2} \cdot \Big(\Pr[\tilde{c} = \mathsf{Enc}_K(m_0)] + \frac{1}{2} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_0)] \Big)\\ &\quad + \ \frac{1}{2} \cdot \frac{1}{2} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_1)]\\ &= \frac{1}{2} \cdot \Big(\Pr[\tilde{c} = \mathsf{Enc}_K(m_0)] + \frac{1}{2} \cdot (1 - \Pr[\tilde{c} = \mathsf{Enc}_K(m_0)]) \Big)\\ &\quad + \ \frac{1}{2} \cdot \frac{1}{2} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_1)]\\ &= \frac{1}{4} + \frac{1}{4} \cdot \Pr[\tilde{c} = \mathsf{Enc}_K(m_0)] + \frac{1}{4} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_1)]\\ &\neq \frac{1}{4} + \frac{1}{4} \cdot \Pr[\tilde{c} = \mathsf{Enc}_K(m_1)] + \frac{1}{4} \cdot \Pr[\tilde{c} \neq \mathsf{Enc}_K(m_1)]\\ &= \frac{1}{4} + \frac{1}{4} = \frac{1}{2}\end{aligned}$$

K16478_SM_Cover.indd 15 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

where the inequality is by Equation (2.2). Since $\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1] \neq 1/2$ we have that Π is not perfectly indistinguishable.

2.6 For each of the following encryption schemes, state whether the scheme is perfectly secret. Justify your answer in each case.

(a) The message space is $\mathcal{M} = \{0, \ldots, 4\}$. Algorithm Gen chooses a uniform key from the key space $\{0, \ldots, 5\}$. $\mathsf{Enc}_k(m)$ returns $[k + m \bmod 5]$, and $\mathsf{Dec}_k(c)$ returns $[c - k \bmod 5]$.

(b) The message space is $\mathcal{M} = \{m \in \{0,1\}^\ell \mid \text{the last bit of } m \text{ is } 0\}$. Gen chooses a uniform key from $\{0,1\}^{\ell-1}$. $\mathsf{Enc}_k(m)$ returns ciphertext $m \oplus (k \| 0)$, and $\mathsf{Dec}_k(c)$ returns $c \oplus (k \| 0)$.

Solution:

(a) The scheme is not perfectly secret. To see this, we can use the equivalent definition of perfect secrecy given by Equation (2.1). If the message is 0, then the ciphertext is 0 if and only if $k \in \{0, 5\}$. So $\Pr[\mathsf{Enc}_K(0) = 0] = 2/5$. On the other hand, if the message is 1, then the ciphertext is 0 if and only if $k = 4$. So

$$\Pr[\mathsf{Enc}_K(1) = 0] = 1/5 \neq \Pr[\mathsf{Enc}_K(0) = 0].$$

(b) One can prove that this is perfectly secret by analogy with the one-time pad. (Essentially the final bit of the message is being ignored here, since it is always 0.)

2.7 When using the one-time pad with the key $k = 0^\ell$, we have $\mathsf{Enc}_k(m) = k \oplus m = m$ and the message is sent in the clear! It has therefore been suggested to modify the one-time pad by only encrypting with $k \neq 0^\ell$ (i.e., to have Gen choose k uniformly from the set of *nonzero* keys of length ℓ). Is this modified scheme still perfectly secret? Explain.

Solution: The modified scheme is *not* perfectly secret. To see this formally, consider the uniform distribution over $\mathcal{M} = \{0,1\}^\ell$. For any fixed message $\alpha \in \{0,1\}^\ell$, we have

$$\Pr[M = \alpha \mid C = \alpha] = 0 \neq \Pr[M = \alpha].$$

This contradicts perfect secrecy.

We conclude that in order to obtain perfect secrecy, it must be possible to encrypt using the key 0^ℓ. This may seem counter-intuitive, since this key does not change the plaintext. However, note that an eavesdropper has no way of knowing if the key is 0^ℓ, so the fact that the ciphertext is the same as the plaintext in this case is really of no help to the adversary.

K16478_SM_Cover.indd 16 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

2.8 Let Π denote the Vigenère cipher where the message space consists of all 3-character strings (over the English alphabet), and the key is generated by first choosing the period t uniformly from $\{1, 2, 3\}$ and then letting the key be a uniform string of length t.

(a) Define $\mathcal{A}$ as follows: $\mathcal{A}$ outputs $m_0 = \texttt{aab}$ and $m_1 = \texttt{abb}$. When given a ciphertext c, it outputs 0 if the first character of c is the same as the second character of c, and outputs 1 otherwise. Compute $\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1]$.

(b) Construct and analyze an adversary $\mathcal{A}'$ for which $\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A}',\Pi} = 1]$ is greater than your answer from part (a).

Solution:

(a) Say aab is encrypted to give ciphertext c. What is the probability that the first and second characters of c are equal? When $t = 1$ (which occurs 1/3 of the time) this always happens. But when $t \in \{2, 3\}$ this happens only if the first and second characters of the key are equal, which occurs with probability 1/26. So

$$\Pr[\mathcal{A} \text{ outputs } 0 \mid m_0 \text{ is encrypted}] = \frac{1}{3} + \frac{2}{3} \cdot \frac{1}{26} \approx 0.359.$$

If instead abb is encrypted, then the first and second characters of c can never be equal when $t = 1$, but are equal with probability 1/26 when $t \in \{2, 3\}$. Thus,

$$\Pr[\mathcal{A} \text{ outputs } 0 \mid m_1 \text{ is encrypted}] = \frac{2}{3} \cdot \frac{1}{26} \approx 0.026.$$

We therefore have

$$\begin{aligned}
&\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1] \\
&= \frac{1}{2} \cdot \Pr[\mathcal{A} \text{ outputs } 0 \mid m_0 \text{ is encrypted}] \\
&\quad + \frac{1}{2} \cdot \Pr[\mathcal{A} \text{ outputs } 1 \mid m_1 \text{ is encrypted}] \\
&\approx \frac{1}{2} \cdot 0.359 + \frac{1}{2} \cdot 0.974 \approx 0.667.
\end{aligned}$$

(b) There are many possible solutions; we present one. Consider the adversary $\mathcal{A}'$ who outputs $m_0 = \texttt{aaa}$ and $m_1 = \texttt{abc}$ and outputs '0' iff the first and second characters in the ciphertext c are the same, or if the first and last characters in the ciphertext are the same. Call this event E.

Say aaa is encrypted. If $t \in \{1, 2\}$ then E always happens. When $t = 3$ all characters in the ciphertext are uniform and independent;

K16478_SM_Cover.indd 17 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

rather than calculate the probability of E in this case, we just let p denote that probability.

Say `abc` is encrypted. If $t = 1$ then E never happens. When $t = 2$ the first and last characters of c are never equal, but the first and second characters are equal with probability $1/26$. When $t = 3$ then all characters in the ciphertext are random and so E occurs with probability p.

Putting everything together gives:

$$\begin{aligned}
&\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A}',\Pi} = 1] \\
&= \frac{1}{2} \cdot \Pr[\mathcal{A}' \text{ outputs } 0 \mid m_0 \text{ is encrypted}] \\
&\quad + \frac{1}{2} \cdot \Pr[\mathcal{A}' \text{ outputs } 1 \mid m_1 \text{ is encrypted}] \\
&= \frac{1}{2} \cdot \left(\frac{2}{3} + \frac{1}{3} \cdot p\right) + \frac{1}{2} \cdot \left(1 - \left(\frac{1}{3} \cdot \frac{1}{26} + \frac{1}{3} \cdot p\right)\right) \approx 0.827.
\end{aligned}$$

2.9 In this exercise, we look at different conditions under which the shift, mono-alphabetic substitution, and Vigenère ciphers are perfectly secret:

(a) Prove that if only a single character is encrypted, then the shift cipher is perfectly secret.

(b) What is the largest message space $\mathcal{M}$ for which the mono-alphabetic substitution cipher provides perfect secrecy?

(c) Prove that the Vigenère cipher using (fixed) period t is perfectly secret when used to encrypt messages of length t.

Reconcile this with the attacks shown in the previous chapter.

Solution:

(a) This can be proved directly (as in the case of the one-time pad) or using Shannon's theorem.

(b) Let $\mathcal{M}$ be the set of all permutations of the alphabet (i.e., all strings of length 26 with no repeated letter). One can prove directly that the mono-alphabetic substitution cipher is perfectly secret for this message space; we prove it using Shannon's theorem. Briefly, we have $|\mathcal{M}| = |\mathcal{K}| = |\mathcal{C}| = 26!$, where the size of $\mathcal{C}$ follows from the fact that the composition of two permutations is a permutation (and furthermore all permutations can be obtained in this way). This latter fact also proves the second requirement of Shannon's theorem: for every two permutations (namely, plaintext and ciphertext) there exists only one permutation π that maps the plaintext to the ciphertext. This is the largest possible plaintext space, because by Theorem 2.7 we must have $|\mathcal{M}| \leq |\mathcal{K}|$ for *any* perfectly secret encryption scheme.

K16478_SM_Cover.indd 18 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

(c) This can be proved directly (as in the case of the one-time pad) or using Shannon's theorem.

The attacks in Chapter 1 rely on longer plaintexts being encrypted.

2.10 Prove that a scheme satisfying Definition 2.5 must have $|\mathcal{K}| \geq |\mathcal{M}|$ without using Lemma 2.6. Specifically, let Π be an arbitrary encryption scheme with $|\mathcal{K}| < |\mathcal{M}|$. Show an $\mathcal{A}$ for which $\Pr\left[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1\right] > \frac{1}{2}$.

Solution: Consider the following $\mathcal{A}$: output uniform $m_0, m_1 \in \mathcal{M}$. Upon receiving ciphertext c, check (by exhaustive search) whether there exists a key k such that $\mathsf{Dec}_k(c) = m_0$. If so, output 0; else output 1.

When m_0 is encrypted then $\mathcal{A}$ always outputs 0. On the other hand, when m_1 is encrypted then there are at most $|\mathcal{K}|$ possible messages that c can decrypt to; since m_0 is uniform and independent of m_1 the probability that m_0 is equal to one of those messages, and so the probability that $\mathcal{A}$ outputs 0 in this case, is at most $|\mathcal{K}|/|\mathcal{M}| < 1$. We conclude that

$$\begin{aligned}\Pr\left[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1\right] &= \frac{1}{2} \cdot \Pr[\mathcal{A}(\mathsf{Enc}_K(m_0)) = 0] + \frac{1}{2} \cdot \Pr[\mathcal{A}(\mathsf{Enc}_K(m_1)) = 1] \\ &= \frac{1}{2} + \frac{1}{2} \cdot (1 - \Pr[\mathcal{A}(\mathsf{Enc}_K(m_1)) = 0]) \;>\; \frac{1}{2}\,.\end{aligned}$$

2.11 Assume we require only that an encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\mathcal{M}$ satisfy the following: For all $m \in \mathcal{M}$, we have $\Pr[\mathsf{Dec}_K(\mathsf{Enc}_K(m)) = m] \geq 2^{-t}$. (This probability is taken over choice of the key as well as any randomness used during encryption.) Show that perfect secrecy can be achieved with $|\mathcal{K}| < |\mathcal{M}|$ when $t \geq 1$. Prove a lower bound on the size of $\mathcal{K}$ in terms of t.

Solution: Let $\mathcal{K} = \{0,1\}^\ell$ and $\mathcal{M} = \{0,1\}^{\ell+t}$. The key-generation algorithm chooses a uniform string from $\mathcal{K}$. To encrypt a message $m \in \mathcal{M}$ using key k, let m' denote the first ℓ bits of m and output $c := m' \oplus k$ (both m' and k have length ℓ). To decrypt a ciphertext c using key k, choose a random string $r \leftarrow \{0,1\}^t$ and output $m := (c \oplus k) \| r$. Note that $\Pr[\mathsf{Dec}_k(\mathsf{Enc}_k(m)) = m] = 2^{-t}$ because decryption is correct if and only if the random string r chosen during decryption happens to equal the last t bits of m (and this occurs with probability 2^{-t}). Perfect secrecy of this scheme follows from the proof of the one-time pad (indeed, this is exactly a one-time pad on the first ℓ bits of the message).

We now prove the following lower bound:

THEOREM *Let* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a perfectly secret encryption scheme over message space* $\mathcal{M}$ *with* $\Pr[\mathsf{Dec}_k(\mathsf{Enc}_k(m)) = m] \geq 2^{-t}$. *Then* $|\mathcal{K}| \geq |\mathcal{M}| \cdot 2^{-t}$.

K16478_SM_Cover.indd 19 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

PROOF Let $\Pr^*$ denote the maximum probability with which an unbounded algorithm can cause the stated event to occur. Consider an experiment in which a message m is chosen uniformly from $\mathcal{M}$ and we are interested in the maximum probability of correctly guessing m. Clearly, $\Pr^*[\text{guess } m] = |\mathcal{M}|^{-1}$. Now consider an extension of this experiment where a key k is generated using Gen and a ciphertext $c \leftarrow \mathsf{Enc}_k(m)$ is computed. Perfect secrecy implies that

$$|\mathcal{M}|^{-1} = \Pr^*[\text{guess } m] = \Pr^*[\text{guess } m \text{ given } c].$$

Continuing, we have

$$\begin{aligned}\Pr^*[\text{guess } m \text{ given } c] &\geq \Pr^*[\text{guess } m \text{ and } k \text{ given } c] \\ &= \Pr^*[\text{guess } k \text{ given } c] \cdot \Pr^*[\text{guess } m \text{ given } k \text{ and } c] \\ &\geq |\mathcal{K}|^{-1} \cdot \Pr^*[\text{guess } m \text{ given } k \text{ and } c] \\ &\geq |\mathcal{K}|^{-1} \cdot 2^{-t},\end{aligned}$$

by the correctness guarantee, and using the fact that m and k are independent. We conclude that $|\mathcal{K}| \geq |\mathcal{M}| \cdot 2^{-t}$. ■

2.12 Let $\varepsilon \geq 0$ be a constant. Say an encryption scheme is *ε-perfectly secret* if for every adversary $\mathcal{A}$ it holds that

$$\Pr\left[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1\right] \leq \frac{1}{2} + \varepsilon\,.$$

(Compare to Definition 2.5.) Show that ε-perfect secrecy can be achieved with $|\mathcal{K}| < |\mathcal{M}|$ when $\varepsilon > 0$. Prove a lower bound on the size of $\mathcal{K}$ in terms of ε.

Solution: (See `http://eprint.iacr.org/2012/053` for further details.) We first show that ε-perfect secrecy can be achieved in this case. Let the message space be $\mathcal{M} = \{0,1\}^\ell$ for some ℓ, and let $\mathcal{K} \subset \{0,1\}^\ell$ be an arbitrary set of size $(1-\varepsilon)\cdot 2^\ell$. One can check that for any $\mathcal{A}$ we have $\Pr\left[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi} = 1\right] \leq \frac{1}{2} + \varepsilon$. One can show that this scheme is optimal, in the sense that any scheme that is ε-perfectly secret must have $|\mathcal{K}| \geq (1-\varepsilon)\cdot|\mathcal{M}|$.

2.13 In this problem we consider definitions of perfect secrecy for the encryption of *two* messages (using the same key). Here we consider distributions over *pairs* of messages from the message space $\mathcal{M}$; we let M_1, M_2 be random variables denoting the first and second message, respectively. (We stress that these random variables are not assumed to be independent.) We generate a (single) key k, sample a pair of messages (m_1, m_2) according to the given distribution, and then compute ciphertexts $c_1 \leftarrow \mathsf{Enc}_k(m_1)$ and $c_2 \leftarrow \mathsf{Enc}_k(m_2)$; this induces a distribution over pairs of ciphertexts and we let C_1, C_2 be the corresponding random variables.

K16478_SM_Cover.indd 20 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm

(a) Say encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is *perfectly secret for two messages* if for all distributions over $\mathcal{M} \times \mathcal{M}$, all $m_1, m_2 \in \mathcal{M}$, and all ciphertexts $c_1, c_2 \in \mathcal{C}$ with $\Pr[C_1 = c_1 \wedge C_2 = c_2] > 0$:

$$\Pr[M_1 = m_1 \wedge M_2 = m_2 \mid C_1 = c_1 \wedge C_2 = c_2] = \Pr[M_1 = m_1 \wedge M_2 = m_2].$$

Prove that *no* encryption scheme can satisfy this definition.

(b) Say encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is *perfectly secret for two distinct messages* if for all distributions over $\mathcal{M} \times \mathcal{M}$ where the first and second messages are guaranteed to be different (i.e., distributions over pairs of *distinct* messages), all $m_1, m_2 \in \mathcal{M}$, and all $c_1, c_2 \in \mathcal{C}$ with $\Pr[C_1 = c_1 \wedge C_2 = c_2] > 0$:

$$\Pr[M_1 = m_1 \wedge M_2 = m_2 \mid C_1 = c_1 \wedge C_2 = c_2] = \Pr[M_1 = m_1 \wedge M_2 = m_2].$$

Show an encryption scheme that provably satisfies this definition.

Solution:

(a) The definition requires the equation to hold for all pairs of plaintexts m, m' and ciphertexts c, c', *even when* $c = c'$ *but* $m \neq m'$. We show that this is impossible. Take the uniform distribution over $\mathcal{M}$ and any c such that $\Pr[C = c \wedge C' = c] > 0$. Let $m, m' \in \mathcal{M}$ be distinct. For any scheme with no decryption error, and any key k, we must have

$$\Pr[C = c \wedge C' = c \mid M = m \wedge M' = m' \wedge K = k] = 0.$$

(If not, then decryption of the ciphertext c using the key k gives an error some of the time.) The above implies

$$\Pr[C = c \wedge C' = c \mid M = m \wedge M' = m'] = 0.$$

So $\Pr[M = m \wedge M' = m' \mid C = c \wedge C' = c] = 0$ but on the other hand $\Pr[M = m \wedge M' = m'] > 0$. This holds for any encryption scheme, and so no scheme can satisfy the given definition.

(b) Let $\mathcal{K}$ be the set of all permutations on $\mathcal{M}$, and let Gen choose a random permutation from $\mathcal{K}$. (Note that Gen is not necessarily efficient; nevertheless, this is allowed in this exercise.) Encryption is carried out by applying the permutation specified by the key to the plaintext, and decryption is carried out by applying the inverse of the permutation to the ciphertext. This scheme satisfies the definition; details omitted.

Note: An efficient scheme meeting this definition can be constructed using *pairwise-independent permutations* (a topic not covered in the book).

K16478_SM_Cover.indd 21 Introduction to Modern Cryptography 2nd edition Katz Solutions Manual 05/04/16 3:43 pm