

Block Ciphers

Block Ciphers and Stream Ciphers

In practical ciphers the plaintext M is divided into fixed-length blocks $M = M_1M_2 \dots M_N$. Then, each block M_i is encrypted to the ciphertext block $C_i = E_K(M_i)$, and the results are concatenated to the ciphertext $C = C_1C_2 \dots C_N$.

There are two major kind of ciphers, which differ in the way the plaintexts are encrypted:

Stream Ciphers

The blocks are encrypted sequentially, each block is encrypted by a distinct transformation, which might depend on

1. the previous encrypted blocks,
2. the previous transformation,
3. the block number,
4. the key.

This information from one block is kept in **memory** between the encryption of this block and the succeeding block, for use during the encryption of the succeeding block.

Usually, stream ciphers use **blocks of either one bit or eight bits** (one character).

Block Ciphers

All the blocks are encrypted in the same way, under exactly the same transformation (no memory): $C_1 = E(M_1)$, $C_2 = E(M_2)$, etc.

Encryption transformation should not be vulnerable to known plaintext attacks. Attacker should not be able to collect (almost) all the plaintext/ciphertext blocks pairs, keep the transformation table $T(M) = C$, and use it to encrypt/decrypt if they do not know the mathematical formulation of the transformation (and in particular the key).

Thus, **the block size should be large**, and the number of distinct possible values in a plaintext block should be larger than the minimal allowed complexity of an attack.

In the past **blocks of 64 bits** were used, which have 2^{64} possibilities, whose table storing costs at least 2^{64} known plaintexts and memory space.

Nowadays, the standard block size is 128 bits.

Block Ciphers

Block ciphers are substitution ciphers in which the plaintext and the ciphertext blocks are binary vectors of length N . When $N = 64$ there are 2^{64} different plaintexts/ciphertexts, and when $N = 128$ there are 2^{128} different plaintexts/ciphertexts.

For each key the encryption function $E_K(\cdot)$ is a permutation from $\{0, 1\}^N$ to itself.

$D_K(\cdot)$ is the decryption function (the inverse permutation), such that $D_K(E_K(\cdot)) = E_K(D_K(\cdot)) = \text{Identity}$.


The Data Encryption Standard - DES

1. The most widely used cipher in civilian applications.
2. Developed by IBM; Evolved from Lucifer.
3. Accepted as an US NBS standard in 1977, and later as an international standard.
4. A block cipher with **$N = 64$ bit blocks**.
5. **56-bit keys** (eight bytes, in each byte seven bits are used; the eighth bit can be used as a parity bit).
6. Exhaustive search requires 2^{56} encryption steps (2^{55} on average).

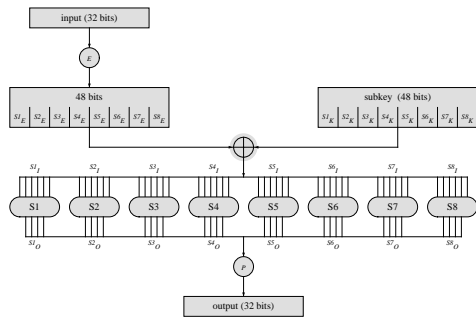
The Data Encryption Standard - DES (cont.)

7. Iterates a round-function 16 times in 16 **rounds** (רונדים). The round-function mixes the data with the key.
8. Each round, the key information entered to the round function is called a **subkey**. The subkeys K_1, \dots, K_{16} are computed by a **key scheduling algorithm**.

DES Outline



The F-Function



The Initial Permutation (IP)

The following tables describe for each output bit the number of the input bit whose value enters to the output bit. For example, in *IP*, the 58'th bit in the input becomes the first bit of the output.

IP:	FP=IP ⁻¹ :
58	40
50	8
42	48
34	16
26	56
18	24
10	64
2	32
60	39
52	7
44	47
36	15
28	55
20	23
12	63
4	31
62	38
54	6
46	46
38	14
30	54
22	22
14	62
6	30
64	37
56	5
48	45
40	13
32	53
24	21
16	61
8	29
57	36
49	44
41	12
33	52
25	20
17	60
9	28
59	35
51	43
43	11
35	51
27	19
19	59
11	27
61	34
53	2
45	42
37	10
29	50
21	18
13	58
5	26
63	33
55	41
47	9
39	49
31	17
23	57
15	25
7	63

The P Permutation and the E Expansion

P Permutes the order of 32 bits. *E* Expands 32 bits to 48 bits by duplicating 16 bits twice.

P:	E:
16	32
7	1
20	2
21	3
29	4
12	5
28	17
17	4
1	5
15	6
23	7
26	8
5	9
18	10
31	11
10	12
2	13
8	14
24	15
14	16
32	17
27	18
3	19
9	20
19	21
13	22
30	23
6	24
22	25
11	26
4	27
25	28
	29
	28
	29
	30
	31
	32
	1

The S Boxes

S box S1:

14	4	13	1	2	15	11	8	3	10	6	12	5	9	0	7
0	15	7	4	14	2	13	1	10	6	12	11	9	5	3	8
4	1	14	8	13	6	2	11	15	12	9	7	3	10	5	0
15	12	8	2	4	9	1	7	5	11	3	14	10	0	6	13

S box S2:

15	1	8	14	6	11	3	4	9	7	2	13	12	0	5	10
3	13	4	7	15	2	8	14	12	0	1	10	6	9	11	5
0	14	7	11	10	4	13	1	5	8	12	6	9	3	2	15
13	8	10	1	3	15	4	2	11	6	7	12	0	5	14	9

The S Boxes (cont.)

S box S3:

10	0	9	14	6	3	15	5	1	13	12	7	11	4	2	8
13	7	0	9	3	4	6	10	2	8	5	14	12	11	15	1
13	6	4	9	8	15	3	0	11	1	2	12	5	10	14	7
1	10	13	0	6	9	8	7	4	15	14	3	11	5	2	12

S box S4:

7	13	14	3	0	6	9	10	1	2	8	5	11	12	4	15
13	8	11	5	6	15	0	3	4	7	2	12	1	10	14	9
10	6	9	0	12	11	7	13	15	1	3	14	5	2	8	4
3	15	0	6	10	1	13	8	9	4	5	11	12	7	2	14

The S Boxes (cont.)

S box S5:

2	12	4	1	7	10	11	6	8	5	3	15	13	0	14	9
14	11	2	12	4	7	13	1	5	0	15	10	3	9	8	6
4	2	1	11	10	13	7	8	15	9	12	5	6	3	0	14
11	8	12	7	1	14	2	13	6	15	0	9	10	4	5	3

S box S6:

12	1	10	15	9	2	6	8	0	13	3	4	14	7	5	11
10	15	4	2	7	12	9	5	6	1	13	14	0	11	3	8
9	14	15	5	2	8	12	3	7	0	4	10	1	13	11	6
4	3	2	12	9	5	15	10	11	14	1	7	6	0	8	13

The S Boxes (cont.)

S box S7:

4	11	2	14	15	0	8	13	3	12	9	7	5	10	6	1
13	0	11	7	4	9	1	10	14	3	5	12	2	15	8	6
1	4	11	13	12	3	7	14	10	15	6	8	0	5	9	2
6	11	13	8	1	4	10	7	9	5	0	15	14	2	3	12

S box S8:

13	2	8	4	6	15	11	1	10	9	3	14	5	0	12	7
1	15	13	8	10	3	7	4	12	5	6	11	0	14	9	2
7	11	4	1	9	12	14	2	0	6	10	13	15	3	5	8
2	1	14	7	4	10	8	13	15	12	9	0	3	5	6	11

The S Boxes (cont.)

How to interpret the S boxes:

The representation of the S boxes use the first and sixth bits of the input as a line index (between 0 and 3), and the four middle bits as the row index (between 0 and 15).

Thus, the input values which correspond to the standard description of the S boxes are

0	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30
1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31
32	34	36	38	40	42	44	46	48	50	52	54	56	58	60	62
33	35	37	39	41	43	45	47	49	51	53	55	57	59	61	63

The S Boxes (cont.)

Note that **all the operations are linear, except for the S boxes**. Thus, **the strength of DES crucially depends on the choice of the S boxes**.

If the S boxes would be affine, the cipher becomes affine, and thus easily breakable.

The S boxes were chosen with some criteria to prevent attacks.

The Key Scheduling Algorithm

The key scheduling algorithm generates the 16 48-bit subkeys from the 56-bit key, by duplicating each key bit into about 14 of the subkeys in a particular order.

PC-1:

57	49	41	33	25	17	9
1	58	50	42	34	26	18
10	2	59	51	43	35	27
19	11	3	60	52	44	36
63	55	47	39	31	23	15
7	62	54	46	38	30	22
14	6	61	53	45	37	29
21	13	5	28	20	12	4

The Key Scheduling Algorithm (cont.)

Number of rotations in the key scheduling algorithm:

Round	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Rotations	1	1	2	2	2	2	2	2	1	2	2	2	2	2	2	1

PC-2:

14	17	11	24	1	5
3	28	15	6	21	10
23	19	12	4	26	8
16	7	27	20	13	2
41	52	31	37	47	55
30	40	51	45	33	48
44	49	39	56	34	53
46	42	50	36	29	32

Decryption

Decryption is done by the **same algorithm** as encryption, except that the order of the subkeys is reversed (i.e., K16 is used instead of K1, K15 instead of K2, ..., and K1 instead of K16).

Known Weaknesses

- Secret design criteria:** IBM were requested by the NSA not to publish the design criteria (later known as differential cryptanalysis).
- Weak keys:** four keys for which $E_K(E_K(X)) = X$ for any X (i.e., $E_K(\cdot) \equiv D_K(\cdot)$).
- Semi-weak keys:** 12 keys (6 pairs) for which $E_{K_1}(E_{K_2}(X)) = X$ for any X (i.e., $E_{K_1}(\cdot) \equiv D_{K_2}(\cdot)$).
- The S boxes are not random:** for example,
 - S4 is 75% redundant: Each line of S4 can be calculated from any other line by a linear function.
 - In all the S boxes ($S(x_1x_2x_3x_4x_5x_6) = y_1y_2y_3y_4$) there is a high correlation between the bit x_2 of the input to the parity of the output $y_1 \oplus y_2 \oplus y_3 \oplus y_4$ (For example, in S5 there are 12 equalities and 52 inequalities).

Known Weaknesses (cont.)

- Short 56-bit keys:** On a single processor exhaustive search takes:
 - In the 70's: thousands of years.
 - In the 90's: On DEC GaAs DES chip: 100 years.

Exhaustive search is not totally infeasible:

 - In the 70's: Diffie and Hellman suggested a 20,000,000\$ special purpose machine which can find a key within a day. They estimate the cost per solution to be 5000\$.
 - In the 90's: Wiener suggested a 1,000,000\$ special purpose machine which can find a key within 3.5 hours.
 - June 1997: Done over the *Internet* using 14000-78000 computers within 90 days.
 - January 1998: Done over the Internet within 39 days.
 - July 1998: A \$210000 search machine, called *deep crack*, was built. Found a key within 56 hours.

The Complementation Property

Let \bar{X} be the 1-complement of X.

In DES if

$$C = E_K(P)$$

then

$$\bar{C} = E_{\bar{K}}(\bar{P}).$$

Exercise: Prove the complementation property.

The Complementation Property (cont.)

This property can be used for a chosen plaintext attack, twice faster than exhaustive search:

- Choose P .
- Ask for the ciphertexts $C_0 = E_K(P)$ and $C_1 = E_K(\bar{P})$.
- Exhaustively try all the 2^{55} keys K' whose most significant bit is zero, whether $E_{K'}(P) \in \{C_0, \bar{C}_1\}$.
- if $E_{K'}(P) = C_0$, then probably $K = K'$.
- if $E_{K'}(P) = \bar{C}_1$, then probably $K = \bar{K}'$.
- Otherwise, neither K' nor \bar{K}' are the key K .
- Since comparisons are much faster than a trial encryptions, this attack is twice faster than exhaustive search.

Contemporary Attacks on DES

Today there are three methods which can reduce the complexity of attacks to less than 2^{55} . All of them require a huge (unrealistic) amount of data (their complexity is the time to create the data, while the analysis is much faster).

Differential cryptanalysis: 2^{47} chosen plaintexts (1990–1991, Biham and Shamir).

Linear cryptanalysis: 2^{43} known plaintexts (1993–1994, Matsui). Slightly improved to $0.75 \cdot 2^{43}$ (1998, Shimoyama, Kaneko).

Improved Davies attack: 2^{50} known plaintexts (1994, Biham and Biryukov).

Statistical Cryptanalysis: $2^{42.9}$ known plaintexts (1996, Vaudenay).

Contemporary Attacks on DES (cont.)

Related Keys: (1993, Biham) proposes that the shifts in the key scheduling algorithm should not be all equal. If they were all equal (like all 2, or all 7) then DES could be broken with either

- $1.43 \cdot 2^{53}$ steps, or
- 2^{33} chosen key known plaintexts, or
- 2^{17} chosen key chosen plaintexts.

Key-Collision Attacks: (Biham, 1996) In some circumstances it is much easier to “forge” messages than to “break” the cipher: 2^{28} .

Other Works

1. Method of Formal Coding (Schaumüller-Bichl)
2. Linear factors (Chaum, Evertse)

Modes of Operation

Long messages of several blocks are encrypted using one of the **modes of operation**. In the modes of operation, the messages M are divided into N -bit blocks $M_1M_2 \dots M_n$, each block M_i is encrypted (as defined by the mode of operation) to C_i , and the results are concatenated into the ciphertext $C = C_1C_2 \dots C_n$.

Many modes of operation actually transform the block ciphers into stream ciphers, by adding memory (external to the block cipher).


Electronic Code Book (ECB) Mode

In this mode, each plaintext block M_i is encrypted simply by

$$C_i = E_K(M_i).$$

Decryption is done by

$$M_i = D_K(C_i).$$



The main drawback of ECB is that $M_i = M_j$ iff $C_i = C_j$. Thus, Eve can easily identify which plaintext blocks are equal.


Cipher Block Chaining (CBC) Mode

Each plaintext block M_i is mixed with the previous ciphertext block before encryption:

$$C_i = E_K(M_i \oplus C_{i-1}).$$

Decryption is done by

$$M_i = D_K(C_i) \oplus C_{i-1}.$$



A (non-secret) initial value $C_0 = IV$ is chosen for each message.

Cipher Block Chaining (CBC) Mode (cont.)


In this mode two equal message blocks are usually encrypted to different ciphertext blocks.

This mode has small error propagation: if C_i is received with errors, only M_i and M_{i+1} have errors after decryption. M_{i+2} is not affected from the error in C_i .

(64-bit) Output FeedBack (OFB) Mode

Generates a pseudo random bit stream from the key K and an initial value $V_0 = IV$ by

$$V_i = E_K(v_{i-1}).$$



Encryption is done by

$$C_i = M_i \oplus V_i.$$

Decryption is done by

$$M_i = C_i \oplus V_i.$$

(64-bit) Output FeedBack (OFB) Mode (cont.)

This mode has no error propagation (one bit error in C causes only one bit error in the message during decryption).

An advantage of this mode is that V_i can be computed in advance, before the plaintext/ciphertext is known.


(64-bit) Cipher FeedBack (CFB) Mode

Similar to the OFB mode, but the bit stream depends on the ciphertext. Encryption is done by

$$C_i = M_i \oplus E_K(C_{i-1}).$$

Decryption is done by

$$M_i = C_i \oplus E_K(C_{i-1}).$$



Triple-DES

Triple-DES is an improvement to DES which use the same DES hardware/software but encrypt the plaintext three times.

Triple DES has two variants:

1. **Three-key triple-DES:** the plaintext is encrypted three times under three different DES keys K_1 , K_2 , and K_3 .

$$C = \text{DES}_{K_3}(\text{DES}_{K_2}^{-1}(\text{DES}_{K_1}(P))).$$

The total key length is $3 \cdot 56 = 168$ bits.

2. **Two-key triple-DES:** the plaintext is encrypted three times under two different DES keys K_1 , and K_2 , where K_1 is used in the first and third application of DES, and K_2 in the second application.

$$C = \text{DES}_{K_1}(\text{DES}_{K_2}^{-1}(\text{DES}_{K_1}(P))).$$

The total key length is $2 \cdot 56 = 112$ bits.

Triple-DES (cont.)

Note that the second application of DES performs decryption. This is done to allow compatibility to older systems which use DES: if the two or three keys are all equal, then the triple encryption actually performs a single encryption with that key.

Speed:


Triple-DES is about three times slower than DES. (It is slightly faster than a third of the speed of DES as the initial and final permutation in the borders between the first/second and the second/third DES applications can be eliminated).

Status:

Triple-DES replaced DES as the de-facto standard a few years ago. AES is now the standard for new applications, while Triple-DES remains in many old applications.

Triple-DES (cont.)

Two-key and three-key Triple-DES:



Strength of Triple-DES

Triple-DES is believed to be secure for any practical purpose.

Its key size of 112 and 168 bits ensures that brute force (exhaustive search) attacks are impossible.


Other attacks are eliminated due to the large number of rounds.

The best theoretical known attack on Three-key Triple-DES has complexity 2^{108} .

Two-key Triple-DES has a theoretical attack with complexity 2^{56} (2^{56} time and 2^{56} chosen plaintext blocks). Although this attack is impractical, it is preferable to use Three-key Triple-DES.


Triple-DES Modes of Operation

Triple-DES can be used with all the standard modes of operations of DES (ECB, CBC, OFB, CFB), where DES is replaced by Triple-DES as the underlying cipher. For example, the ECB mode (TECB):



Triple-DES Modes of Operation (cont.)

And the CBC mode (TCBC):



Triple-DES Modes of Operation (cont.)

In addition Triple-DES has some additional interleaved CBC and CFB modes of operations, which are used to improve efficiency of encryption, by removing the sequentiality of the CBC (and CFB) modes, and using three times the DES hardware. These modes actually perform three parallel TCBC modes, each encrypt every third block of the original plaintext.

Idea

Idea (The International Data Encryption Algorithm) is a block cipher designed by Lai and Massey, as an improvement of their PES cipher.

Idea encrypts 64-bit blocks under 128-bit keys. The blocks are divided into four 16-bit words. In the first halves of the rounds (and the last half-round of the cipher) each word is mixed with key material. In the second halves of the rounds, the words are mixed together, with some additional key material.


Idea uses three operations on 16-bit words:

1. XOR
2. Addition modulo 2^{16}
3. Multiplication modulo the prime $2^{16} + 1$, where zero is treated as 2^{16} .

The combination of the operations ensures fast mixing of the plaintext and key.

Idea (cont.)

The following figure describes one round of Idea:



The best known attack on Idea analyzes up to 4.5 rounds faster than exhaustive search.

Skipjack and the Clipper Chip

In 1993 the US government started an initiative to give the industry an NSA-developed cipher, called **Skipjack**, with a military strength (used for non-secret information). The catch is that the design of the cipher will not be published, and that the users of the system will have to disclose their keys to the government agencies if the court decides to. To ensure that, the cryptographic chip, called the **Clipper chip**, will send the key encrypted under a secret key which can be recovered if several government agencies collaborate (but not without collaboration).


This initiative started a long debate, and did not survive.

Skipjack encrypts blocks of 64 bits, and have 80-bit keys.

In June 1998, Skipjack was published by NIST in order to allow implementation in software. Consecutively, it was analyzed, and shown that a reduced cipher with 31 rounds, rather than the 32 rounds of Skipjack, is theoretically breakable faster than exhaustive search.

Skipjack and the Clipper Chip (cont.)

The first half of Skipjack: (the second half is similar)



The Advanced Encryption Standard (AES)

AES is the new standard replacement for DES. It has blocks of 128 bits. It supports key sizes of 128, 192, and 256 bits.

In 1997 NIST (the institute which standardized DES) called for proposals for a successor for DES. 15 proposals were submitted from all over the world, and in August 1999 the list was reduced to 5 finalists:

- Mars (IBM; USA)
- RC6 (RSA data security; USA)
- Rijndael (Daemen, Rijmen; Belgium)
- Serpent (Anderson, Biham, Knudsen; UK, Israel, Norway)
- Twofish (Counterpane; USA)

In October 2000 **Rijndael** was selected to be the AES, and in November 2001 it became a formal standard (FIPS 197).

The AES – Rijndael

1. A block cipher with $N = 128$ bit blocks.
2. Supports three key sizes: **128 bits**, **192 bits**, and **256 bits**
3. Exhaustive search requires 2^{128} (2^{192} , 2^{256} , respectively) encryption steps (2^{127} , 2^{191} , 2^{255} on average).
4. Iterates a round-function 10–14 times in 10–14 **rounds** (רונדן). The round-function mixes the data with the key.
5. Each round, the key information entered to the round function is called a **subkey**. The subkeys K_0, \dots, K_{10} (or K_0, \dots, K_{14}) are computed by a **key scheduling algorithm**.

Rijndael

Let the 128-bit blocks be divided into 32-bit words, and each word be divided into bytes.

Let the block be viewed in a rectangle (square) of 4×4 bytes, where each column have the bytes of a word

0	4	8	12
1	5	9	13
2	6	10	14
3	7	11	15

128-bit keys are viewed in the same way. Longer keys have additional columns, e.g., 256-bit keys

0	4	8	12	16	20	24	28
1	5	9	13	17	21	25	29
2	6	10	14	18	22	26	30
3	7	11	15	19	23	27	31

Rijndael (cont.)

Rijndael uses four invertible operations:

1. Byte substitution (S box; 8-bit to 8-bit)
2. Shift row (rotating order of bytes in each row)
3. Mix column (linear mixing of a word column)
4. Key mixing

Byte Substitution

The byte substitution is a fixed S box from 8 bits to 8 bits. It substitutes all the bytes of the input by the following table

```

words S[256] = {
  99, 124, 119, 123, 242, 107, 111, 197, 48, 1, 103, 43, 254, 215, 171, 118,
  202, 130, 201, 125, 250, 89, 71, 240, 173, 212, 162, 175, 156, 164, 114, 192,
  183, 253, 147, 38, 54, 63, 247, 204, 52, 165, 229, 241, 113, 216, 49, 21,
  4, 199, 35, 195, 24, 150, 5, 154, 7, 18, 128, 226, 235, 39, 178, 117,
  9, 131, 44, 26, 27, 110, 90, 160, 82, 59, 214, 179, 41, 227, 47, 132,
  83, 209, 0, 237, 32, 252, 177, 91, 105, 303, 190, 57, 74, 76, 88, 207,
  208, 239, 170, 251, 67, 77, 51, 133, 69, 249, 2, 127, 80, 60, 159, 168,
  81, 163, 64, 143, 146, 157, 56, 245, 188, 182, 218, 33, 16, 255, 243, 210,
  205, 12, 19, 236, 95, 151, 68, 23, 196, 167, 126, 61, 100, 93, 25, 115,
  96, 129, 79, 220, 34, 42, 144, 136, 70, 238, 184, 20, 222, 94, 11, 219,
  224, 50, 58, 10, 73, 6, 36, 92, 194, 211, 172, 98, 145, 149, 228, 121,
  231, 200, 55, 109, 141, 213, 78, 169, 108, 86, 244, 234, 101, 122, 174, 8,
  186, 120, 37, 46, 28, 166, 180, 198, 232, 221, 116, 31, 75, 189, 139, 138,
  112, 62, 181, 102, 72, 3, 246, 14, 97, 53, 87, 185, 134, 193, 29, 158,
  225, 248, 152, 17, 105, 217, 142, 148, 155, 30, 135, 233, 206, 85, 40, 223,
  140, 161, 137, 13, 191, 230, 66, 104, 65, 153, 45, 15, 176, 84, 187, 22
};
    
```

Shift Row

The Shift row operation rotates the order of bytes in the i 'th row of the rectangle by i bytes to the left



Mix Column

Mix column is a linear mixing of a word column (four bytes) (a_0, a_1, a_2, a_3) into a word column (b_0, b_1, b_2, b_3) . The mixing is defined by operations in $GF(2^8)$.

It is equivalent to

$$\begin{pmatrix} b_0 \\ b_1 \\ b_2 \\ b_3 \end{pmatrix} = \begin{pmatrix} f(a_0, a_1, a_2, a_3) \\ f(a_1, a_2, a_3, a_0) \\ f(a_2, a_3, a_0, a_1) \\ f(a_3, a_0, a_1, a_2) \end{pmatrix}$$

where

$$f(a_0, a_1, a_2, a_3) = (a_1 \oplus a_2 \oplus a_3) \oplus g((a_0 \oplus a_1) \ll 1)$$

$$g(x) = x \& 100_x ? x \oplus 11B_x : x$$

(\ll , \gg , $\&$, $?$: are shifts, logical AND, and conditional expression as in the C programming language).

Note that the left shift might turn the 9th bit on. The XOR with $11B_x$ in $g()$ ensures that the result always fits in a byte ($g()$ enforces the Galois field $GF(2^8)$ modulo the polynomial denoted by $11B_x$, i.e., $x^8 + x^4 + x^3 + x + 1$).

Mix Column (cont.)

The **inverse** of the Mix Column operation is also linear, and is similar to the form of Mix Column itself, where f is replaced by d where

$$d_1(a_0, a_1, a_2, a_3) = a_0 \oplus a_1 \oplus a_2 \oplus a_3$$

$$d_2(a_0, a_1, a_2, a_3) = g(d_1(a_0, a_1, a_2, a_3) \ll 1) \oplus (a_0 \oplus a_2)$$

$$d_3(a_0, a_1, a_2, a_3) = g(d_2(a_0, a_1, a_2, a_3) \ll 1) \oplus (a_0 \oplus a_1)$$

$$d(a_0, a_1, a_2, a_3) = g(d_3(a_0, a_1, a_2, a_3) \ll 1) \oplus (a_1 \oplus a_2 \oplus a_3)$$

Key Mixing

The key mixing operation XORs the data with a subkey. The subkey has the same size of the blocks.

The subkeys are computed from the key by a key scheduling algorithm, shown later.

Encryption

Encryption is performed by

1. First key mixing (using subkey K_0)
2. Nr rounds consisting of
 - (a) Byte substitution
 - (b) Shift row
 - (c) Mix column (except for the last round)
 - (d) Key mixing (using subkey K_i in round $i \in \{1, \dots, Nr\}$)

Decryption

Decryption applies the inverses of the operations in the reverse order, using the reverse order of the subkeys.

It requires the application of the inverse S boxes, inverse Mix column, and inverse shift rows.

The key mixing is equivalent to itself.

The Number of Rounds

The number of rounds Nr depends on the key size

Key Size	Rounds
128	10
192	12
256	14

The Key Schedule

The key schedule takes an Nk -word key ($32 \cdot Nk$ bits) and computes the $Nr+1$ subkeys of a total size of $4(Nr+1)$ words.

Let W be a word array of size $W[4 \cdot (Nr+1)]$. The first 4 words of W become the subkey K_0 , the next 4 words become K_1 , etc.

The key schedule is as follows

```

W[0..Nk-1] = key[*];
for i := Nk to 4*(Nr+1)-1 do {
  temp = W[i-1]
  if ((i%Nk)==0)
    temp = bytesubstitution(temp<<<8) ^ RCON[i/Nk];
  W[i] = W[i-Nk] ^ temp;
}
    
```

The Key Schedule (cont.)

where

- `bytesubstitution()` apply the S box (byte substitution) of Rijndael on each byte of `temp<<<8`, and
- `RCON` is fixed by

```

word32 RCON[] = {
  01x, 02x, 04x, 08x, 10x, 20x, 40x, 80x,
  1bx, 36x, 6cx, d8x, abx, 4dx, 9ax, 2fx
};
    
```

(these are powers of 2 in $GF(2^8)$).

Efficient Implementations

Efficient implementations of Rijndael combine the S boxes, Shift Rows and Mix Column operations together.

The Mix Column operation can be performed by four table lookups (8-bit indices, 32-bit output), and XORing the results.

The Shift row operation can be eliminated by carefully indexing the indices of the input bytes to the succeeding Mix Column operation.

As the S boxes are also implemented as table lookups, the three operations together can be implemented by 16 table lookups in total (and 12 XORs), using 4 precomputed tables of

Mix Column(Shift Row(S boxes(-)))

Security

There are no complementation properties, nor equivalent keys.

The security of the cipher highly depends on the choice of the S boxes. If the S boxes were affine, the whole cipher would be affine, thus easily breakable. (linear: $f(x) = ax$; affine: $f(x) = ax + b$).

The best theoretic attack breaks up to 8 rounds with over 2^{120} complexity for 128-bit keys, and 2^{204} for 256-bit keys.

No shortcut attacks on the full Rijndael are known.

The Square Attack

Observe the following property:

If one byte is modified in the plaintext, then exactly 4 are modified after one round, and all the 16 are modified after two rounds.

The same property holds in the decryption direction.

Thus, a one-byte difference cannot lead to a one-byte difference after three rounds.

This observation can be used for an attack on 5 rounds of Rijndael. Using a few more tricks, up to 7 rounds can be attacked.

Speed of Widely-Used Ciphers

Encrypt/Decrypt speed is given in units of cycles/byte on Pentium III. Key setup time is given in cycles. (source: NESSIE performance document)

Cipher	Key Size	Block size	Encrypt	Decrypt	Key setup
DES	56	64	59	59	883
Triple-DES	168	64	154	155	2636
Idea	128	64	55	55	450
Kasumi	128	64	73	74	297
Skipjack	80	64	114	120	16
AES - Rijndael	128	128	25	26	504
	192	128	30	31	601
	256	128	34	35	949
Rijndael	192	256	33	34	1785
(not part of AES)	256	256	33	35	2298
RC4 (stream cipher)	128	8 stream generation			2659
		7.3 cycles/byte			