

Lecture 3: Block Ciphers and the Data Encryption Standard

Lecture Notes on “Computer and Network Security”

by Avi Kak (kak@purdue.edu)

January 19, 2012

6:22pm

©2012 Avinash Kak, Purdue University

Goals:

- To introduce the notion of a block cipher in the modern context.
- To talk about the infeasibility of **ideal block ciphers**
- To introduce the notion of the **Feistel Cipher Structure**
- To go over **DES**, the Data Encryption Standard

CONTENTS

	<i>Section Title</i>	<i>Page</i>
3.1	Ideal Block Cipher	3
3.1.1	Size of the Encryption Key for the Ideal Block Cipher	6
3.2	The Feistel Structure for Block Ciphers	7
3.2.1	Mathematical Description of Each Round in the Feistel Structure	9
3.2.2	Decryption in Ciphers Based on the Feistel Structure	11
3.3	DES: The Data Encryption Standard	14
3.3.1	One Round of Processing in DES	16
3.3.2	The S-Box for the Substitution Step in Each Round	20
3.3.3	The Substitution Tables	24
3.3.4	The P-Box Permutation in the Feistel Function	26
3.3.5	The DES Key Schedule: Generating the Round Keys	27
3.3.6	Initial Permutation of the Encryption Key	29
3.3.7	Contraction-Permutation that Generates the 48-Bit Round Key from the 56-Bit Key	30
3.4	What Makes DES a Strong Cipher (to the Extent it is a Strong Cipher)	31
3.5	Homework Problems	33

3.1: Ideal Block Cipher

- In a modern block cipher (but still using a classical encryption method), we replace a block of N bits from the plaintext with a block of N bits from the ciphertext. This general idea is illustrated in Figure 1 for the case of $N = 4$. (In general, though, N is set to 64 or multiples thereof.)
- To understand Figure 1, note that there are 16 different possible 4-bit patterns. We can represent each pattern by an integer between 0 and 15. So the bit pattern 0000 could be represented by the integer 0, the bit pattern 0001 by integer 1, and so on. The bit pattern 1111 would be represented by the integer 15.
- In an ideal block cipher, the relationship between the input blocks and the output block is completely random. But it must be invertible for decryption to work. Therefore, it has to be one-to-one, meaning that each input block is mapped to a unique output block.
- The mapping from the input bit blocks to the output bit blocks can also be construed as a mapping from the *integers* corresponding to the input bit blocks to the *integers* corresponding to the

output bit blocks.

- The encryption key for the ideal block cipher is the codebook itself, meaning the table that shows the relationship between the input blocks and the output blocks.
- Figure 1 depicts an ideal block cipher that uses blocks of size 4. Each block of 4 bits in the plaintext is transformed into a block of 4 ciphertext bits.

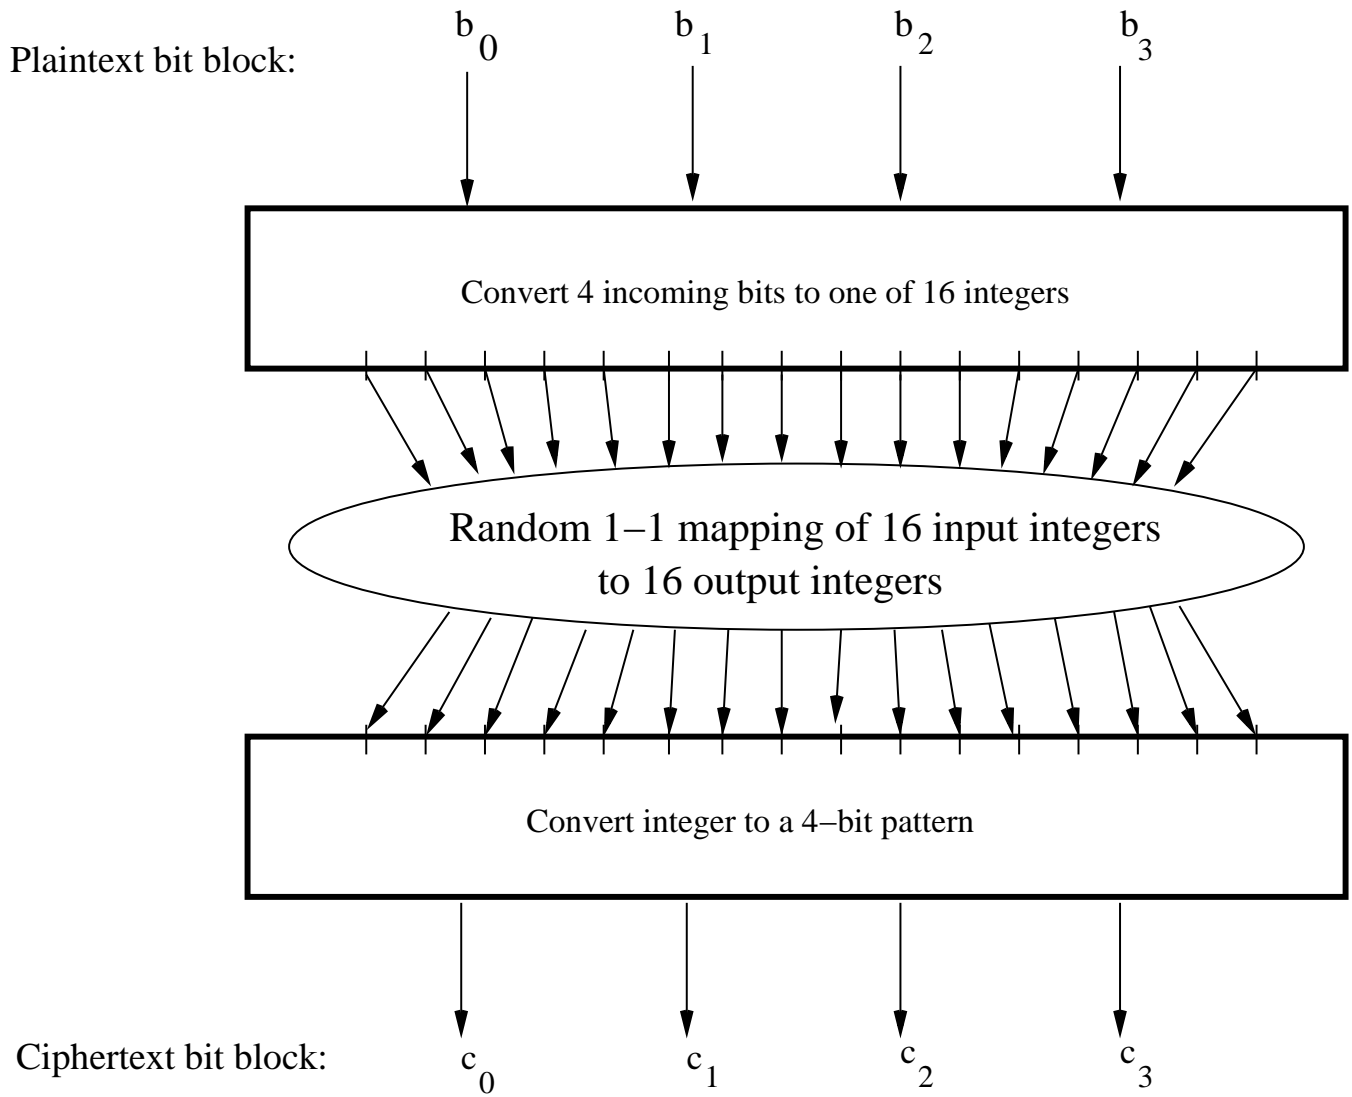


Figure 1: *This figure is from Lecture 3 of “Lecture Notes on Computer and Network Security” by Avi Kak*

3.1.1: The Size of the Encryption Key for the Ideal Block Cipher

- Consider the case of 64-bit block encryption.
- With a 64-bit block, we can think of each possible input block as one of 2^{64} integers and for each such integer we can specify an output 64-bit block. We can construct the codebook by displaying just the output blocks in the order of the integers corresponding to the input blocks. Such a code book will be of size $64 \times 2^{64} \approx 10^{21}$.
- That implies that the encryption key for the ideal block cipher using 64-bit blocks will be of size 10^{21} .
- The size of the encryption key would make the ideal block cipher an impractical idea. **Think of the logistical issues related to the transmission, storage, and processing of such large keys.**

3.2: The Feistel Structure for Block Ciphers

- Named after the IBM cryptographer Horst Feistel and first implemented in the Lucifer cipher by Horst Feistel and Don Coppersmith.
- A cryptographic system based on Feistel structure uses the same basic algorithm for both encryption and decryption.
- As shown in Figure 2, the Feistel structure consists of multiple rounds of processing of the plaintext, with each round consisting of a “substitution” step followed by a permutation step.
- The input block to each round is divided into two halves that we can denote \mathbf{L} and \mathbf{R} for the left half and the right half.
- In each round, the right half of the block, \mathbf{R} , goes through unchanged. But the left half, \mathbf{L} , goes through an operation that depends on \mathbf{R} and the encryption key.
- The permutation step at the end of each round consists of swapping the modified \mathbf{L} and \mathbf{R} . *Therefore, the \mathbf{L} for the next round would be \mathbf{R} of the current round. And \mathbf{R} for the next round be the output \mathbf{L} of the current round.*

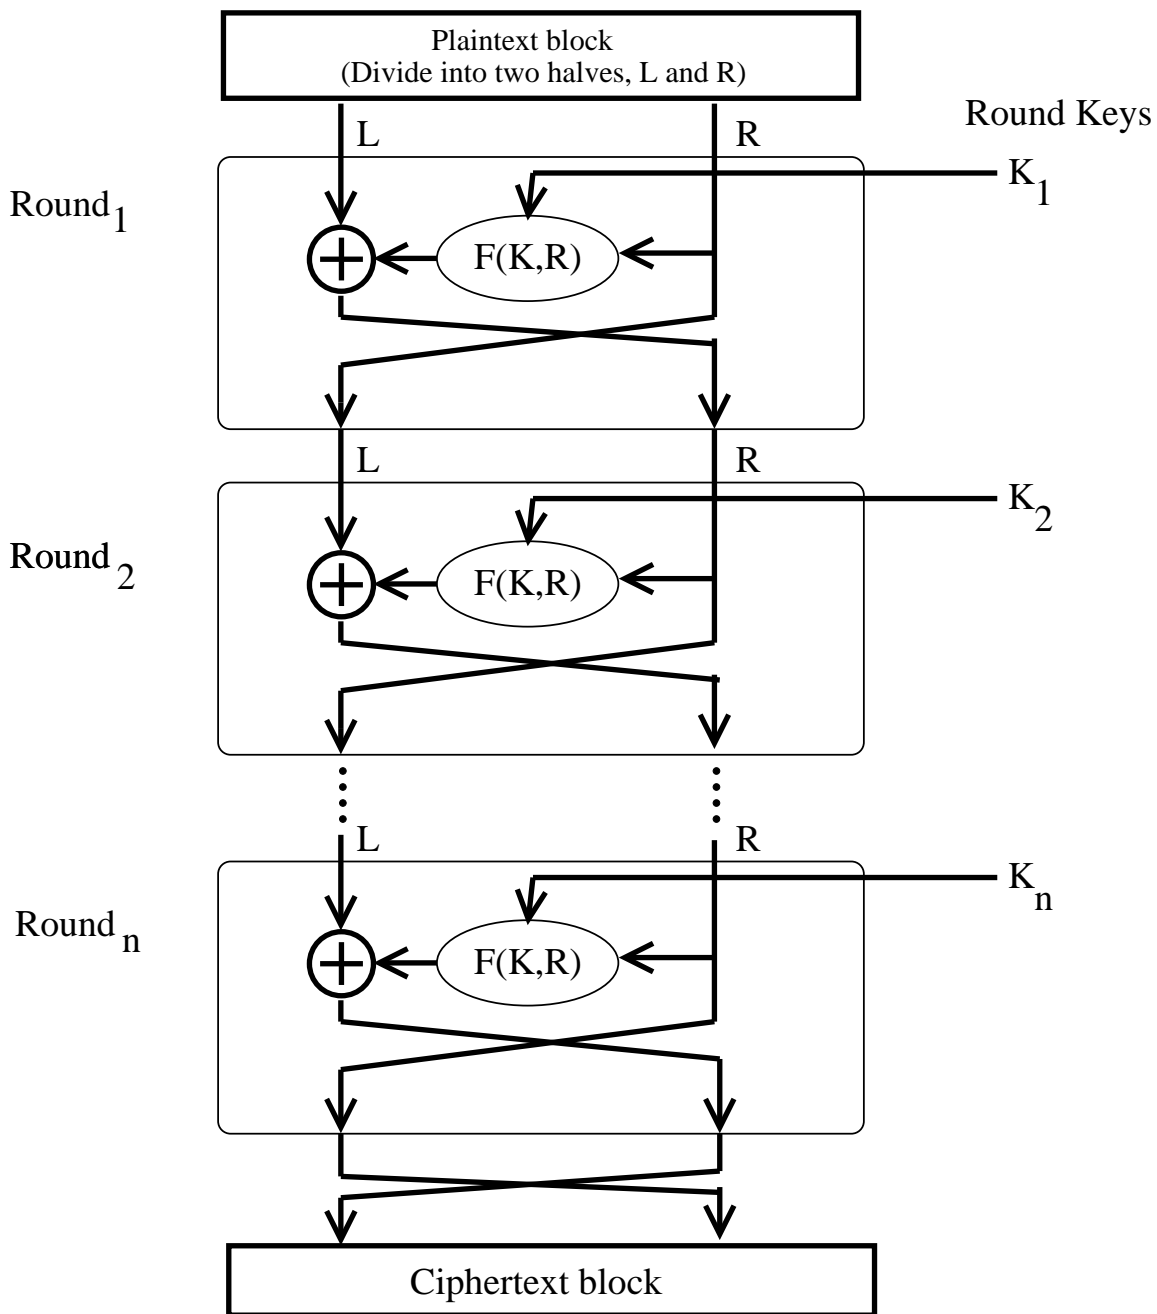


Figure 2: This figure is from Lecture 3 of “Lecture Notes on Computer and Network Security” by Avi Kak

3.2.1: Mathematical Description of Each Round in the Feistel Structure

- Let LE_i and RE_i denote the output half-blocks at the end of the i^{th} round of processing. The letter 'E' denotes encryption.
- In the Feistel structure, the relationship between the output of the i^{th} round and the output of the previous round, that is, the $(i - 1)^{th}$ round, is given by

$$\begin{aligned}LE_i &= RE_{i-1} \\RE_i &= LE_{i-1} \oplus F(RE_{i-1}, K_i)\end{aligned}$$

where \oplus denotes the bitwise EXCLUSIVE OR operation. The symbol F denotes the operation that “scrambles” RE_{i-1} of the previous round with what is shown as the **round key** K_i in Figure 2. The round key K_i is derived from the main encryption key as we will explain later.

- F is referred to as the Feistel function, after Horst Feistel naturally.

- Assuming 16 rounds of processing (which is typical), the output of the last round of processing is given by

$$\begin{aligned}LE_{16} &= RE_{15} \\RE_{16} &= LE_{15} \oplus F(RE_{15}, K_{16})\end{aligned}$$

3.2.2: Decryption in Ciphers Based on the Feistel Structure

- As shown in Figure 3, the decryption algorithm is exactly the same as the encryption algorithm with the only difference that the round keys are used in the reverse order.
- **The output of each round during decryption is the input to the corresponding round during encryption. This property holds true regardless of the choice of the Feistel function F .**
- To prove the above claim, let LD_i and RD_i denote the left half and the right half of the output of the i^{th} round.
- That means that the output of the first decryption round consists of LD_1 and RD_1 . So we can denote the input to the first decryption round by LD_0 and RD_0 . The relationship between the two halves that are input to the first decryption round and what is output by the encryption algorithm is

$$\begin{aligned}LD_0 &= RE_{16} \\RD_0 &= LE_{16}\end{aligned}$$

- We can write the following equations for the output of the first decryption round

$$\begin{aligned}
 LD_1 &= RD_0 \\
 &= LE_{16} \\
 &= RE_{15}
 \end{aligned}$$

$$\begin{aligned}
 RD_1 &= LD_0 \oplus F(RD_0, K_{16}) \\
 &= RE_{16} \oplus F(LE_{16}, K_{16}) \\
 &= [LE_{15} \oplus F(RE_{15}, K_{16})] \oplus F(RE_{15}, K_{16}) \\
 &= LE_{15}
 \end{aligned}$$

This shows that the output of the first round of decryption is the same as the input to the last stage of the encryption round since we have $LD_1 = RE_{15}$ and $RD_1 = LE_{15}$

- The following equalities are used in the above derivation. Assume that A , B , and C are bit arrays.

$$\begin{aligned}
 [A \oplus B] \oplus C &= A \oplus [B \oplus C] \\
 A \oplus A &= 0 \\
 A \oplus 0 &= A
 \end{aligned}$$

- **The above result is independent of the precise nature of F .**

Encryption

Decryption

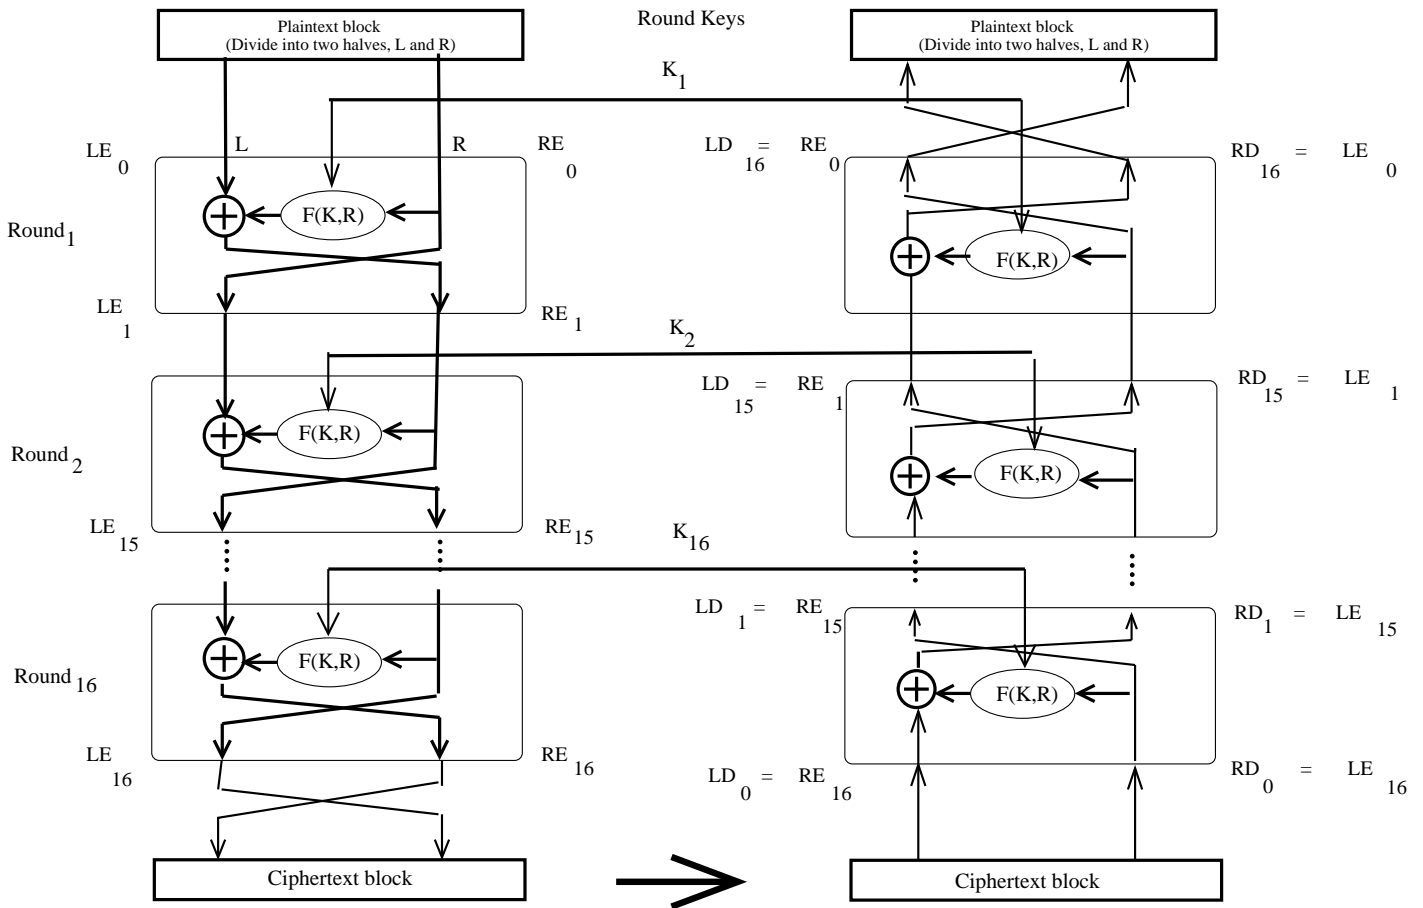


Figure 3: This figure is from Lecture 3 of “Lecture Notes on Computer and Network Security” by Avi Kak

3.3: DES: The Data Encryption Standard

- Adopted by NIST in 1977.
- Based on a cipher (Lucifer) developed earlier by IBM for Lloyd's of London for cash transfer.
- DES uses the Feistel cipher structure with 16 rounds of processing.
- DES uses a 56-bit encryption key. (The key size was apparently dictated by the memory and processing constraints imposed by a single-chip implementation of the algorithm for DES.) The key itself is specified with 8 bytes, but one bit of each byte is used as a parity check.
- **DES encryption was broken in 1999 by Electronics Frontiers Foundation (EFF, www.eff.org).** This resulted in NIST issuing a new directive that year that required organizations to use **Triple DES**, that is, three consecutive applications of **DES**. (That DES was found to be not as strong as originally

believed also prompted NIST to initiate the development of new standards for data encryption. The result is **AES** that we will discuss later.)

- **Triple DES** continues to enjoy wide usage in commercial applications even today. To understand Triple DES, you must first understand the basic **DES** encryption.
- As mentioned, DES uses the Feistel structure with 16 rounds.
- What is specific to DES is the implementation of the F function in the algorithm and how the round keys are derived from the main encryption key.
- The round keys are generated from the main key by a sequence of permutations. Each round key is 48 bits in length.

3.3.1: One Round of Processing in DEA

- The algorithmic implementation of DES is known as **DEA** for **Data Encryption Algorithm**.
- Figure 4 shows a single round of processing in DEA. The dotted rectangle constitutes the F function.
- The 32-bit right half of the 64-bit input data block is expanded by into a 48-bit block. This is referred to as the **expansion permutation** step, or the **E-step**.
- The above-mentioned E-step entails the following:
 - first divide the 32-bit block into eight 4-bit words
 - attach an additional bit on the left to each 4-bit word that is the last bit of the previous 4-bit word
 - attach an additional bit to the right of each 4-bit word that is the beginning bit of the next 4-bit word.

Note that what gets prefixed to the first 4-bit block is the last bit of the last 4-bit block. By the same token, what gets appended

to the last 4-bit block is the first bit of the first 4-bit block. The reason for why we expand each 4-bit block into a 6-bit block in the manner explained will become clear shortly.

- The 56-bit key is divided into two halves, each half shifted separately, and the combined 56-bit key **permuted/contracted** to yield a 48-bit **round** key. How this is done will be explained later.
- The 48 bits of the expanded output produced by the E-step are XORed with the round key. This is referred to as **key mixing**.
- The output produced by the previous step is broken into eight six-bit words. Each six-bit word goes through a substitution step; its replacement is a 4-bit word. The substitution is carried out with an **S-box**, as explained in greater detail in Section 3.3.2. [The name “S-Box” stands for “Substitution Box”.]
- So after all the substitutions, we again end up with a 32-bit word.
- The 32-bits of the previous step then go through a P-box based permutation, as shown in Figure 4.

- What comes out of the P-box is then XORed with the left half of the 64-bit block that we started out with. The output of this XORing operation gives us the right half block for the next round.
- Note that the goal of the substitution step implemented by the **S-box** is to introduce **diffusion** in the generation of the output from the input. *Diffusion means that each plaintext bit must affect as many ciphertext bits as possible.*
- The strategy used for creating the different round keys from the main key is meant to introduce **confusion** into the encryption process. *Confusion in this context means that the relationship between the encryption key and the ciphertext must be as complex as possible.* Another way of describing confusion would be that each bit of the key must affect as many bits as possible of the output ciphertext block.
- Diffusion and confusion are the two cornerstones of block cipher design.

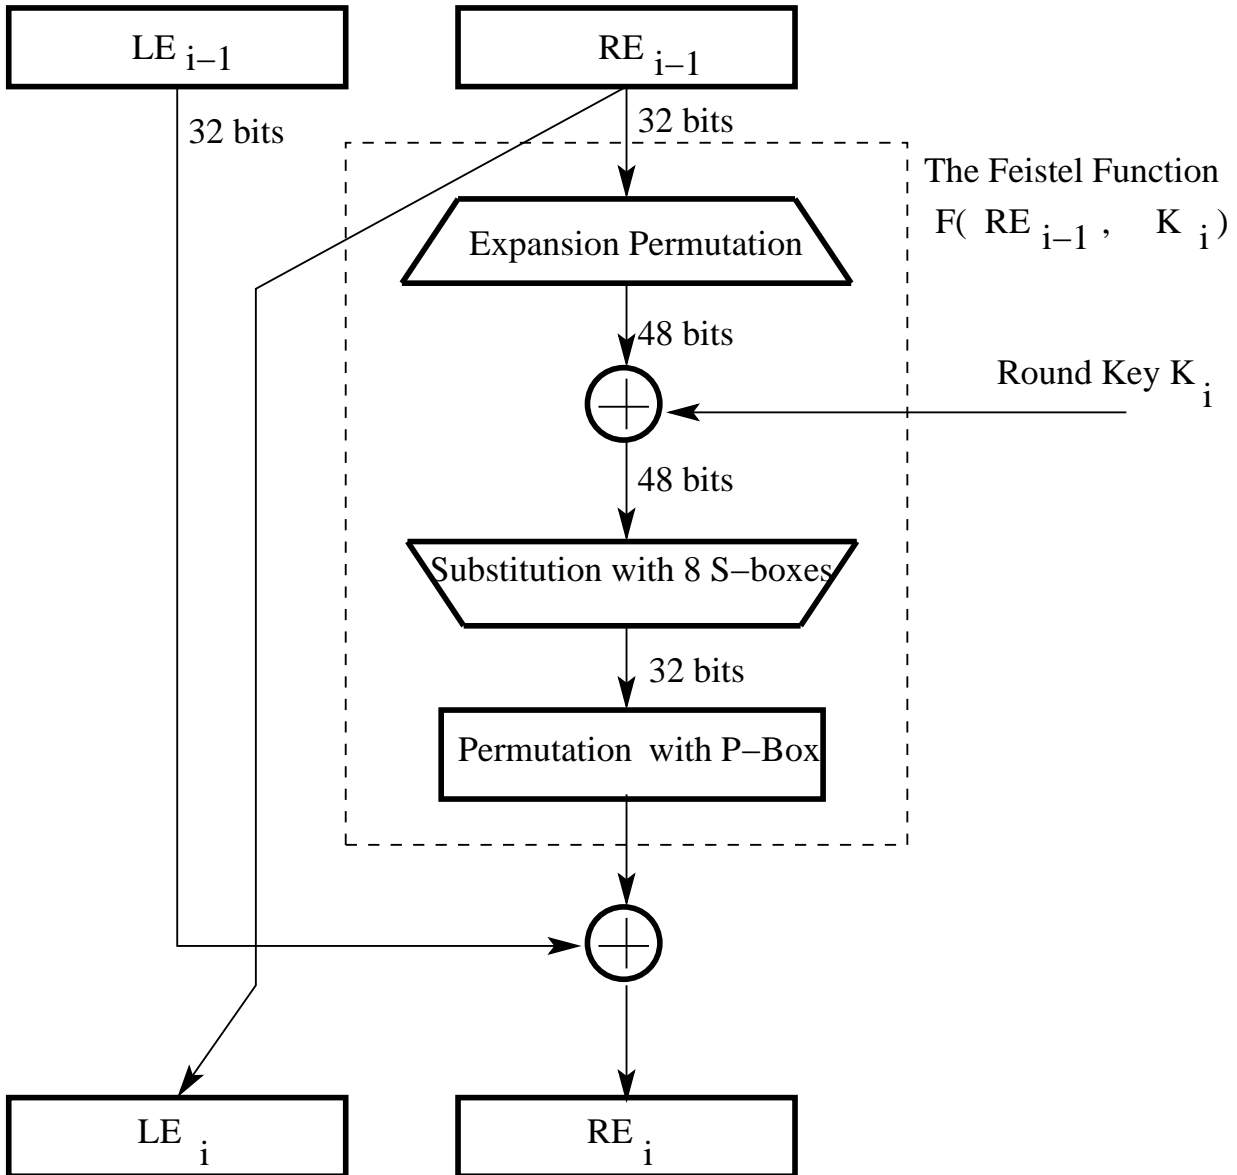


Figure 4: *This figure is from Lecture 3 of “Lecture Notes on Computer and Network Security” by Avi Kak*

3.3.2: The S-Box for the Substitution Step in Each Round

- As shown in Figure 5, the 48-bit input word is divided into eight 6-bit words and each 6-bit word fed into a separate S-box. Each S-box produces a 4-bit output. Therefore, the 8 S-boxes together generate a 32-bit output. As you can see, the overall substitution step takes the 48-bit input back to a 32-bit output.
- Each of the eight S-boxes consists of a 4×16 table lookup for an output 4-bit word. The first and the last bit of the 6-bit input word are decoded into one of 4 rows and the middle 4 bits decoded into one of 16 columns for the table lookup.
- The goal of the substitution carried out by an S-box is to enhance **diffusion**, as mentioned previously. As you will recall from the E-step described in Section 3.3.1, the expansion-permutation step (the E-step) expands a 32-bit block into a 48-bit block by attaching a bit at the beginning and a bit at the end of each 4-bit sub-block, the two bits needed for these attachments belonging to the adjacent blocks.
- Thus, the row lookup for each of the eight S-boxes becomes a

function of the input bits for the previous S-box and the next S-box.

- In the design of the DES, the S-boxes were tuned to enhance the resistance of DES to what is known as the **differential cryptanalysis** attack. [For block ciphers, differential cryptanalysis consists of presenting to the encryption algorithm pairs of plaintext bit patterns **with known differences** between them and examining the differences between the corresponding ciphertext outputs. The outputs are usually recorded at the input to the last round of the cipher. Let's represent one plaintext bit block by $X = [X_1, X_2, \dots, X_n]$ where X_i denotes the i^{th} bit in the block, and let's represent the corresponding output bit block by $Y = [Y_1, Y_2, \dots, Y_n]$. By the difference between two plaintext bit blocks X' and X'' we mean $\Delta X = X' \oplus X''$. The difference between the corresponding outputs Y' and Y'' is given by $\Delta Y = Y' \oplus Y''$. The pair $(\Delta X, \Delta Y)$ is known as a **differential**. In an ideally randomizing block cipher, the probability of ΔY being a particular value for a given ΔX is $1/2^n$ for an n -bit block cipher. What is interesting is that the probabilities of ΔY taking on different values for a given ΔX can be shown to be independent of the encryption key because of the properties of the XOR operator, but these probabilities are strongly dependent on the S-box tables. By feeding into a cipher several pairs of plaintext blocks with known ΔX and observing the corresponding ΔY , it is possible to establish constraints on the round key bits encountered along the different paths in the encryption processing chain. (By constraints I mean the following: Speaking hypothetically for the purpose of illustrating a point and focusing on just one round of DES, suppose we can show that the following condition can be expected to be obeyed with high probability: $\Delta X_i \oplus \Delta Y_i \oplus K_i = 0$ for some bit K_i of the encryption key, then it must be the case that $K_i = \Delta X \oplus \Delta Y$.) Note that differential cryptanalysis is a **chosen plaintext attack**, meaning that the attacker will feed known plaintext bit patterns into the cipher and analyze the corresponding outputs in order to figure out the encryption key. In a theoretical analysis of an attack based on differential cryptanalysis, it was shown by Eli Biham and Adi Shamir in 1990 that the DES's encryption key could be figured out provided one could feed known

2^{47} plaintext blocks into the cipher. For a tutorial by Howard Heys on differential cryptanalysis, see http://www.engr.mun.ca/~howard/PAPERS/ldc_tutorial.pdf.]




Figure 5: *This figure is from Lecture 3 of “Lecture Notes on Computer and Network Security” by Avi Kak*

3.3.3: The Substitution Tables

- Shown on the next page are the eight S-boxes, S_1 through S_8 , each S-box being a 4×16 substitution table that is used to convert 6 incoming bits into 4 outgoing bits.
- As mentioned earlier, each row of a substitution table is indexed by the two outermost bits of a six-bit block and each column by the remaining inner 4 bit.

The 4×16 substitution table for S-box S_1															
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 S_2															
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
S-box S_3															
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 S_4															
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
S-box S_5															
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 S_6															
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
S-box S_7															
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 S_8															
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

3.3.4: The P-Box Permutation in the Feistel Function

The last step in the Feistel function shown in Figure 4 is labeled “Permutation with P-Box”. The permutation table is shown below.

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

This permutation table says that the first output bit will be the 16th bit of the input, the second output bit the 7th bit of the input, and so on, for all of the 32 bits of the output that are obtained from the 32 bits of the input. *Note that bit indexing starts with 1 and not with 0.*

3.3.5: The DES Key Schedule: Generating the Round Keys

- The initial 56-bit key may be represented as 8 bytes, with the last bit of each byte used as a parity bit.
- The relevant 56 bits are subject to a permutation at the beginning before any round keys are generated. This is referred to as Permutation Choice 1 that is shown in Section 3.3.6.
- At the beginning of each round, we divide the 56 relevant key bits into two 28 bit halves and circularly shift to the left each half by one or two bits, depending on the round, as shown in the table on the next page.
- For generating the round key, we join together the two halves and apply a 56 bit to 48 bit contracting permutation (this is referred to as Permutation Choice 2, as shown in Section 3.3.7) to the joined bit pattern. **The resulting 48 bits constitute our round key.**

- The contraction permutation shown in Permutation Choice 2, along with the one-bit or two-bit rotation of the two key halves prior to each round, is meant to ensure that each bit of the original encryption key is used in roughly 14 of the 16 rounds.
- The two halves of the encryption key generated in each round are fed as the two halves going into the next round.
- The table shown below tells us how many positions to use for the left circular shift that is applied to the two key halves at the beginning of each round:

<i>Round Number</i>	<i>Number of left shifts</i>
1	1
2	1
3	2
4	2
5	2
6	2
7	2
8	2
9	1
10	2
11	2
12	2
13	2
14	2
15	2
16	1

3.3.6: Initial Permutation of the Encryption Key

Permutation Choice 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

Note that the bit positions assume that the key bits are addressed 1 through 64 in an 8-byte bit pattern. But note that the last bit of each byte is used as a parity bit. Also note that the permutation shown is not a table, in the sense that the rows and the columns do not carry any special and separate meanings. The permutation order for the bits is given by reading the entries shown from the upper left corner to the lower right corner.

This table tells us that the first bit of the output will be the 57th bit of the input (in a 64 bit representation of the 56-bit encryption key), the second bit of the output the 49th bit of the input, and so on, until finally we have for the 56th bit of the output the 4th bit of the input.

3.3.7: Contraction-Permutation that Generates the 48-Bit Round Key from the 56-Bit Key

Permutation Choice 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

As with the permutation shown on the previous page, note that the bit positions assume that the key bits are addressed 1 through 64 in an 8-byte bit pattern. But note that the last bit of each byte is used as a parity bit. Also note that the permutation shown is not a table, in the sense that the rows and the columns do not carry any special and separate meanings. The permutation order for the bits is given by reading the entries shown from the upper left corner to the lower right corner.

3.4: What Makes DES a Strong Cipher (to the Extent it is a Strong Cipher)

- The substitution step is very effective as far as **diffusion** is concerned. It has been shown that if you change just one bit of the 64-bit input data block, on the average that alters 34 bits of the ciphertext block.
- The manner in which the round keys are generated from the encryption key is also very effective as far as **confusion** is concerned. It has been shown that if you change just one bit of the encryption key, on the average that changes 35 bits of the ciphertext.
- Both effects mentioned above are referred to as the **avalanche effect**.
- And, of course, the 56-bit encryption key means a key space of size $2^{56} \approx 7.2 \times 10^{16}$.

- Assuming that, on the average, you'd need to try half the keys in a brute-force attack, a machine able to process 1000 keys per microsecond would need roughly 13 months to break the code. However, a parallel-processing machine trying 1 million keys simultaneously would need only about 10 hours. **(EFF took three days on a specially architected machine to break the code.)**
- The official document that presents the DES standard can be found at:

<http://www.itl.nist.gov/fipspubs/fip46-2.htm>

3.5: HOMEWORK PROBLEMS

1. In a block cipher, we replace N bits from the plaintext with N bits of ciphertext. What defines an ideal block cipher?
2. Whereas it is true that the relationship between the input and the output is completely random for an ideal block cipher, it must nevertheless be invertible for decryption to work. That implies that the mapping between the input blocks and the output blocks must be one-to-one. If we had to express this mapping in the form of a table lookup, what will be the size of the table?
3. What would be the encryption key for an ideal block cipher?
4. What makes ideal block ciphers impractical?
5. What do we mean by a “Feistel Structure for Block Ciphers”?
6. Are there any constraints on the Feistel function F in a Feistel structure?
7. Explain the concepts of diffusion and confusion as used in DES.
8. If we have all the freedom in the world for choosing the Feistel function F , how should we specify it?

9. How does the permutation/expansion step in DES enhance diffusion? This is the step in which we expand by permutation and repetition the 32-bit half-block into a 48-bit half-block
10. DES encryption was broken in 1999. Why do you think that happened?
11. Since DES was cracked, does that make this an unimportant cipher?
12. **Programming Assignment 1:**

Write a Perl or Python script that implements the full DES. Use the S-boxes that are specified for the DES standard (See Section 3.3.3). Make sure you implement all of the key generation steps outlined in Section 3.3.5. For the encryption key, your script should prompt the user for a keyboard entry that consists of at least 8 printable ASCII characters. (You may choose to either use the first seven or the last seven bits of each character byte for the 56-bit key you need for DES.)

What makes this homework not as difficult as you think is that once you write the code that carries out one round of processing, you basically use the same code in a loop for the whole encryption chain and the decryption chain. Obviously, you will have to reverse the order in which the round keys are used for the decryption chain.

Although you are free to write your own code from scratch, here are some recommendations: If using Python, you might want to

start with the Python BitVector class. To help you get started with the Python implementation, please see the `hw2_starter.py` file. If using Perl, use the `Bit::Vector` module from www.cpan.org. It is a popular Perl module for manipulating bit arrays. It is also well documented. To help you get started with the Perl implementation, please see the `hw2_starter.pl` file.

13. Programming Assignment 2:

Now modify the implementation you created for the previous homework by filling the 4×16 tables for the S-boxes with randomly generated integers. Obviously, each randomly generated entry will have to be between 0 and 15, both ends inclusive. Calculate the avalanche effect for this implementation of DES and compare it with the same effect for your previous implementation. (See Section 3.3.1 for the avalanche effect.)