

Camellia: A 128-Bit Block Cipher Suitable for Multiple Platforms – Design and Analysis

Kazumaro Aoki¹, Tetsuya Ichikawa², Masayuki Kanda¹, Mitsuru Matsui²,
Shiho Moriai¹, Junko Nakajima², and Toshio Tokita²

¹ Nippon Telegraph and Telephone Corporation,
1-1 Hikarinooka, Yokosuka, Kanagawa, 239-0847 Japan
{maro,kanda,shiho}@isl.ntt.co.jp

² Mitsubishi Electric Corporation,
5-1-1 Ofuna, Kamakura, Kanagawa, 247-8501 Japan
{ichikawa,matsui,june15,tokita}@iss.isl.melco.co.jp

Abstract. We present a new 128-bit block cipher called *Camellia*. *Camellia* supports 128-bit block size and 128-, 192-, and 256-bit keys, i.e., the same interface specifications as the Advanced Encryption Standard (AES). Efficiency on both software and hardware platforms is a remarkable characteristic of *Camellia* in addition to its high level of security. It is confirmed that *Camellia* provides strong security against differential and linear cryptanalyses. Compared to the AES finalists, i.e., MARS, RC6, Rijndael, Serpent, and Twofish, *Camellia* offers at least comparable encryption speed in software and hardware. An optimized implementation of *Camellia* in assembly language can encrypt on a Pentium III (800MHz) at the rate of more than 276 Mbits per second, which is much faster than the speed of an optimized DES implementation. In addition, a distinguishing feature is its small hardware design. The hardware design, which includes encryption and decryption and key schedule, occupies approximately 11K gates, which is the smallest among all existing 128-bit block ciphers as far as we know.

1 Introduction

This paper presents a 128-bit block cipher called *Camellia*, which was jointly developed by NTT and Mitsubishi Electric Corporation. *Camellia* supports 128-bit block size and 128-, 192-, and 256-bit key lengths, and so offers the same interface specifications as the Advanced Encryption Standard (AES). The design goals of *Camellia* are as follows.

High Level of Security. The recent advances in cryptanalytic techniques are remarkable. A quantitative evaluation of security against powerful cryptanalytic techniques such as differential cryptanalysis [4] and linear cryptanalysis [18] is considered to be essential in designing any new block cipher. We evaluated the security of *Camellia* by utilizing state-of-art cryptanalytic techniques. We have confirmed that *Camellia* has no differential and linear characteristics that hold

with probability more than 2^{-128} . Moreover, Camellia was designed to offer security against other advanced cryptanalytic attacks including higher order differential attacks [13,10], interpolation attacks [10,2], related-key attacks [5,15], truncated differential attacks [13,23], boomerang attacks [26], and slide attacks [6,7].

Efficiency on Multiple Platforms. As cryptographic systems are needed in various applications, encryption algorithms that can be implemented efficiently on a wide range of platforms are desirable, however, few 128-bit block ciphers are suitable for both software and hardware implementation. Camellia was designed to offer excellent efficiency in hardware and software implementations, including gate count for hardware design, memory requirements in smart card implementations, as well as performance on multiple platforms.

Camellia consists of only 8-by-8-bit substitution tables (*s*-boxes) and logical operations that can be efficiently implemented on a wide variety of platforms. Therefore, it can be implemented efficiently in software, including the 8-bit processors used in low-end smart cards, 32-bit processors widely used in PCs, and 64-bit processors. Camellia doesn't use 32-bit integer additions and multiplications, which are extensively used in some software-oriented 128-bit block ciphers. Such operations perform well on platforms providing a high degree of support, e.g., Pentium II/III or Athlon, but not as well on others. These operations can cause a longer critical path and larger hardware implementation requirements.

The *s*-boxes of Camellia are designed to minimize hardware size. The four *s*-boxes are affine equivalent to the inversion function in the finite field $\text{GF}(2^8)$. Moreover, we reduced the inversion function in $\text{GF}(2^8)$ to a few $\text{GF}(2^4)$ arithmetic operations. It enabled us to implement the *s*-boxes by fewer gate counts.

The key schedule is simple and shares part of its procedure with encryption. It supports on-the-key subkey generation and subkeys are computable in any order. The memory requirement for generating subkeys is quite small; an efficient implementation requires about 32-byte RAM for 128-bit keys and about 64-byte RAM for 192- and 256-bit keys.

Outline of the Paper. This paper is organized as follows: Sect. 2 describes the notations and high-level structure of Camellia. Section 3 defines each components of the cipher. Section 4 describes the rationale behind Camellia's design. In Sect. 5 we evaluate Camellia's strength against known attacks. Section 6 contains the performance of Camellia. We conclude in Sect. 7.

2 Structure of Camellia

Camellia uses an 18-round Feistel structure for 128-bit keys, and a 24-round Feistel structure for 192- and 256-bit keys, with additional input/output whitenings and logical functions called the *FL*-function and *FL*⁻¹-function inserted every 6 rounds. Figure 1 shows an overview of encryption using 128-bit keys. An element with the suffix _(*n*) shows that the element is *n*-bit long.

The key schedule generates 64-bit subkeys kw_t ($t = 1, 2, 3, 4$) for input/output whitenings, k_u ($u = 1, 2, \dots, r$) for round functions and kl_v ($v = 1, 2, \dots, r/3-2$)

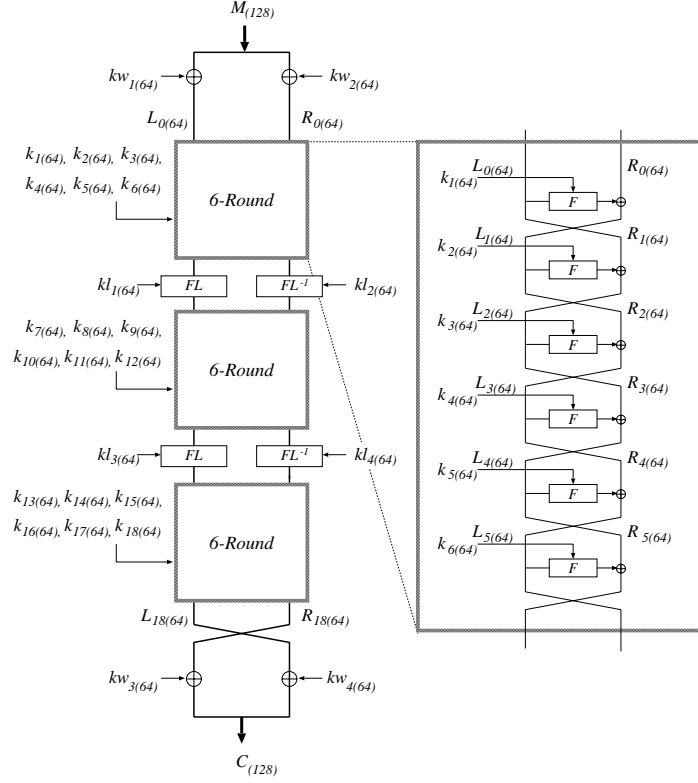


Fig. 1. Encryption procedure of Camellia for 128-bit keys

for FL - and FL^{-1} -functions from the secret key K , where r is the number of rounds.

2.1 Notations

- X_L, X_R : the left-half and the right-half data of X , respectively.
- \oplus, \cap, \cup : bitwise exclusive-OR (XOR), AND and OR operation, respectively.
- \parallel : concatenation of two operands.
- \ggg_n, \lll_n : rotation to the right and the left by n bits, respectively.
- $0x$: hexadecimal representation.

2.2 Encryption for 128-Bit Keys

First a 128-bit plaintext M is XORed with $kw_1 \parallel kw_2$ and separated into two 64-bit data L_0 and R_0 , i.e., $M \oplus (kw_1 \parallel kw_2) = L_0 \parallel R_0$. Then, the following operations are performed from $r = 1$ to 18, except for $r = 6$ and 12;

$$L_r = R_{r-1} \oplus F(L_{r-1}, k_r), R_r = L_{r-1}.$$

For $r = 6$ and 12 , the following is carried out;

$$\begin{aligned} L'_r &= R_{r-1} \oplus F(L_{r-1}, k_r), & R'_r &= L_{r-1}, \\ L_r &= FL(L'_r, kl_{r/3-1}), & R_r &= FL^{-1}(R'_r, kl_{r/3}). \end{aligned}$$

Lastly, R_{18} and L_{18} are concatenated and XORed with $kw_3||kw_4$. The resultant value is the 128-bit ciphertext, i.e., $C = (R_{18}||L_{18}) \oplus (kw_3||kw_4)$.

2.3 Encryption for 192- and 256-Bit Keys

Similarly to the encryption for 128-bit keys, first a 128-bit plaintext M is XORed with $kw_1||kw_2$ and separated into two 64-bit data L_0 and R_0 , i.e., $M \oplus (kw_1||kw_2) = L_0||R_0$. Then, the following operations are performed from $r = 1$ to 24 , except for $r = 6, 12$, and 18 ;

$$L_r = R_{r-1} \oplus F(L_{r-1}, k_r), \quad R_r = L_{r-1}.$$

For $r = 6, 12$, and 18 , the following are performed;

$$\begin{aligned} L'_r &= R_{r-1} \oplus F(L_{r-1}, k_r), & R'_r &= L_{r-1}, \\ L_r &= FL(L'_r, kl_{r/3-1}), & R_r &= FL^{-1}(R'_r, kl_{r/3}). \end{aligned}$$

Lastly, R_{24} and L_{24} are concatenated and XORed with $kw_3||kw_4$. The resultant value is the 128-bit ciphertext, i.e., $C = (R_{24}||L_{24}) \oplus (kw_3||kw_4)$.

2.4 Decryption

The decryption procedure of Camellia can be done in the same way as the encryption procedure by reversing the order of the subkeys, which is one of merits of Feistel networks. In Camellia, FL/FL^{-1} -function layers are inserted every 6 rounds, but this property is still preserved.

2.5 Key Schedule

Figure 2 shows the key schedule of Camellia. Two 128-bit variables K_L and K_R are defined as follows. For 128-bit keys, the 128-bit key K is used as K_L and K_R is 0. For 192-bit keys, the left 128-bit of the key K is used as K_L , and concatenation of the right 64-bit of K and the complement of the right 64-bit of K is used as K_R . For 256-bit keys, the left 128-bit of the key K is used as K_L and the right 128-bit of K is used as K_R .

Two 128-bit variables K_A and K_B are generated from K_L and K_R as shown in Fig. 2. Note that K_B is used only if the length of the secret key is 192 or 256 bits. The 64-bit constants Σ_i ($i = 1, 2, \dots, 6$) are used as “keys” in the Feistel network. They are defined as continuous values from the second hexadecimal place to the seventeenth hexadecimal place of the hexadecimal representation of the square root of the i -th prime. These constant values are shown in Table 1.

The 64-bit subkeys kw_t , k_u , and kl_v are generated from K_L , K_R , K_A , and K_B . The subkeys are generated by rotating K_L , K_R , K_A , and K_B and taking the left- or right-half of them. Details are shown in Table 2.

Table 1. The key schedule constants

Σ_1	0xA09E667F3BCC908B	Σ_4	0x54FF53A5F1D36F1C
Σ_2	0xB67AE8584CAA73B2	Σ_5	0x10E527FADE682D1D
Σ_3	0xC6EF372FE94F82BE	Σ_6	0xB05688C2B3E6C1FD

Table 2. Subkeys for 128-bit keys and 192/256-bit keys

128-bit keys	subkey	value	192/256-bit keys	subkey	value
Prewhitening	kw_1	$(K_L \lll 0)_L$	Prewhitening	kw_1	$(K_L \lll 0)_L$
	kw_2	$(K_L \lll 0)_R$		kw_2	$(K_L \lll 0)_R$
F (Round1)	k_1	$(K_A \lll 0)_L$	F (Round1)	k_1	$(K_B \lll 0)_L$
F (Round2)	k_2	$(K_A \lll 0)_R$	F (Round2)	k_2	$(K_B \lll 0)_R$
F (Round3)	k_3	$(K_L \lll 15)_L$	F (Round3)	k_3	$(K_R \lll 15)_L$
F (Round4)	k_4	$(K_L \lll 15)_R$	F (Round4)	k_4	$(K_R \lll 15)_R$
F (Round5)	k_5	$(K_A \lll 15)_L$	F (Round5)	k_5	$(K_A \lll 15)_L$
F (Round6)	k_6	$(K_A \lll 15)_R$	F (Round6)	k_6	$(K_A \lll 15)_R$
FL	kl_1	$(K_A \lll 30)_L$	FL	kl_1	$(K_R \lll 30)_L$
FL^{-1}	kl_2	$(K_A \lll 30)_R$	FL^{-1}	kl_2	$(K_R \lll 30)_R$
F (Round7)	k_7	$(K_L \lll 45)_L$	F (Round7)	k_7	$(K_B \lll 30)_L$
F (Round8)	k_8	$(K_L \lll 45)_R$	F (Round8)	k_8	$(K_B \lll 30)_R$
F (Round9)	k_9	$(K_A \lll 45)_L$	F (Round9)	k_9	$(K_L \lll 45)_L$
F (Round10)	k_{10}	$(K_L \lll 60)_R$	F (Round10)	k_{10}	$(K_L \lll 45)_R$
F (Round11)	k_{11}	$(K_A \lll 60)_L$	F (Round11)	k_{11}	$(K_A \lll 45)_L$
F (Round12)	k_{12}	$(K_A \lll 60)_R$	F (Round12)	k_{12}	$(K_A \lll 45)_R$
FL	kl_3	$(K_L \lll 77)_L$	FL	kl_3	$(K_L \lll 60)_L$
FL^{-1}	kl_4	$(K_L \lll 77)_R$	FL^{-1}	kl_4	$(K_L \lll 60)_R$
F (Round13)	k_{13}	$(K_L \lll 94)_L$	F (Round13)	k_{13}	$(K_R \lll 60)_L$
F (Round14)	k_{14}	$(K_L \lll 94)_R$	F (Round14)	k_{14}	$(K_R \lll 60)_R$
F (Round15)	k_{15}	$(K_A \lll 94)_L$	F (Round15)	k_{15}	$(K_B \lll 60)_L$
F (Round16)	k_{16}	$(K_A \lll 94)_R$	F (Round16)	k_{16}	$(K_B \lll 60)_R$
F (Round17)	k_{17}	$(K_L \lll 111)_L$	F (Round17)	k_{17}	$(K_L \lll 77)_L$
F (Round18)	k_{18}	$(K_L \lll 111)_R$	F (Round18)	k_{18}	$(K_L \lll 77)_R$
Postwhitening	kw_3	$(K_A \lll 111)_L$	FL	kl_5	$(K_A \lll 77)_L$
	kw_4	$(K_A \lll 111)_R$	FL^{-1}	kl_6	$(K_A \lll 77)_R$
			F (Round19)	k_{19}	$(K_R \lll 94)_L$
			F (Round20)	k_{20}	$(K_R \lll 94)_R$
			F (Round21)	k_{21}	$(K_A \lll 94)_L$
			F (Round22)	k_{22}	$(K_A \lll 94)_R$
			F (Round23)	k_{23}	$(K_L \lll 111)_L$
			F (Round24)	k_{24}	$(K_L \lll 111)_R$
			Postwhitening	kw_3	$(K_B \lll 111)_L$
				kw_4	$(K_B \lll 111)_R$




Fig. 2. Key schedule

3 Components of Camellia

3.1 F -Function

The F -function is shown in Fig. 3. The F -function uses the SPN (Substitution-Permutation Network) structure. The S -function is the non-linear layer and the P -function is the linear layer.

3.2 S -Function, s -Boxes

The S -function consists of eight s -boxes, and four different s -boxes, s_1 , s_2 , s_3 , and s_4 are used. All of them are affine equivalent to the inversion function in $\text{GF}(2^8)$. The data of s_2 , s_3 , and s_4 can be generated from the s_1 table. The tables are shown in [1].

$$\begin{aligned}
 s_1 &: \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, x \mapsto \mathbf{h}(\mathbf{g}(\mathbf{f}(0\mathbf{x}c5 \oplus x))) \oplus 0\mathbf{x}6\mathbf{e} \\
 s_2 &: \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, x \mapsto s_1(x) \lll 1 \\
 s_3 &: \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, x \mapsto s_1(x) \ggg 1 \\
 s_4 &: \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, x \mapsto s_1(x \lll 1)
 \end{aligned}$$

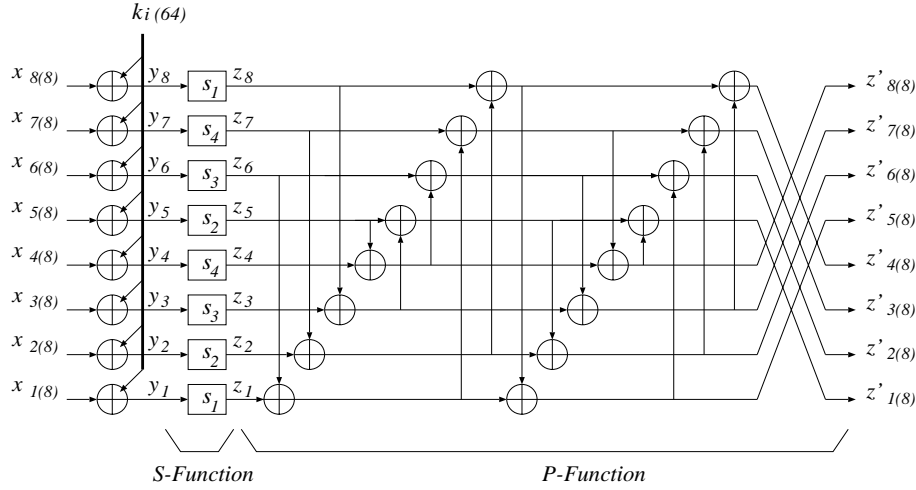


Fig. 3. F-function

where functions \mathbf{f} and \mathbf{h} are affine functions and function \mathbf{g} is the inversion function in $\text{GF}(2^8)$ as given below.

$$\mathbf{f} : \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, (a_1, a_2, \dots, a_8) \mapsto (b_1, b_2, \dots, b_8),$$

where

$$\begin{aligned} b_1 &= a_6 \oplus a_2, & b_2 &= a_7 \oplus a_1, & b_3 &= a_8 \oplus a_5 \oplus a_3, & b_4 &= a_8 \oplus a_3, \\ b_5 &= a_7 \oplus a_4, & b_6 &= a_5 \oplus a_2, & b_7 &= a_8 \oplus a_1, & b_8 &= a_6 \oplus a_4. \end{aligned}$$

$$\mathbf{h} : \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, (a_1, a_2, \dots, a_8) \mapsto (b_1, b_2, \dots, b_8),$$

where



$$\begin{aligned} b_1 &= a_5 \oplus a_6 \oplus a_2, & b_2 &= a_6 \oplus a_2, & b_3 &= a_7 \oplus a_4, & b_4 &= a_8 \oplus a_2, \\ b_5 &= a_7 \oplus a_3, & b_6 &= a_8 \oplus a_1, & b_7 &= a_5 \oplus a_1, & b_8 &= a_6 \oplus a_3. \end{aligned}$$

$$\mathbf{g} : \text{GF}(2)^8 \rightarrow \text{GF}(2)^8, (a_1, a_2, \dots, a_8) \mapsto (b_1, b_2, \dots, b_8),$$

where

$$\begin{aligned} & (b_8 + b_7\alpha + b_6\alpha^2 + b_5\alpha^3) + (b_4 + b_3\alpha + b_2\alpha^2 + b_1\alpha^3)\beta \\ &= ((a_8 + a_7\alpha + a_6\alpha^2 + a_5\alpha^3) + (a_4 + a_3\alpha + a_2\alpha^2 + a_1\alpha^3)\beta)^{-1}. \end{aligned}$$

This inversion is performed in $\text{GF}(2^8)$ assuming $0^{-1} = 0$, where β is an element in $\text{GF}(2^8)$ that satisfies $\beta^8 + \beta^6 + \beta^5 + \beta^3 + 1 = 0$, and $\alpha = \beta^{238} = \beta^6 + \beta^5 + \beta^3 + \beta^2$ is an element in $\text{GF}(2^4)$ that satisfies $\alpha^4 + \alpha + 1 = 0$.

Fig. 4. FL -functionFig. 5. FL^{-1} -function

3.3 P -Function

The P -function is defined as follows:

$$P : (\text{GF}(2)^8)^8 \rightarrow (\text{GF}(2)^8)^8, (z_1, z_2, \dots, z_8) \mapsto (z'_1, z'_2, \dots, z'_8),$$

where

$$\begin{aligned} z'_1 &= z_1 \oplus z_3 \oplus z_4 \oplus z_6 \oplus z_7 \oplus z_8, & z'_2 &= z_1 \oplus z_2 \oplus z_4 \oplus z_5 \oplus z_7 \oplus z_8, \\ z'_3 &= z_1 \oplus z_2 \oplus z_3 \oplus z_5 \oplus z_6 \oplus z_8, & z'_4 &= z_2 \oplus z_3 \oplus z_4 \oplus z_5 \oplus z_6 \oplus z_7, \\ z'_5 &= z_1 \oplus z_2 \oplus z_6 \oplus z_7 \oplus z_8, & z'_6 &= z_2 \oplus z_3 \oplus z_5 \oplus z_7 \oplus z_8, \\ z'_7 &= z_3 \oplus z_4 \oplus z_5 \oplus z_6 \oplus z_8, & z'_8 &= z_1 \oplus z_4 \oplus z_5 \oplus z_6 \oplus z_7. \end{aligned}$$

3.4 FL -Function and FL^{-1} -Function

The FL -function is shown in Fig. 4, and is defined as follows.

$$FL : \text{GF}(2)^{64} \times \text{GF}(2)^{64} \rightarrow \text{GF}(2)^{64}, (X_L || X_R, kl_L || kl_R) \mapsto Y_L || Y_R,$$

where

$$Y_R = ((X_L \cap kl_L) \lll_1) \oplus X_R, \quad Y_L = (Y_R \cup kl_R) \oplus X_L.$$

The FL^{-1} -function is shown in Fig. 5. The following equation holds.

$$FL^{-1}(FL(x, k), k) = x.$$

4 Design Rationale

4.1 F -Function

The design strategy of the F -function of Camellia follows that of the F -function of E2 [14]. The main difference between E2 and Camellia is the adoption of

the 1-round (conservative) SPN, not the 2-round SPN, i.e., S-P-S. When the 1-round SPN is used as the round function in a Feistel cipher, the theoretical evaluation of the upper bound of differential and linear characteristic probability becomes more complicated, but the speed under the same level of “real” security is expected to be improved. See Sect. 6 for detailed discussions on security.

4.2 *P*-Function

The design rationale of the *P*-function is similar to that of the *P*-function of E2 [16]. That is, for computational efficiency, it should be represented using only bitwise XORs and for security against differential and linear cryptanalyses, its branch number should be optimal. From among the linear transformations that satisfy these conditions, we chose one considering highly efficient implementation on 32-processors [3] and high-end smart cards, as well as 8-bit processors.

4.3 *s*-Boxes

As the *s*-boxes we adopted functions affine equivalent to the inversion function in $\text{GF}(2^8)$ for enhanced security and small hardware design.

There is a function affine equivalent to the inversion function in $\text{GF}(2^8)$ that achieves the best known of the maximum differential and linear probabilities, 2^{-6} . We choose this kind of functions as *s*-boxes. Moreover, the high degree of the Boolean polynomial of every output bit of the *s*-boxes makes it difficult to attack Camellia by higher order differential attacks. The two affine functions that are performed at the input and output of the inversion function in $\text{GF}(2^8)$ complicates the expressions of the *s*-boxes in $\text{GF}(2^8)$, which is expected to make interpolation attacks ineffective. Making the four *s*-boxes different slightly improves security against truncated differential cryptanalysis [23].

For small hardware design, the elements in $\text{GF}(2^8)$ can be represented as polynomials with coefficients in the subfield $\text{GF}(2^4)$. In other words, we can implement the *s*-boxes by using a few operations in the subfield $\text{GF}(2^4)$ [22]. Two affine functions at the input and output of the inversion function in $\text{GF}(2^8)$ also play a role in complicating the expressions of the *s*-boxes in $\text{GF}(2^4)$.

4.4 *FL*- and *FL*⁻¹-Functions

FL- and *FL*⁻¹-functions are “inserted” between every 6 rounds of a Feistel network to provide non-regularity across rounds. One of the goals for such a design is to thwart future unknown attacks. It is one of merits of regular Feistel networks that encryption and decryption procedures are the same except for the order of the subkeys. In Camellia, *FL/FL*⁻¹-function layers are inserted every 6 rounds, but this property is still preserved.

The design criteria of these functions are similar to those of the *FL*-function of MISTY [20]. The difference between MISTY and Camellia is the addition of 1-bit rotation. This is expected to make bitwise cryptanalysis harder, but it

has no negative impact on hardware size or speed. The design criteria are that these functions must be linear for any fixed key and that their forms depend on key values. Since these functions are linear as long as the key is fixed, they do not make the average differential and linear probabilities of the cipher higher. Moreover, these functions are fast in both software and hardware since they are constructed by logical operations such as AND, OR, XOR, and rotations.

4.5 Key Schedule

The design criteria of the key schedule are as follows.

1. It should be simple and share part of its procedure with encryption (and decryption).
2. Subkey generation for 128-, 192- and 256-bit keys can be performed by using the same key schedule (circuit). Moreover, the key schedule for 128-bit keys can be performed by using a part of this circuit.
3. Key setup time should be shorter than encryption time. In cases where large amounts of data are processed with a single secret key, the setup time for key scheduling may be unimportant. On the other hand, in applications in which the key is changed frequently, key agility is a factor. One basic component of key agility is key setup time.
4. It should support on-the-fly subkey generation.
5. On-the-fly subkey generation should be computable in the same way in both encryption and decryption. Some ciphers have separate key schedules for encryption and decryption. In other ciphers, e.g., Rijndael or Serpent, subkeys are computable in the forward direction only and require unwinding for decryption.
6. There should be no equivalent keys.
7. There should be no related-key attacks or slide attacks.

Criteria 1 and 2 mainly address small hardware requirements, Criteria 3, 4, and 5 are advantageous in terms of practical applications, and Criteria 6 and 7 are for security.

The memory requirement for generating subkeys is quite small. An efficient implementation of Camellia for 128-bit keys requires 16 bytes (=128 bits) for the original secret key, K_L , and 16 bytes (=128 bits) for the intermediate key, K_A . Thus the required memory is 32 bytes. Similarly, an efficient implementation of Camellia for 192- and 256-bit keys needs only 64 bytes.

5 Security

This section discusses the security of Camellia. Hereafter, we call Camellia without FL - and FL^{-1} -functions Camellia*.

5.1 Differential and Linear Cryptanalysis

The most well-known and powerful approaches to attacking many block ciphers are differential cryptanalysis [4] and linear cryptanalysis [18]. There are several methods of evaluating security against these attacks, where there is a kind of “duality” relation between them [19,8]: in other words, the security against both attacks can be evaluated in similar ways.

It is known that the upper bounds of differential and linear characteristic probabilities can, for several block ciphers, be estimated using the minimum numbers of differential and linear active s -boxes in some consecutive rounds, respectively. Kanda [11] shows the minimum numbers of differential and linear active s -boxes for Feistel ciphers with conservative SPN (S-P) round function.

Definition 1. ([25]) *The branch number \mathcal{B} of linear transformation P is defined by*

$$\mathcal{B} = \min_{x \neq 0} (w_{\text{H}}(x) + w_{\text{H}}(P(x))),$$

where $w_{\text{H}}(x)$ denotes the bitwise Hamming weight of x .

Theorem 1. *The minimum number of differential/linear active s -boxes in any eight consecutive rounds is equal or larger than $2\mathcal{B} + 1$.*

Theorem 2. *Let p_s and q_s be the maximum differential and linear probabilities of all s -boxes, and \mathcal{D} and \mathcal{L} be the minimum numbers of total differential and linear active s -boxes, respectively. Then, the maximum differential and linear characteristic probabilities are bounded by $p_s^{\mathcal{D}}$ and $q_s^{\mathcal{L}}$, respectively.*

In the case of Camellia, the maximum differential and linear probabilities of the s -boxes are $p_s = q_s = 2^{-6}$. The branch number of the linear transformation (P -function) is 5, i.e., $\mathcal{B} = 5$. Letting p, q be the maximum differential and linear characteristic probabilities of Camellia* reduced to 16-round, respectively, we have $p \leq p_s^{2(2\mathcal{B}+1)} = (2^{-6})^{22} = 2^{-132}$ and $q \leq q_s^{2(2\mathcal{B}+1)} = (2^{-6})^{22} = 2^{-132}$ from Theorems 1 and 2. Both probabilities are below the security threshold of 128-bit block ciphers: 2^{-128} . It follows that there is no effective differential characteristic or linear characteristic for Camellia* reduced to more than 15 rounds. Since FL - and FL^{-1} -functions are linear for any fixed key, they do not make the average differential and linear probabilities of the cipher higher. Hence, it is proven that Camellia offers enough security against differential and linear cryptanalyses.

Note that the result above are based on Theorems 1 and 2. Both theorems deal with general cases of Feistel ciphers with SPN round function, so we expect that Camellia is actually more secure than shown by the result above. As supporting evidence, we counted the number of active s -boxes of Camellia and Camellia* with reduced rounds. The counting algorithm is similar to that described in [21] except following three items.

- Prepare the table for the number of active s -boxes instead of transition probability table.

Table 3. Upper bounds of differential characteristic probability of Camellia

# of rounds	1	2	3	4	5	6	7	8	9	10	11	12
Based on			2^{-12}	2^{-30}		2^{-42}		2^{-66}				2^{-96}
Th. 1 and 2			(2)	(5)		(7)		(11)				(16)
Camellia	1	2^{-6}	2^{-12}	2^{-42}	2^{-54}	2^{-66}	2^{-72}	2^{-72}	2^{-78}	2^{-108}	2^{-120}	2^{-132}
	(0)	(1)	(2)	(7)	(9)	(11)	(12)	(12)	(13)	(18)	(20)	(22)
Camellia*	1	2^{-6}	2^{-12}	2^{-36}	2^{-54}	2^{-66}	2^{-78}	2^{-90}	2^{-108}	2^{-126}	2^{-132}	
	(0)	(1)	(2)	(6)	(9)	(11)	(13)	(15)	(18)	(21)	(22)	

Note: The numbers in brackets are the number of active s -boxes.

Table 4. Upper bounds of linear characteristic probability of Camellia

# of rounds	1	2	3	4	5	6	7	8	9	10	11	12
Based on			2^{-12}	2^{-30}		2^{-42}		2^{-66}				2^{-96}
Th. 1 and 2			(2)	(5)		(7)		(11)				(16)
Camellia	1	2^{-6}	2^{-12}	2^{-36}	2^{-54}	2^{-66}	2^{-72}	2^{-72}	2^{-78}	2^{-102}	2^{-120}	2^{-132}
	(0)	(1)	(2)	(6)	(9)	(11)	(12)	(12)	(13)	(17)	(20)	(22)
Camellia*	1	2^{-6}	2^{-12}	2^{-36}	2^{-54}	2^{-66}	2^{-78}	2^{-84}	2^{-108}	2^{-120}	2^{-132}	
	(0)	(1)	(2)	(6)	(9)	(11)	(13)	(14)	(18)	(20)	(22)	

Note: The numbers in brackets are the number of active s -boxes.

- Count the number of active s -boxes instead of computing transition probability.
- FL - and FL^{-1} -functions set all elements to the minimum number of active s -boxes in the table. This means that the algorithm gives consideration to existence of weak subkeys inserted to FL - and FL^{-1} -functions, since there may be some possibility of connecting every later differential and linear characteristic with the previous one with the highest probability, which is equivalent to the minimum number of active s -boxes.

As a result, we confirmed that 12-round Camellia has no differential and linear characteristic with probability higher than 2^{-128} (see Tables 3 and 4).

5.2 Truncated Differential and Linear Cryptanalysis

The attacks using truncated differentials were introduced by Knudsen [13]. He defined them as differentials where only a part of the difference can be predicted. The notion of truncated differentials introduced by him is wide, but with a byte-oriented cipher it is natural to study bitwise differentials as truncated differentials [23].

The maximum differential probability is considered to provide the strict evaluation of security against differential cryptanalysis, but computing its value is impossible in general, since a differential is a set of all differential characteristics with the same input difference and the same output difference for a Markov cipher [17]. On the other hand, a truncated differential can be regarded as a subset

of the differential characteristics which are exploitable in cryptanalysis. For some ciphers, e.g., byte-oriented ciphers, the probability of truncated differential can be computed easily and correctly, and it gives a more strict evaluation than the maximum differential characteristic probability.

A truncated differential cryptanalysis of reduced-round variants of E2 was presented by Matsui and Tokita at FSE'99 [23]. Their analysis was based on the “byte characteristic,” where the values to the difference in a byte are distinguished between non-zero and zero. They found a 7-round byte characteristic, which leads to a possible attack on an 8-round variant of E2 without *IT*-Function (the initial transformation) and *FT*-Function (the final transformation). The best attack of E2 shown in [24] breaks an 8-round variant of E2 with either *IT*-Function or *FT*-Function using 2^{94} chosen plaintexts. In [24] we also show the attack which distinguishes a 7-round variant of E2 with *IT*- and *FT*-Functions from a random permutation using 2^{91} chosen plaintexts.

Camellia is a byte-oriented cipher similar to E2, and it is important to evaluate its security against truncated differential cryptanalysis. We searched for truncated differentials using an algorithm similar to the one described in [23,24]. The main difference of the round function between E2 and Camellia is the adoption of the 1-round SPN not the 2-round SPN, i.e., S-P-S. In the search for truncated differentials of E2, we used about 2^{-8} as the probability of difference cancellation in one byte at the XOR of Feistel network. However, the round function of Camellia doesn't have the second *s*-boxes-layer, and the difference cancellation in plural bytes sometimes occurs with the same probability. Accordingly, we changed the difference cancellation rule at the XOR of Feistel network in the search algorithm. As a result, Camellia with more than 10 rounds is indistinguishable from a random permutation, in both cases with/without *FL*/*FL*⁻¹-function layers.

Next, we introduce a new cryptanalysis called truncated linear cryptanalysis. Due to the duality between differential and linear cryptanalyses, we can evaluate security against truncated linear cryptanalysis by using a similar algorithm to that above. To put it concretely, we can perform the search by replacing the matrix of *P*-function with the transposed matrix. As a result, Camellia* with more than 10 rounds is indistinguishable from a random permutation.

5.3 Boomerang Attack

Boomerang attack [26] requires two differentials. Let the probability of the differentials be p_{Δ} and p_{∇} . An boomerang attack that is superior than exhaustive key search requires

$$p_{\Delta}p_{\nabla} \geq 2^{-64}. \quad (1)$$

Using Table 3, there is no combination that satisfies (1) for Camellia*. The best boomerang probability for Camellia* reduced to 8-round is bounded by 2^{-66} that is obtained by $p_{\Delta} = 2^{-12}$ (3 rounds) and $p_{\nabla} = 2^{-54}$ (5 rounds). Since the attackable rounds is bounded by much shorter than the specification of Camellia, 18 or 24, Camellia seems secure against a boomerang attack.

Table 5. Smallest number of unknown coefficients for 128-, 192-, and 256-bit keys

whitening×1 + round× r ($r < 4$)	1
whitening×1 + round×4	255
More rounds	256

5.4 Higher Order Differential Attack

Higher order differential attack is generally applicable to ciphers that can be represented as Boolean polynomials of low degree. All intermediate bits in the encryption process can be represented as Boolean polynomials, i.e., polynomials $\text{GF}(2)[x_1, x_2, \dots, x_n]$ in the bits of the plaintext: $\{x_1, x_2, \dots, x_n\}$. In the higher order differential attack described in [10, Theorem 1], if the intermediate bits are represented by Boolean polynomials of degree at least d , the $(d+1)$ -th order differential of the Boolean polynomial becomes 0.

For the degrees of Boolean polynomials of the s -boxes of Camellia, the functions affine equivalent to the inversion function in $\text{GF}(2^8)$ are adopted as the s -boxes. We confirmed that the degree of the Boolean polynomial of every output bit of the s -boxes is 7 by finding Boolean polynomial for every output bit of the s -boxes. In Camellia, it is expected that the degree of an intermediate bit in the encryption process increases as the data pass through many s -boxes. For example, the degree becomes $7^3 > 128$ after passing through three s -boxes. Therefore, we expect that higher order differential attacks fail against Camellia with full rounds.

5.5 Interpolation Attack and Linear Sum Attack

The interpolation attack proposed in [10] is typically applicable to attacking ciphers that use simple algebraic functions. Linear sum attack [2] is a generalization of the interpolation attack.

A practical algorithm that evaluates the security against linear sum attack was proposed in [2]. We searched for linear relations between any plaintext byte and any ciphertext byte over $\text{GF}(2^8)$ using the algorithm. Table 5 summarizes the results, and shows that Camellia is secure against linear sum attack including interpolation attack. It also implies that Camellia is secure against SQUARE attack [9] followed by [2, Theorem 3].

5.6 Security of Key Schedule

No Equivalent Keys: Since the set of subkeys generated by the key schedule contain the original secret key, there is no equivalent set of subkeys generated from distinct secret keys. Therefore, we expect that there are no distinct secret keys both of which encrypt each of many plaintexts into the same ciphertext.

Slide Attack: In [6,7] the slide attacks were introduced, based on earlier work in [5,12]. In particular it was shown that iterated ciphers with identical round

functions, that is, equal structures and equal subkeys in the round functions, are susceptible to slide attacks.

In Camellia, FL - and FL^{-1} -functions are “inserted” between every 6 rounds of a Feistel network to provide non-regularity across rounds. Moreover, from the viewpoint of the key schedule, slide attacks seems to be very unlikely to succeed.

Related-Key Attack: We are convinced that the key schedule of Camellia makes related-key attacks [5,15] very difficult. In these attacks, an attacker must be able to get encryptions using several related keys. However, since the subkeys depend on K_A and K_B , which are the results of encryption of a secret key, and if an attacker wants to change the secret key, he can’t get K_A and K_B desired, and vice versa, these subkey relations will be very hard to control and predict.

6 Performance

6.1 Software Implementations

Table 6 summarizes the current software implementations of Camellia. The table shows that Camellia can be efficiently implemented on low-end smart cards, and 32-bit and 64-bit processors. We use the abbreviations M (mega) for 10^6 and m (milli) for 10^{-3} in the table.

6.2 Hardware Performance

We measured the hardware performance of Camellia for 128-bit keys on ASIC (Application Specific Integrated Circuit) and FPGA (Field Programmable Gate Array). Table 7 shows the environment of our hardware design and evaluation. We evaluated hardware performance of the three types: Type 1, Type 2 and Type 3 logic. The hardware design policy of each type is as follows.

Type 1 Fast implementation from the viewpoint of encryption speed

Type 2 Small implementation from the viewpoint of total logic size

Type 3 Small implementation (special case for FPGA)

Tables 8 through 11 summarize the hardware performance of Camellia for 128-bit keys on ASIC and FPGA.

7 Conclusion

We have presented Camellia, the rationale behind its design, its suitability for both software and hardware implementation, and the results of our cryptanalyses. For further information, please refer to the specification of Camellia [1] or full paper, which are available on the Camellia home page: <http://info.is1.ntt.co.jp/camellia/>.

The performances shown in this paper leave room for further optimizations. The latest performance results will be posted on the Camellia home page.

Table 6. Camellia software performance

Processor	Lang.	Key [bits]	Timing [cycles]		Dynamic [bytes]		Code [bytes]		Table [bytes]
			Setup ^a (^b)	Enc. ^c (^d)	Setup ^a	Enc. ^c	Setup ^a	Enc. ^c	
P III ^e	Asm	128	160 (4.4M)	371 (242M)	28	36	1,046	2,150	8,224
		192	222 (3.2M)	494 (181M)	28	36	1,469	3,323	8,240
		256	226 (3.1M)	494 (181M)	28	36	1,485	3,323	8,240
P II ^f	C ^g	128	263 (1.1M)	577 (67M)	44	64	1,600	3,733	4,128
Alpha ^h	Asm	128	118 (5.7M)	339 (252M)	48	48	1,132	3,076	16,528
		192	176 (3.7M)	445 (192M)	48	48	1,668	4,000	16,528
		256	176 (3.7M)	445 (192M)	48	48	1,676	4,000	16,528
		128	158 (4.2M)	326 (262M)	48	48	1,600	2,928	16,512
8051 ⁱ	Asm	128	0 (0)	10217 (10m)	0	32	0	702	288

^a Key schedule may be included.^b Seconds for 8051, and keys/s for other processors.^c Numbers of this column is the same as decryption.^d Seconds for 8051, and b/s for other processors.^e Intel Pentium III (700MHz), 256KB on-die L2 cache, FreeBSD 4.0R, 128MB main memory.^f Intel Pentium II (300MHz), 512KB L2 cache, MS-Windows 95, 160MB main memory.^g ANSI C, Microsoft Visual C++ 6 with the optimization options /G6 /Zp16 /ML /Ox /Ob2.^h Alpha 21264 (667MHz), Compaq Tru64 UNIX 4.0F, 2GB main memory.ⁱ Intel 8051 (12MHz; 1 cycle = 12 oscillator periods) simulator on Unix.**Table 7.** Hardware evaluation environment (ASIC, FPGA)

Language	(ASIC, FPGA) Verilog-HDL
Simulator	(ASIC, FPGA) Verilog-XL
Design library	(ASIC) Mitsubishi Electric 0.35 μ CMOS ASIC library (FPGA) Xilinx XC4000XL series
Login synthesis	(ASIC) Design Compiler version 1998.08 (FPGA) Synplify version 5.3.1 and ALLIANCE version 2.1i

Table 8. Hardware performance (Type 1: [ASIC(0.35 μ CMOS)])

Algorithm name	Area [Gate]			Key setup time [ns]	Critical-path [ns]	Throughput [Mb/s]
	Enc.&Dec. ^a	Key expan. ^b	Total logic ^c			
DES	42,204	12,201	54,405	—	55.11	1161.31
Triple-DES	124,888	23,207	128,147	—	157.09	407.40
MARS	690,654	2,245,096	2,935,754	1740.99	567.49	225.55
RC6	741,641	901,382	1,643,037	2112.26	627.57	203.96
Rijndael	518,508	93,708	612,834	57.39	65.64	1950.03
Serpent	298,533	205,096	503,770	114.07	137.40	931.58
Twofish	200,165	231,682	431,857	16.38	324.80	394.08
Camellia	216,911	55,907	272,819	24.36	109.35	1170.55

^a including output registers^b including subkey registers^c including buffers for fan-out adjustment

Table 9. Hardware performance (Type 2: [ASIC(0.35 μ CMOS)])

Algorithm name	Area [Gate]		Key setup time [ns]	Critical-path [ns]	Throughput [Mb/s]
	Enc.&Dec. ^a	Key sched. ^b Total logic ^c			
Camellia	6,367	4,979 11,350	110.2	27.67	220.28

^a including output registers and data selector^b including subkey registers and a part of key expansion logic^c including buffers for fan-out adjustment**Table 10.** Hardware performance (Type 2: [FPGA(XC4000XL series)])

Algorithm	Total Area [CLBs]	Critical-path [ns]	Throughput [Mb/s]
Camellia	1,296	78.815	77.34

Table 11. Hardware performance (Type 3: [FPGA(XC4000XL series)])

Algorithm	Total Area [CLBs]	Critical-path [ns]	Throughput [Mb/s]
Camellia	874	49.957	122.01

We have analyzed Camellia and found no important weakness. The cipher has a conservative design and any practical attacks against Camellia would require a major breakthrough in the area of cryptanalysis. We think that Camellia is a very strong cipher, which matches the security of the existing best block ciphers.

References

1. Aoki, Ichikawa, Kanda, Matsui, Moriai, Nakajima, Tokita, "Specification of Camellia — a 128-bit Block Cipher," <http://info.is1.ntt.co.jp/camellia/>, 2000.
2. Aoki, "Practical Evaluation of Security against Generalized Interpolation Attack," IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, Vol.E83-A, No.1, pp.33–38, 2000.
3. Aoki, Ueda, "Optimized Software Implementations of E2," IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, Vol.E83-A, No.1, pp.101–105, 2000.
4. Biham, Shamir, "Differential Cryptanalysis of the Data Encryption Standard," Springer-Verlag, 1993.
5. Biham, "New Types of Cryptanalytic Attacks Using Related Keys," Journal of Cryptology, Vol.7, No.4, pp.229–246, Springer-Verlag, 1994.
6. Biryukov, Wagner, "Slide Attacks," Fast Software Encryption, FSE'99, LNCS 1636, pp.245–259, 1999.
7. Biryukov, Wagner, "Advanced Slide Attacks," Advances in Cryptology — EUROCRYPT2000, LNCS 1807, pp.589–606, 2000.
8. Chabaud, Vaudenay, "Links Between Differential and Linear Cryptanalysis," Advances in Cryptology — EUROCRYPT'94, LNCS 950, pp.356–365, 1995.
9. Daemen, Knudsen, Rijmen, "The Block Cipher SQUARE," Fast Software Encryption, FSE'97, LNCS 1267, pp.54–68, 1997.

10. Jakobsen, Knudsen, "The Interpolation Attack on Block Ciphers," Fast Software Encryption, FSE'97, LNCS 1267, pp.28–40, 1997.
11. Kanda, "Practical Security Evaluation against Differential and Linear Cryptanalyses for Feistel Ciphers with SPN Round Function," Selected Areas in Cryptography, SAC2000, LNCS in this proceeding.
12. Knudsen, "Cryptanalysis of LOKI91," Advances in Cryptology — AUSCRYPT'92, LNCS 718, pp.196–208, 1993.
13. Knudsen, "Truncated and Higher Order Differentials," Fast Software Encryption — Second International Workshop, LNCS 1008, pp.196–211, 1995.
14. Kanda, Moriai, Aoki, Ueda, Takashima, Ohta, Matsumoto, "E2 — A New 128-Bit Block Cipher," IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, Vol.E83-A, No.1, pp.48–59, 2000.
15. Kelsey, Schneier, Wagner, "Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES," Advances in Cryptology — CRYPTO'96, LNCS 1109, pp.237–251, 1996.
16. Kanda, Takashima, Matsumoto, Aoki, Ohta, "A Strategy for Constructing Fast Round Functions with Practical Security against Differential and Linear Cryptanalysis," Selected Areas in Cryptography, SAC'98, LNCS 1556, pp.264–279, 1999.
17. Lai, Massey, Murphy, "Markov Ciphers and Differential Cryptanalysis," Advances in Cryptology — EUROCRYPT'91, LNCS 547, pp.17–38, 1991.
18. Matsui, "Linear Cryptanalysis Method for DES Cipher," Advances in Cryptology — EUROCRYPT'93, LNCS 765, pp.386–397, 1994.
19. Matsui, "On Correlation Between the Order of S-boxes and the Strength of DES," Advances in Cryptology — EUROCRYPT'94, LNCS 950, pp.366–375, 1995.
20. Matsui, "New Block Encryption Algorithm MISTY," Fast Software Encryption, FSE'97, LNCS 1267, pp.54–68, 1997.
21. Matsui, "Differential Path Search of the Block Cipher E2," Technical report of IEICE, ISEC99-19, pp.57–64, The Institute of Electronics, Information and Communication Engineers, 1999. (in Japanese)
22. Matsui, Inoue, Yamagishi, Yoshida, "A note on calculation circuits over $GF(2^{2n})$," Technical Report of IEICE, IT88-14, The Institute of Electronics, Information and Communication Engineers, 1988. (in Japanese)
23. Matsui, Tokita, "Cryptanalysis of a Reduced Version of the Block Cipher E2," Fast Software Encryption, FSE'99, LNCS 1636, pp.71–80, 1999.
24. Moriai, Sugita, Aoki, Kanda, "Security of E2 against Truncated Differential Cryptanalysis," Selected Areas in Cryptography, SAC'99, LNCS 1758, pp.106–117, 2000.
25. Rijmen, Daemen, Preneel, Bosselaers, Win, "The cipher SHARK," Fast Software Encryption — Third International Workshop, LNCS 1039, pp.99–111, 1996.
26. Wagner, "The Boomerang Attack," Fast Software Encryption, FSE'99, LNCS 1636, pp.156–170, 1999.