Block ciphers
And modes of operation

Foundations of Cryptography
Computer Science Department
Wellesley College

Fall 2016

Table of contents

Keyed permutations: Some definitions

Definition. Let $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ be an efficient, length-preserving, keyed function. We call F a *keyed permutation* if for every k, the function $F_k(\cdot)$ is *one-to-one*.

Definition. We say that a keyed permutation is *efficient* if there is a polynomial time algorithm computing $F_k(x)$ given k and x, as well as a polynomial-time algorithms computing $F_k^{-1}(x)$ given k and x.

Remark. The input and output lengths, called the *block size* are the same, but the key length may be smaller or larger than the block size.

Pseudorandom permutations

Definition. Let $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ an efficient keyed permutation. We say that F is a *pseudorandom permutation* if for all probabilistic polynomial-time distinguishers D, there exists a negligible function negl such that:

$$\left| \Pr[D^{F_k(\cdot)}(1^n) = 1] - \Pr[D^{f(\cdot)}(1^n) = 1] \right| \leq \mathsf{negl}(n),$$

where $k \leftarrow \{0,1\}^n$ is chosen uniformly at random and f is chosen uniformly at random from the set of permutations mapping n-bit strings to n-bit strings.

Pseudorandom functions and permutations are polynomially indistinguishable

Theorem 3.27. If F is a pseudorandom permutation then it is also a pseudorandom function.

Proof. The basic idea behind the proof is that a random function f looks identical to a random permutation unless a distinct pair of values x and y are found for which $f(x) = f(y)$. The probability of finding such points x, y using a polynomial number of queries is negligible. $\square$

A new security concern

- If F is an efficient pseudorandom permutation then cryptographic schemes based on F might require honest parties to compute both F_k and F_k^{-1}.
- We may wish that F_k is indistinguishable from a random permutation even if the distinguisher is given oracle access to the inverse permutation.

"Instead of doing jail time, I'm performing eighteen hundred hours of community service."

Strong pseudorandom permutations

Definition 3.28. Let $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ an efficient keyed permutation. We say that F is a *strong pseudorandom permutation* if for all probabilistic polynomial-time distinguishers D, there exists a negligible function negl such that:

$$\left| \Pr[D^{F_k(\cdot), F_k^{-1}(\cdot)}(1^n) = 1] - \Pr[D^{f(\cdot), f^{-1}(\cdot)}(1^n) = 1] \right| \leq \mathrm{negl}(n),$$

where $k \leftarrow \{0,1\}^n$ is chosen uniformly at random and f is chosen uniformly at random from the set of permutations mapping n-bit strings to n-bit strings.

Block ciphers

- From your reading you know that stream ciphers can be modeled as pseudorandom generators.
- The analogue for the case of a strong pseudorandom permutation is a *block cipher*.
- Block ciphers are not secure encryption schemes. Rather, they are building blocks that can be used to construct secure schemes.

*Modes of operations**

- A mode of operation is essentially a way of encrypting arbitrary-length messages using a block cipher (i.e., pseudorandom permutation).
- Note that messages can be unambiguously padded to a total length that is a multiple of the block size by appending a 1 followed by sufficiently-may 0's. Our notation: ℓ blocks each of size n.

*Donald has volunteered to help demonstrate.

Electronic code book (ECB) mode

Electronic code book. Given a plaintext message $m = m_1, \ldots, m_\ell$, the ciphertext is obtained by "encrypting" each block separately, i.e., $c = \langle F_k(m_1), \ldots, F_k(m_\ell) \rangle$.

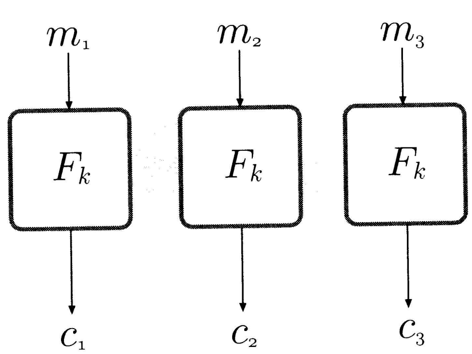

Donald does ECB

- ECB is deterministic and therefore cannot be CPA-secure.
- Worse, ECB-mode encryption does not even have indistinguishable encryptions in the presence of eavesdroppers.

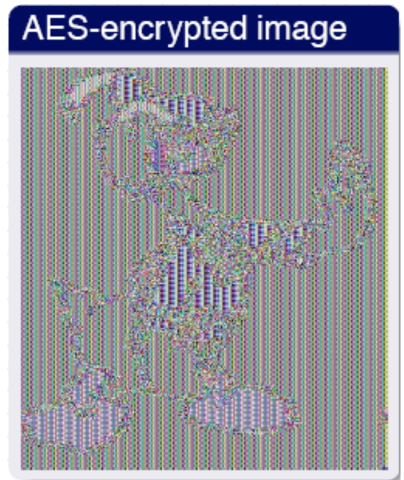

*Uncompressed bitmap format encrypted using AES in ECB mode.

Cipher block chaining (CBC) mode

Cipher block chaining. We choose a random initial vector (IV) of length n. The ciphertext is obtained by applying the pseudorandom permutation to the XOR of the current plaintext block and the previous ciphertext block. That is, we set $c_0 = IV$ and then, for $i = 1$ to ℓ, set $c_i = F_k(c_{i-1} \oplus m_i)$.

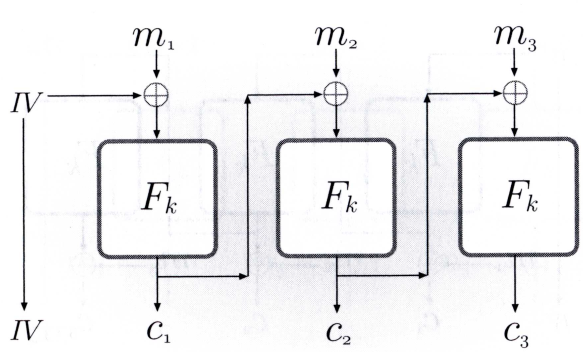

Donald disappears*

- Encryption in CBC is probabilistic and it has been proven that if F is a pseudorandom permutation then CBC-mode encryption is CPA-secure.
- All is not rosie in cipherland however: Encryption must be carried out sequentially. If parallel processing is available CBC may not be the most efficient choice.

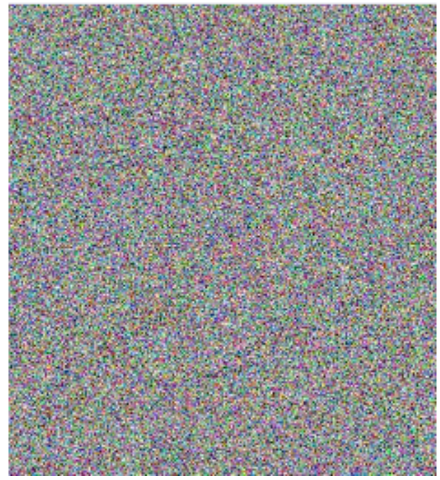

*Uncompressed bitmap format encrypted using AES in CBC mode.

Output feedback (OFB) mode

Output Feedback mode. Again a random initial vector (IV) of length n is chosen and a stream is generated from IV as follows: Define $r_0 := IV$ and set the ith block of the stream $r_i = F_k(r_{i-1})$. Then, for $i = 1$ to ℓ, set $c_i = m_i \oplus r_i$.

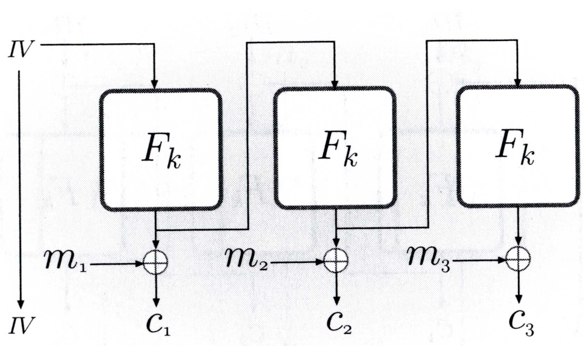

Donald stays away

- This mode is also probabilistic and it and can be shown to be CPA-secure.
- Both encryption and decryption must be carried out sequentially, but the bulk of the computation* can be carried out independently of the message.

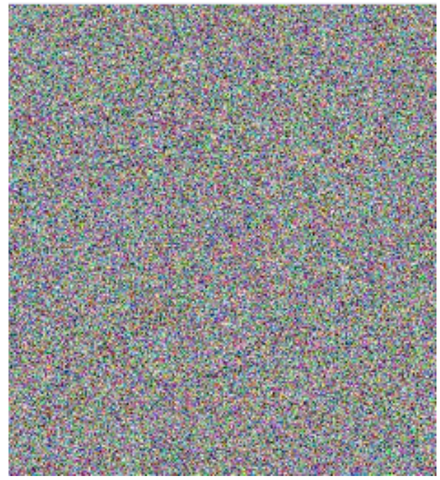

*Namely computing the pseudorandom stream.

Counter (CTR) mode

Randomized counter mode. A random initial vector (IV) of length n is chosen, this is referred to as ctr. Then a stream is generated from IV by computing $r_i := F_k(\text{ctr} + i)$ where ctr and i are viewed as binary numbers and addition is performed modulo 2^n. Finally, the ith ciphertext block is computed as $c_i = m_i \oplus r_i$.

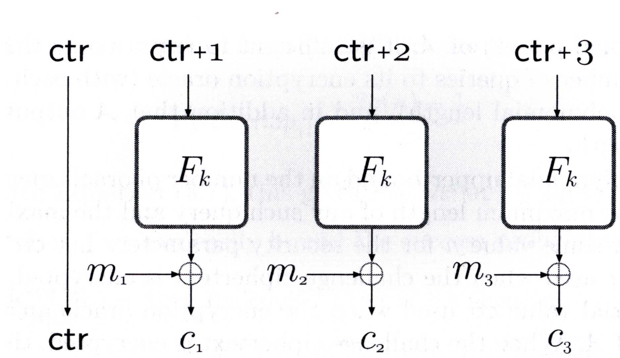

*There are a number of different types of counter modes.

Donald gets tired of waiting in the wings

- Again this mode is probabilistic and it and can be shown to be CPA-secure.
- Both encryption and decryption can be fully parallelized and, as with OFB mode, it is possible to generate the pseudorandom stream ahead of time.
- Finally, it is possible to decrypt the ith block of ciphertext without decrypting anything else*.

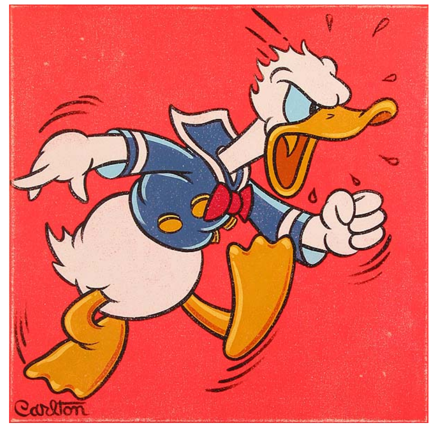

*A property known as *random access*.

Randomized counter (CTR) mode is CPA-secure

Theorem 3.32. If F is a pseudorandom function, then randomized counter mode has indistinguishable encryption under a chosen-plaintext attack.

Proof. As previously, we prove counter mode is CPA-secure when a truly random function is used. We then prove that replacing the random function by a pseudorandom function cannot make the scheme insecure.

Let ctr^* denote the initial value used when the challenge ciphertext is encrypted in the $\mathsf{PrivK}^{\mathsf{cpa}}$ experiment. When a random function is used in CTR mode, security is achieved as long as each block c_i is encrypted using a value $\mathrm{ctr}^* + i$ that was never used by the encryption oracle to answer any of its queries since in this case $f(\mathrm{ctr}^* + i)$ is completely random.

A first step

Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ denote the randomized counter mode encryption scheme, and let $\widetilde{\Pi} = (\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{Enc}}, \widetilde{\mathsf{Dec}})$ be the encryption scheme that is identical to Π except that a truly random permutation f is used in place of F_k.

Fix an arbitrary PPT adversary $\mathcal{A}$, and let $q(n)$ be a polynomial-upper bound on the number of oracle queries made by $\mathcal{A}(1^n)$ as well as on the maximum number of blocks of any such query and the maximum number of blocks in m_0, m_1.

First note that as in the proof of Theorem 3.31 there exists a negligible function negl such that

$$\left|\Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n) = 1\right] - \Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n) = 1\right]\right| \leq \mathsf{negl}.$$

The tilde experiment

Next we claim that

$$\Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n) = 1\right] \leq \frac{1}{2} + \frac{2q(n)^2}{2^n}.$$

Combining this with the previous inequality we see that

$$\Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n) = 1\right] \leq \frac{1}{2} + \frac{2q(n)^2}{2^n} + \mathsf{negl}(n).$$

Since q is a polynomial, $\frac{2q(n)^2}{2^n}$ is negligible and we're done (well done once we establish the claim.

The adversary is polynomially bounded

Let $\ell^* \leq q(n)$ denote the length (in blocks) of the messages m_0, m_1 output by $\mathcal{A}$ in experiment $\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n)$, and let ctr^* denote the initial value used when the challenge ciphertext is encrypted.

Similarly, let $\ell_i \leq q(n)$ be the length (in blocks) of the ith encryption-oracle query and ctr_i denote the initial value used when answering this query.

When the challenge ciphertext is encrypted, f is applied to the values

$$\mathsf{ctr}^* + 1, \ldots, \mathsf{ctr}^* + \ell^*$$

while when the ith is answered, the function f is applied to the values

$$\mathsf{ctr}_i + 1, \ldots, \mathsf{ctr}_i + \ell_i.$$

There are two cases to consider

1. *Suppose there do not exist any $i, j, j' \geq 1$ for which $\mathsf{ctr}_i + j = \mathsf{ctr}^* + j'$.* Then, the values $f(\mathsf{ctr}^* + 1), \ldots, f(\mathsf{ctr}^* + \ell^*)$ are independently and uniformly distributed since f was not applied to any of these when encrypting oracle queries.
   The challenge text is encrypted with a random string and the probability that $\mathcal{A}$ outputs $b' = b$ is exactly $1/2$ as in the one-time pad.
2. *There exists $i, j, j' \geq 1$ for which $\mathsf{ctr}_i + j = \mathsf{ctr}^* + j'$.* In this case $\mathcal{A}$ has it made in the shade since it can easily determine the value $f(\mathsf{ctr}_i + j) = f(\mathsf{ctr}^* + j')$ from the answer to its ith oracle query. We analyze the probability that this occurs.

Probability of overlaps continued

The probability is as large as possible when $\ell^* = \ell_i = q(n)$ for all i. Let $\mathrm{Overlap}_i$ denote the event that the sequence $\mathrm{ctr}_i + 1, \ldots, \mathrm{ctr}_i + q(n)$ overlaps with $\mathrm{ctr}^* + 1, \ldots, \mathrm{ctr}^* + q(n)$ and let Overlap denote the event that $\mathrm{Overlap}_i$ occurs from some i. By the union bound

$$\Pr[\mathrm{Overlap}] \leq \sum_{i=1}^{q(n)} \Pr[\mathrm{Overlap}_i].$$

Probability of overlaps

Fixing ctr^*, event $\mathrm{Overlap}_i$ occurs exactly when ctr_i satisfies

$$\mathrm{ctr}^* + 1 - q(n) \leq \mathrm{ctr}_i \leq \mathrm{ctr}^* + q(n) - 1.$$

Since there are $2q(n) - 1$ values of ctr_i for which $\mathrm{Overlap}_i$ can occur,

$$\Pr[\mathrm{Overlap}_i] = \frac{2q(n)-1}{2^n} < \frac{2q(n)}{2^n}.$$

Combining this with the union bound given on the previous slide, we have

$$\Pr[\mathrm{Overlap}] \leq \sum_{i=1}^{q(n)} \frac{2q(n)}{2^n} = \frac{2q(n)^2}{2^n}.$$

Bounds on the success probability

Given the above, we can now bound the success probability of $\mathcal{A}$

$$
\begin{aligned}
\Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n)=1\right] &= \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\widetilde{\Pi}}(n)=1 \wedge \mathsf{Overlap}] \\
&\quad + \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\widetilde{\Pi}}(n)=1 \wedge \overline{\mathsf{Overlap}}] \\
&\le \Pr[\mathsf{Overlap}] + \Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\widetilde{\Pi}}(n)=1 \mid \overline{\mathsf{Overlap}}] \\
&\le \frac{2q(n)^2}{2^n} + \frac{1}{2}.
\end{aligned}
$$

$\square$

*Block length and security**

Remark. If an input to the block cipher is used more than once then security can be violated. Thus, it is not only the *key length* of a block cipher than is important, but also its *block length*

Example. Suppose we use a block cipher with block length 64-bits. Even if a completely random function with this block length is used, an adversary can achieve success with probability roughly $\frac{1}{2} + \frac{q^2}{2^{63}}$ in a chosen-plaintext attack with it makes q queries, each q blocks long.

*Ghost of Disney past.

Bad news for 

Plaintext	Bois	owes	Bullwinkle	$10,000
Ciphertext	tsrwmqp	wplsmaka	os02%sb@	gwom*bxz

Altered Ciphertext	os02%sb@	wplsmaka	tsrwmqp	gwom*bxz
Ciphertext	Bullwinkle	owes	Bois	$10,000

*Well yes, but not our job. Issues of *message integrity* or *message authentication* should be dealt with separately from encryption.