University of Maryland
CMSC456 — Introduction to Cryptography
Professor Jonathan Katz

Lecture 17

1 Modes of Encryption

In the example of the previous lecture, encryption was "technically" only defined for messages of length n (where our keyed function F was such that $F : \{0,1\}^k \times \{0,1\}^m \rightarrow \{0,1\}^n$). But it should be clear from the fact that the scheme above achieves security in the sense of indistinguishability — both on an intuitive and on a formal level — that the scheme immediately extends to give a secure encryption scheme for arbitrary-length messages. In fact, one could state the following general theorem:

Theorem 1 *(Informally:) Given an encryption scheme $\mathcal{E}$ for n-bit messages which is secure in the sense of left-or-right indistinguishability, we can construct an encryption scheme $\mathcal{E}'$ for arbitrary-length messages which is secure in the sense of left-or-right indistinguishabiity. Namely, $\mathcal{E}'_s(M)$ parses M as $M_1, \ldots, M_\ell$ (where $|M_i| = n$) and outputs $\mathcal{E}_s(M_1) \circ \mathcal{E}_s(M_2) \circ \cdots \circ \mathcal{E}_s(M_\ell)$.*

Thus, using our original encryption scheme for n-bit messages, we may encrypt arbitrary-length messages M as follows (as usual, s is the shared key):

Parse M as $M_1, M_2, \ldots M_\ell$, where $|M_i| = n$
For $i = 1$ to ℓ:
  $r_i \leftarrow \{0,1\}^m$
  $C_i = F_s(r_i) \oplus M_i$
Output $r_1, C_1, \ldots, r_\ell, C_\ell$

(In our entire discussion above, we assume that $|M|$ is a multiple of n; this restriction is easy to remove.)

Note: We need to be careful with how we define security in the sense of left-or-right indistinguishability when encryption of arbitrary-length messages is allowed. As usual, we allow the adversary to access the left-or-right oracle LR as many times as it likes. However, we *must* restrict the adversary as follows: when the adversary submits a query (M_0, M_1) to the oracle we require that $|M_0| = |M_1|$. (In the encryption scheme above, it should be clear that it is easy to break the encryption scheme if this restriction is removed; in fact, it is *impossible* to construct any encryption scheme which is secure when the adversary is allowed to submit two different-length messages.) We stress that we impose *no* restriction on the lengths of messages in *different* queries; thus, the first query (M_0, M_1) and the second query (M'_0, M'_1) can have $|M_0| \neq |M'_0|$.

While the scheme above is secure in the sense of indistinguishability, it is not entirely satisfying since the ciphertext is twice as long as the message for typical F (typical block ciphers have $m = n$). The natural question is: can we do better? This brings us to the subject of *modes of encryption* which are techniques used to encrypt long messages using a block cipher/PRP *with fixed input lengths* as a building block.

In all of what follows we assume a PRF $F : \{0,1\}^k \times \{0,1\}^m \to \{0,1\}^m$ and the message M to be encrypted is parsed into a sequence of m-bit blocks $M_1, \ldots, M_\ell$. Almost the most natural suggestion for a mode of encryption is the following:

$$\mathcal{E}_s(M) = F_s(M_1) \circ \cdots \circ F_s(M_\ell).$$

This mode (known as *electronic codebook* or *ECB* mode) was actually one of the four modes of encryption specified along with the DES block cipher. Note that for decryption to work properly, F must be a keyed *permutation* and furthermore must be easy to invert (given s).

This mode seems good, since the ciphertext is the same length as the message and hence there is no ciphertext expansion! But...is this scheme secure? It most certainly is not secure in the sense of indistinguishability, since it is a *deterministic* encryption scheme. In fact, it does not even achieve the weaker form of security (cf. previous lectures) where the adversary is allowed only a single query to the LR oracle. Do you see why?

This has led to numerous other suggestions (in fact, designing "good" modes of encryption is an active area of research). We list here two other modes defined as part of the DES standard.

Cipher-block chaining (CBC) mode.

$\mathcal{E}_s(M)$
  $C_0 \leftarrow \{0,1\}^m$
  For $i = 1$ to ℓ
    $C_i = F_s(C_{i-1} \oplus M_i)$
  Output $C_0, C_1, \ldots, C_\ell$

$\mathcal{D}_s(C_0, \ldots, C_\ell)$
  For $i = 1$ to ℓ
   $M_i = C_{i-1} \oplus F_s^{-1}(C_i)$
  Output $M_1, \ldots, M_\ell$

The random block C_0 is called an *initialization vector*, and is chosen randomly each time a new message is encrypted. Since this is a *randomized* encryption scheme, it might potentially be secure in the sense of indistinguishability. In fact, this *is* a secure mode of encryption; see [1] for details.

The ciphertext here is longer than the message by an *additive* factor (as opposed to the multiplicative factor we had in our initial mode of encryption). In fact, the ciphertext is longer than the message by m bits which is essentially optimal (a scheme with no ciphertext expansion would be deterministic and hence insecure). Again, we need here for F to be a keyed permutation, and it must additionally be efficient to invert.

Cipher feedback (CFB) mode.

$\mathcal{E}_s(M)$
  $C_0 \leftarrow \{0,1\}^m$
  For $i = 1$ to ℓ
    $C_i = M_i \oplus F_s(C_{i-1})$
  Output $C_0, C_1, \ldots, C_\ell$

$\mathcal{D}_s(C_0, \ldots, C_\ell)$
For $i = 1$ to ℓ
$M_i = C_i \oplus F_s(C_{i-1})$
Output $M_1, \ldots, M_\ell$

Again we have a random initialization vector chosen each time a new message is encrypted. Like CBC mode, this scheme is secure and the ciphertext is longer than the message by only m bits. An advantage of this scheme over CBC mode is that F does not need to be a keyed permutation (in fact, we never need to invert F in order to decrypt). This is useful in some contexts for reasons of efficiency.

Counter mode. Another mode is motivated by the example with which we opened this section. Recall that our initial suggestion was to encrypt M as:

$$r_1 \circ (M_1 \oplus F_s(r_1)) \circ \cdots \circ r_\ell \circ (M_\ell \oplus F_s(r_\ell)),$$

but this has the disadvantage of having a ciphertext twice as long as the plaintext. But there is no need for the r_i to all be different! (New random values should be chosen every time a message is encrypted; but one random value per message might be enough.) This suggests the following improvement: pick random r and then compute the ciphertext as:

$$r \circ (M_1 \oplus F_s(r)) \circ (M_2 \oplus F_s(r+1)) \circ \cdots \circ (M_\ell \oplus F_s(r+\ell-1)).$$

This scheme has additive expansion factor, and is also secure. See [1] for further details.

References

[1] M. Bellare, A. Desai, E. Jokipii, and P. Rogaway. A Concrete Security Treatment of Symmetric Encryption: Analysis of the DES Modes of Operation. FOCS '97. Available from `http://www-cse.ucsd.edu/users/mihir/papers/sym-enc.html`.