Camellia: A 128-Bit Block Cipher Suitable for Multiple Platforms

Kazumaro Aoki† Tetsuya Ichikawa‡ Masayuki Kanda†
Mitsuru Matsui‡ Shiho Moriai† Junko Nakajima‡ Toshio Tokita‡

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

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

Ver 1.0: July 13, 2000
Ver 2.0: September 26, 2001

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 cryptanalysis. 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 (1.13GHz) at the rate of 471 Mbits per second. In addition, a distinguishing feature is its small hardware design. A hardware implementation, which includes encryption, decryption, and the key schedule for 128-bit keys, occupies only 8.12K gates using a 0.18μm CMOS ASIC library. This is in the smallest class among all existing 128-bit block ciphers.

Contents

1 Introduction **1**

2 Design Rationale **3**
2.1 F-function . . . 3
2.2 P-function . . . 3
2.3 s-boxes . . . 3
2.4 FL- and FL^{-1}-functions . . . 3
2.5 Key Schedule . . . 4

3 Performance Figures **5**
3.1 Software Performance . . . 5
3.2 Hardware Performance . . . 5

4 Software Implementation Techniques **11**
4.1 Setup . . . 11
4.2 Data Randomization . . . 12
4.3 General Guidelines . . . 18

5 Hardware Evaluations **20**
5.1 Type 1: Fast Implementation-1 (Fully loop unrolled architecture) . . . 20
5.2 Type 2: Small Implementation-1 (Loop architecture) . . . 21
5.3 Type 3: Small Implementation-2 (Special Case for FPGA, Loop architecture) . . . 22
5.4 Type 4: Fast Implementation-2 (Pipeline architecture) . . . 24

6 Security **26**
6.1 Differential and Linear Cryptanalysis . . . 26
6.2 Truncated Differential Cryptanalysis . . . 27
6.3 Truncated Linear Cryptanalysis . . . 29
6.4 Cryptanalysis with Impossible Differential . . . 29
6.5 Boomerang Attack . . . 29
6.6 Higher Order Differential Attack . . . 30
6.7 SQUARE Attack . . . 30
6.8 Interpolation Attack and Linear Sum Attack . . . 31
6.9 No Equivalent Keys . . . 31
6.10 Slide Attack . . . 31
6.11 Related-key Attack . . . 32
6.12 Statistical Tests . . . 32
6.13 Implementation Attacks . . . 32
6.14 Brute Force Attacks . . . 33

7 Conclusion **35**

A History **41**

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 [BS93] and linear cryptanalysis [M94] 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 [K95, JK97], interpolation attacks [JK97, A00], related-key attacks [B94, KSW96], truncated differential attacks [K95, MT99], boomerang attacks [W99], and slide attacks [BW99, BW00].

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 GF(2^8). Moreover, we reduced the inversion function in GF(2^8) to a few GF(2^4) arithmetic operations. It enabled us to implement the s-boxes by fewer gate counts.

The key schedule is very 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.

Standardization activities. In March 2000 NTT and Mitsubishi Electric Corporation proposed Camellia in response to the call for contributions from ISO/IEC JTC 1/SC 27, aiming at its being adopted as an international standard. In September 2000, we submitted Camellia to

NESSIE (New European Schemes for Signature, Integrity, and Encryption) project as a strong cryptographic primitive. In September 2001, Camellia was selected as a candidates for the 2nd Phase of the NESSIE project.

Outline of the paper. This paper is organized as follows: Section 2 describes the rationale behind Camellia's design. Section 3 discusses the performance of Camellia. Section 4 contains the techniques for software implementation. In Section 5 we discuss our hardware evaluations. In Section 6 we evaluated Camellia's strength against known attacks. We conclude in Section 7.

For the specification of Camellia, please see the separate document titled "Specification of Camellia – a 128-bit Block Cipher." We will follow the definitions and notation given in this separate paper.

2 Design Rationale

2.1 *F*-function

The design strategy of the F-function of Camellia follows that of the F-function of E2 [KMA+98]. The main difference between E2 and Camellia is the adoption of the 1-round (conservative) SPN (Substitution-Permutation Network), 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 Section 6 for detailed discussions on security.

2.2 *P*-function

The design rationale of the P-function is similar to that of the P-function of E2. That is, for computational efficiency, it should be represented using only bytewise exclusive-ORs and for security against differential and linear cryptanalysis, its branch number should be optimal [KTM+99]. From among the linear transformations that satisfy these conditions, we chose one considering highly efficient implementation on 32-processors [AU00] and high-end smart cards, as well as 8-bit processors.

2.3 *s*-boxes

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

It is well known that the smallest of the maximum differential probability of functions in $\mathrm{GF}(2^8)$ was proven to be 2^{-6}, and the smallest of the maximum linear probability of functions in $\mathrm{GF}(2^8)$ is conjectured to be 2^{-6}. There is a function affine equivalent to the inversion function in $\mathrm{GF}(2^8)$ 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 $\mathrm{GF}(2^8)$ complicates the expressions of the s-boxes in $\mathrm{GF}(2^8)$, which makes interpolation attacks ineffective. Making the four s-boxes different slightly improves security against truncated differential cryptanalysis [MT99].

For small hardware design, the elements in $\mathrm{GF}(2^8)$ can be represented as polynomials with coefficients in the subfield $\mathrm{GF}(2^4)$. In other words, we can implement the s-boxes by using a few operations in the subfield $\mathrm{GF}(2^4)$ [MIYY88]. Two affine functions at the input and output of the inversion function in $\mathrm{GF}(2^8)$ also play a role in complicating the expressions of the s-boxes in $\mathrm{GF}(2^4)$.

2.4 *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 FL- and FL^{-1}-functions are similar to those of the FL-function of MISTY [M97]. The difference between MISTY and Camellia is the addition of 1-bit rotation. This is expected to make bytewise 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.

2.5 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/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 [DR98] or Serpent [ABK98], 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.

3 Performance Figures

3.1 Software Performance

Table 1 summarizes the current software performance of Camellia on the commonly-used 32-bit and 64-bit processors. Table 2 shows the software performance on the microprocessors used for smart cards and embedded systems, which are equipped with the restricted memory. Generally speaking, the first priority of the former is "Speed," while that of the latter is "RAM Usage and ROM Usage." Some of the data are published in [AIK+00a, C01, ISKM01, AIK+00b, Y01a, Y01b], but the others have not been published yet.

The tables show that Camellia can be efficiently implemented on low-end smart cards, 32-bit and 64-bit processors. We use the abbreviations M (mega) for 10^6 and m (milli) for 10^{-3} in the tables.

Optimization level. When we coded programs using assembly language, we tried to use many techniques described in Section 4 to achieve the best performance. However, there is a room for further improvement.

On the other hand, depending on the C compiler used, different assembly codes are produced from the same C code. This means that the assembly codes are not guaranteed to be optimal, even if the C code is optimized. Thus, we did not spend a long time on optimizing C code.

How to measure speed. It is difficult to measure speed on modern processors since there are many elements, for example, status of cache, that are beyond the users control and that influence speed. We decided to measure speed under the following conditions and assumptions:

- All codes and data are correctly aligned.
- Input and output texts and codes are preloaded to the first level cache.
- Branch predictions are correct.
- Setup function (except for on-the-fly implementations) generates subkey-dependent constants from the secret key, and the constants are used by encryption or decryption function.
- Encryption (decryption) function except for on-the-fly implementations can encrypt (decrypt) an integral number of blocks.
- We measured the speed many times, and chose the best result to eliminate cache hit misses and other uncontrollable factors.
- We averaged the speed numbers for large block encryption, but the values include all overheads including loop and function calls.

3.2 Hardware Performance

Table 3 represents the recent results on hardware performance of Camellia on ASIC (Application Specific Integrated Circuit) and FPGA (Field Programmable Gate Array). Table 4 shows the environment of our hardware design and evaluation.

Table 4: Hardware evaluation environment (ASIC, FPGA)

Language	(ASIC, FPGA) Verilog-HDL
Design library	(ASIC) Mitsubishi Electric 0.35μm CMOS ASIC library Mitsubishi Electric 0.18μm CMOS ASIC library 0.25 μm CMOS ASIC library (reported by CRYPTREC Report 2000) (FPGA) Xilinx XC4000XL series Xilinx VirtexE series
Simulator	(ASIC, FPGA) Verilog-XL (except for 0.25μm) (ASIC) VCS5.1 (used for 0.25μm)
Logic synthesis	(ASIC) Design Compiler version 1998.08 (used for 0.35μm) Design Compiler version 2000.11-SP1 (used for 0.18μm) Design Compiler version 2000.05-1 (used for 0.25μm) (FPGA) Synplify version 5.3.1 and ALLIANCE version 2.1i (used for XC4000XL series) Synplify version 6.1.3 and ALLIANCE version 3.3.07i (used for VirtexE series)

Table 5: Hardware design policies (outline)

Type	Top priority	Outline of logic
Type 1	Fast implementation from the viewpoint of Enc(Dec) speed	Figure 1
Type 2	Small implementation from the viewpoint of total logic size	Figure 2
Type 3	Small implementation (special case for FPGA)	Figure 3
Type 4	Pipeline implementation	Figure 4

We evaluated Type 1 through Type 4 logic. Table 5 shows the top priorities of the logic types. The details of each type are described in Section 5.

Table 1: Camellia Software Performance (updated on Aug.31, 2001)

Processor	Language	Key Size	Speed			RAM Usage(*1)		ROM Usage				Reference
			Key setup(*2)	Enc. / Dec.		Key setup(*2)	Enc. / Dec.	Total size	Key setup(*2)	Enc. / Dec.	Table	
		[bits]	[cycles]	[cycles]	[Mbps]	[bytes]	[bytes]	[bytes]	[bytes]	[bytes]	[bytes]	
Pentium III (*4)	Assembly	128	1,570	308	290.9	288	20	15,012	6,788	0	8,224	[AIK+00b]
		128	160	371	241.5	28	36	11,420	1,046	2,150	8,224	[AIK+00b]
		192	222	494	181.4	28	36	13,032	1,469	3,323	8,240	[AIK+00b]
		256	226	494	181.4	28	36	13,048	1,485	3,323	8,240	[AIK+00b]
Pentium III (*5)	Assembly	128	-	326	255.2	-	-	29,285	-	-	-	[C01]
		128	(Enc) 467 (*3)		0.72msec	-	-	20,110	-	-	-	[C01]
		128	(Dec) 474 (*3)		0.73msec	-	-	20,236	-	-	-	[C01]
Pentium II (*6)	ANSI C(*7)	128	263	577	66.6	44	64	9,461	1,600	3,733	4,128	[AIK+00b]
Pentium III (*8)	Java(*9)	128	9,091	793	161.4	-	-	-	-	-	-	Not published
Alpha 21264 (*10)	Assembly	128	158	326	261.9	48	48	21,040	1,600	2,928	16,512	[AIK+00b]
		128	118	339	251.8	48	48	20,736	1,132	3,076	16,528	[AIK+00b]
		192	176	445	191.9	48	48	22,196	1,668	4,000	16,528	[AIK+00b]
		256	176	445	191.9	48	48	22,204	1,676	4,000	16,528	[AIK+00b]
Alpha 21264 (*11)	Assembly	128	-	282	210.2	-	-	31,552	-	-	-	[C01]
		128	(Enc) 448 (*3)		0.97msec	-	-	25,792	-	-	-	[C01]
		128	(Dec) 435 (*3)		0.94msec	-	-	25,792	-	-	-	[C01]
UltraSPARCIIi (*12)	Assembly	128	-	355	144.2	-	-	15,240	-	-	-	[C01]
		128	403 (*3)		1.01msec	-	-	23,992	-	-	-	[C01]

(*1) The figure includes stack area, and excludes text area and key area.

(*2) Key schedule may be included.

(*3) The figure includes key generation and one block encryption. This is achieved by using the on-the-fly subkey generation.

(*4) IBM PC/AT compatible PC, Intel Pentium III (700MHz), 256KB on-die L2 cache, FreeBSD 4.0R, 128MB main memory.

(*5) IBM PC/AT compatible PC, Intel Pentium III (650MHz), 256KB on-die L2 cache, Windows98 SE, 64MB main memory.

(*6) IBM PC/AT compatible PC, Intel Pentium II (300MHz), 512KB L2 cache, Windows95, 160MB main memory.

(*7) Microsoft Visual C++ 6 with the optimization options **/G6 /Zp16 /ML /Ox /Ob2**.

(*8) IBM PC/AT compatible PC, Intel Pentium III (1GHz), 256KB on-die L2 cache, Windows2000, 512MB main memory.

(*9) IBM Java Compiler 1.2.2 and IBM Java VM 1.2.2.

(*10) Alpha 21264 (667MHz), Tru64 UNIX 4.0F, 2GB main memory.

(*11) Alpha 21264 (463MHz), Tru64 UNIX V5.1, 512MB main memory.

(*12) Ultra SPARC IIi (400MHz), Solaris 7, 256MB main memory.

Table 2: Camellia Software Performance for Smart Cards and Embedded Systems (updated on Aug.31, 2001)

Processor	Language	Key size [bits]	Speed		RAM Usage		ROM Usage					Reference
			Key setup(*1) [cycles]	Enc. / Dec. [cycles]	Key setup(*1) [bytes]	Enc. / Dec. [bytes]	Total size [bytes]	Key setup(*1) [bytes]	Enc. / Dec. [bytes]	Table [bytes]	Sharing size(*2) [bytes]	
8051 (*7)	Assembly	128		10,217 (*3) 10.22msec	0	32 (*4)	990	0	702	288	0	[AIK+00b]
Z80 (*8)	Assembly	128	5,146 1.03msec	28,382 5.68msec	44 (*5)	62 (*5)	1,698	358	1,183	288	-131	Not published
		128		(Enc) 35,951 (*3) 7.19msec	0	60 (*5)	1,268	-	1,023	-	-797	Not published
				(Dec) 37,553 (*3) 7.51msec	0	60 (*5)		-	1,042	-		
H8/3113 (*9)	Assembly	128	2,380 0.95msec	4,100 1.64msec	208 (*6)	0	-	-	-	-	-	[Y01a]
MC68HC705B16 (*10)	Assembly	128	7,500 3.57msec	9,900 4.71msec	208 (*6)	0	-	-	-	-	-	[Y01a]
MC68HC908AB32 (*11)	Assembly	128	5,679 0.71msec	8,430 1.05msec	208 (*6)	0	-	-	-	-	-	[Y01b]
M32Rx/D (*12)	Assembly	128	642 6.42msec	1,236 12.36msec	44 (*5)	44 (*5)	8,684	1,392	3,164	4,128	0	[AIK+00b]

(*1) Key schedule may be included.

(*2) Some ROM size may be reduced, since some functions can be shared among key generation, encryption and decryption.

(*3) The figure includes key generation and one block encryption. This is achieved by using the on-the-fly subkey generation.

(*4) The figure includes stack area, and excludes text area and key area.

(*5) The figure includes stack area, text area and key area.

(*6) The figure shows the size of round keys.

(*7) Intel 8051 (12MHz; 1cycle=12 oscillator periods) simulator on Unix.

(*8) Z80 (5MHz) simulator on Windows.

(*9) Hitachi H8/3113 (5MHz; 1cycle=2 oscillator periods) on Hitachi's E6000 Emulator.

(*10) Motorola 6805 series MC68HC705B16 (2.1MHz) on Motorola's In-Circuit Simulator Kits.

(*11) Motorola 6805 series MC68HC908AB32 (8MHz) on Motorola's In-Circuit Simulator Kits.

(*12) Mitsubishi 32-bit microcomputer M32Rx/D (100MHz) on MSA2310 evaluation board.

Table 3: Camellia Hardware Performance (updated on Aug.31, 2001)

Architecture			Design Library	Key size [bits]	Speed				Area Size				Efficiency	Reference
					Key setup [nsec]	Max. delay(*1) [nsec]	Latency [cycles]	Throughput [Mbps]	Unit	Total(*2) [unit]	Key expan.(*3) [unit]	Enc./Dec.(*4) [unit]	Throughput/Area [Kbps/unit]	
ASIC	Unrolled	Type 1	Mitsubishi 0.35μm	128	24.36	109.35	1	1,170.55	Kgate	272.82	55.91	216.91	4.29	[AIK+00b]
			Mitsubishi 0.18μm	128	40.00	40.00	1	3,200.00		355.10	-	-	9.01	Not published
			Mitsubishi 0.18μm	128	45.96	45.96	1	2,785.00		244.90	-	-	11.37	Not published
	Loop	Type 2	Mitsubishi 0.35μm	128	110.20	27.67	21	220.28	Kgate	11.35	4.98	6.37	19.41	[AIK+00a]
			Mitsubishi 0.35μm	128	117.04	28.73	21	212.16		9.66	5.75	3.91	21.96	[AIK+00b]
			Mitsubishi 0.18μm	128	144.88	36.22	21	168.28		8.51	-	-	19.77	Not published
			Mitsubishi 0.18μm	128	25.92	6.48	21	940.62		27.46	-	-	34.25	Not published
			Mitsubishi 0.18μm	128	28.20	7.05	21	864.57		21.45	-	-	40.31	Not published
			Mitsubishi 0.18μm	128	23.20	5.80	21	1,050.90		11.87	-	-	88.52	Not published
			Mitsubishi 0.18μm	128	137.24	34.31	21	177.65		8.12	-	-	21.87	Not published
			Mitsubishi 0.18μm	128	12.96	3.24	21	1,881.25		44.30	-	-	42.47	Not published
			0.25μm	256	-	5.46	-	837.00		39.35	22.76	16.33	21.27	[C01]
			0.25μm	256	-	11.51	-	397.00		23.12	13.30	9.67	17.17	[C01]
FPGA	Loop	Type 2	Xilinx XC4000XL	128	362.83	78.82	21	77.34	CLB	1,296	-	-	59.68	[AIK+00b]
		Type 3	Xilinx XC4000XL	128	-	50.00	21	122.01		874	-	-	139.60	[AIK+00b]
		Type 2	Xilinx VirtexE	128	135.03	30.56	21	199.46	Slice	1,816	-	-	109.83	[ISKM01]
			Xilinx VirtexE	128	126.00	28.80	21	211.90		1,816	-	-	116.69	[ISKM01]
			Xilinx VirtexE	128	127.04	26.80	21	227.42		1,780	-	-	127.76	[ISKM01]
	Unrolled	Type 1	Xilinx VirtexE	128	97.70	318.50	1	401.89		9,426	-	-	42.64	[ISKM01]
	Pipeline	Type 4	Xilinx VirtexE	128	83.25	18.96	20	6,749.99		9,692	-	-	696.45	[ISKM01]

(*1) Critical path of data encryption (or decryption)

(*2) The figure includes key scheduling logic, encryption/decryption logic, data selector (if necessary), output register, subkey register and buffers for fan-out adjustment.

(*3) The figure includes subkey register.

(*4) The figure includes output register and data selector (if necessary).

4 Software Implementation Techniques

This section describes how to implement Camellia efficiently in software. In most cases, an implementation can be divided into two parts: *setup* including key schedule and *data randomization*, that is, encryption or decryption. We first describe how to optimize the setup code, and then describe how to optimize the data randomization code.

This section describes specific techniques for 8-, 32-, or 64-bit processors. However, a technique for 8-bit processors may be applicable to 32- or 64-bit processors and a technique for 32-bit processors may be applicable to 64-bit processors. Other word sizes may need to be considered.

We assume that you first implement Camellia using the specification as it is. This section will optimize the resulting code.

Note that in this section "word" means the natural size of the target processor. For example, the words of IA-32 without MMX technology, IA-32 with MMX technology and Alpha are 32-, 64-, and 64-bits long respectively.

4.1 Setup

4.1.1 Store All Subkeys

Store all subkeys into memory once you generate them if you have sufficient memory, and use the stored subkeys for data randomization.

4.1.2 Subkey Generation Order

You do not have to compute subkeys in order. For example, when you compute subkeys for a 128-bit key, first compute the subkeys that only depend on K_L, and then compute subkeys that only depend on K_A. This reduces the number of registers or memory for storing K_A.

4.1.3 XOR Cancellation Property in Key Schedule

The key schedule of Camellia is based on the Feistel structure. Between the 2nd round and the 3rd round, K_L is XORed to an intermediate value. This structure causes cancellations of K_L. More precisely, the input of the 3rd round can be computed by the following equations.

$$\begin{cases} (\text{right half}) &= F(K_{LL}, \Sigma_1) \\ (\text{left half}) &= F(K_{LR} \oplus (\text{right half}), \Sigma_2) \end{cases} \quad \text{for 128-bit keys}$$

$$\begin{cases} (\text{right half}) &= K_{RR} \oplus F(K_{LL} \oplus K_{RL}, \Sigma_1) \\ (\text{left half}) &= K_{RL} \oplus F(K_{LR} \oplus (\text{right half}), \Sigma_2) \end{cases} \quad \text{for 192- and 256-bit keys}$$

Using the above equations, we can eliminate 3 and 2 XORs in **L** for 128- and 192/256-bit keys, respectively, compared to the straightforward implementation of the specification.

4.1.4 Rotation Bits for K_L, K_R, K_A, and K_B

You do not need to keep K_L, K_R, K_A, and K_B, but you should keep their rotated values when generating subkeys. You can generate subkeys by rotating the kept values by a sum of integral multiples of 16 ± 1 bits.

4.1.5 kl_5 and kl_6 Generation from k_{11} and k_{12}

For 192- and 256- bit keys, you can use word-oriented rotation to generate (kl_5, kl_6) from (k_{11}, k_{12}), since (kl_5, kl_6) equals $(k_{11}, k_{12}) \lll_{32}$. This saves a few instructions compared to general rotation.

4.1.6 On-the-fly Subkey Generation

You can generate subkeys *on-the-fly*. All subkeys are one of the rotated values of K_L, K_R, K_A, and K_B. Thus, you first generate K_L, K_R, K_A, and K_B, and then rotate them to get the subkeys. Refer to Section 4.1.4 for the rotated numbers of bits for K_L, K_R, K_A, and K_B.

4.1.7 128-bit key and 192/256-bit key

If your code does not need to use key sizes larger than 128 bits, you do not need to generate K_B. That is, you can omit the computations for the last two F-functions.

4.1.8 How to Rotate an Element in Q

8-bit processor. As stated in Section 4.1.4, the amount of rotation in bits is a sum of integral multiples of 16 ± 1. Thus, you can rotate an element in **Q** by 16 ± 1 bits by rotating 1-bit left or right followed by a 2-byte move.

32-bit processor. Consider the use of a double precision shift instruction: `shrd` or `shld` if you are programming on IA-32.

4.1.9 F-function

Key schedule includes F-functions, but the main usage of the F-function is for data randomization. Refer to Section 4.2.

4.1.10 Keyed Functions

Camellia has three keyed functions: bitwise XOR, bitwise OR, and bitwise AND. Consider the use of a self-modifying code, if possible.

4.2 Data Randomization

4.2.1 Endian Conversion

Camellia prefers big endian. Thus, the code for little endian processors needs additional code for endian conversions.

The most straightforward implementation converts the endian when loading a register from memory and storing a register to memory. Only FL- and FL^{-1}-functions are endian dependent. More precisely, only the 1-bit rotation in FL- or FL^{-1}-function is endian dependent. This means that you can convert endians just before or just after the 1-bit rotation with the appropriate

subkey generation scheme. A combination of computing endian conversion and 1-bit rotation may increase the performance of Camellia. Details are described in Section 4.2.2.

Some processors have a special instruction for endian conversion. For example, IA-32 (after 80486) has `bswap` instruction. Use these instructions. However, do not use the byte swap technique described in [C98, Appendix A]. The technique reduces the code size, but it is not fast, since the memory load and store instruction incurs long latency.

As described above, the endian problem only effects the 1-bit rotation of a 32-bit word. Thus, we do not need full 64-bit word endian conversion.

The following are general methods to realize endian conversion for 32-bit register x. In the following techniques, you can use either $\cup$ or $\oplus$ instead of $+$ in the equations, and you can switch the computational order between shifts including rotations and ANDs with an appropriate conversion of masked constants.

Straightforward.

$$x \leftarrow (x \lll_{24}) + ((x \cap \texttt{0xff00}) \lll_{8}) + ((x \ggg_{8}) \cap \texttt{0xff00}) + (x \ggg_{24})$$

The technique has high parallelism.

Minimum operations without rotation.

$$\begin{aligned} x &\leftarrow (x \lll_{16}) + (x \ggg_{16}) \\ x &\leftarrow ((x \cap \texttt{0xff00ff}) \lll_{8}) + ((x \ggg_{8}) \cap \texttt{0xff00ff}) \end{aligned}$$

Using rotations.

$$x \leftarrow ((x \cap \texttt{0xff00ff}) \ggg_{8}) + ((x \lll_{8}) \cap \texttt{0xff00ff})$$

Using SSE. New Intel Pentium family processors including Pentium III have a very effective instruction for reordering data, which is called `pshufw` [I99]. 5 instructions including `pshufw` are sufficient to convert endian for 64-bit data.

4.2.2 1-bit Rotation in Little Endian Interpretation

As described in Section 4.2.1, we do not need endian conversion when loading and storing texts if we can efficiently implement 1-bit rotation in FL- and FL^{-1}-functions.

Assuming x to be a 32-bit register that contains little endian data to be rotated by 1-bit, we can compute 1-bit rotation by the following equation.

$$x \leftarrow ((2x) \cap \texttt{0xfefefefe}) + ((x \ggg_{15}) \cap \overline{\texttt{0xfefefefe}}) \tag{1}$$

Of course, this technique requires an appropriate changes to subkey setup and other functions.

Note that $+$ in Equation (1) can be replaced with $\cup$ or $\oplus$, and computing $2x$ can be done by $\lll_1$, $\lll_1$ or addition with x itself, and you can switch the computational order between shifts including rotations and ANDs with an appropriate conversion of masked constants.

Confirm whether your processor has ANDNOT instruction, such as `pandn` in IA-32 and `bic` in Alpha. In this case, you do not need to prepare the constant, $\overline{\texttt{0xfefefefe}}$.

4.2.3 Whitening

The key additions kw_2 and kw_4 can be combined into other keyed operations using the following equations.

$$\begin{aligned}
(x \oplus k) \oplus y &= (x \oplus y) \oplus k, \\
(x \oplus k) \oplus l &= x \oplus (k \oplus l), \\
(x \oplus k) \cap l &= (x \cap l) \oplus (k \cap l), \\
(x \oplus k) \lll_1 &= (x \lll_1) \oplus (k \lll_1), \\
(x \oplus k) \cup l &= (x \cup l) \oplus (k \cap \bar{l}),
\end{aligned} \tag{2}$$

where x, y, k, l are bit strings. Adjust subkeys at setup to eliminate 2 XORs in **L**.

4.2.4 Key XOR

Using Equations (2), you can move key XORs to any place if the movement does not go through the S-function. For example, changing F-function computation $P(S(X \oplus k))$ to $P(S(X)) \oplus k'$ may improve instruction scheduling.

4.2.5 *S*-function

s_1 is defined by the arithmetics in GF(2^8). However, do not compute GF(2^8) arithmetics; instead precompute and hard-code a table in your program, see Table 4 in the specification.

We strongly suggest that you also precompute and hard-code s_2, s_3, and s_4 tables in addition to s_1, if you have sufficient memory and 8-bit rotation is expensive. If you do not have sufficient memory, the data of s_2, s_3, and s_4 can be generated from the table for s_1 using one rotation (See Section 4.5 in the specification).

If you have sufficient memory, and cost of table lookup is heavy, as is true for the current Java virtual machines, consider the use of a two s-box combined table, for example $(s_1(y_1), s_2(y_2))$.

4.2.6 *P*-function

32-bit processor. Let $(Z_L, Z_R) = ((z_1, z_2, z_3, z_4), (z_5, z_6, z_7, z_8))$ be the input of P-function and $(Z'_L, Z'_R) = ((z'_1, z'_2, z'_3, z'_4), (z'_5, z'_6, z'_7, z'_8))$ be the output of P-function.

From Figure 5 in the specification, you can see that P-function can be computed as follows.

$$\begin{aligned}
Z_L &\leftarrow Z_L \oplus (Z_R \lll_8) \\
Z_R &\leftarrow Z_R \oplus (Z_L \lll_{16}) \\
Z_L &\leftarrow Z_L \oplus (Z_R \ggg_8) \\
Z_R &\leftarrow Z_R \oplus (Z_L \ggg_8) \\
Z'_L &\leftarrow Z_R \\
Z'_R &\leftarrow Z_L
\end{aligned}$$

The critical path of this computation is long. We can modify the computation as follows.

$$
\begin{array}{lcllcl}
 & & & Z_R & \leftarrow & Z_R \lll_8 \\
Z_L & \leftarrow & Z_L \oplus Z_R & Z_R & \leftarrow & Z_R \lll_8 \\
Z_L & \leftarrow & Z_L \ggg_8 & Z_R & \leftarrow & Z_R \oplus Z_L \\
Z_L & \leftarrow & Z_L \oplus Z_R & Z_R & \leftarrow & Z_R \lll_{16} \\
Z_L & \leftarrow & Z_L \lll_8 & Z_R & \leftarrow & Z_R \oplus Z_L \\
Z'_L & \leftarrow & Z_R & Z'_R & \leftarrow & Z_L
\end{array}
$$

The critical path of the above computation is decreased. It seems that the technique requires one additional rotation, however, you can probably combine the first step of the above computation and S-function without any additional cost.

8-bit processor (orthogonal mnemonics). If the instruction in your processor can XOR any combination of registers and has sufficient registers, you can compute P-function by using just 16 XORs using Figure 5 in the specification.

8-bit processor (accumulator based). If your processor is accumulator based, minimizing the number of XORs is not always a good idea, since the computation may require register load from memory and store into memory many times. The following computation is optimized for an accumulator based processor.

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

When indexing z'_i costs many operations, the following is useful.

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

4.2.7 Substitution and Permutation

This section describes how to efficiently compute $P \circ S$ compared to independently computing S and P.

64-bit processor. If your processor has a sufficiently large first level cache, use the technique described in [RDP+96]. The technique prepares the following tables defined by Equations (3).

$$\begin{array}{rcl}
SP_1(y_1) & = & (s_1(y_1),\ s_1(y_1),\ s_1(y_1),\ 0,\ s_1(y_1),\ 0,\ 0,\ s_1(y_1)) \\
SP_2(y_2) & = & (0,\ s_2(y_2),\ s_2(y_2),\ s_2(y_2),\ s_2(y_2),\ s_2(y_2),\ 0,\ 0) \\
SP_3(y_3) & = & (s_3(y_3),\ 0,\ s_3(y_3),\ s_3(y_3),\ 0,\ s_3(y_3),\ s_3(y_3),\ 0) \\
SP_4(y_4) & = & (s_4(y_4),\ s_4(y_4),\ 0,\ s_4(y_4),\ 0,\ 0,\ s_4(y_4),\ s_4(y_4)) \\
SP_5(y_5) & = & (0,\ s_2(y_5),\ s_2(y_5),\ s_2(y_5),\ 0,\ s_2(y_5),\ s_2(y_5),\ s_2(y_5)) \\
SP_6(y_6) & = & (s_3(y_6),\ 0,\ s_3(y_6),\ s_3(y_6),\ s_3(y_6),\ 0,\ s_3(y_6),\ s_3(y_6)) \\
SP_7(y_7) & = & (s_4(y_7),\ s_4(y_7),\ 0,\ s_4(y_7),\ s_4(y_7),\ s_4(y_7),\ 0,\ s_4(y_7)) \\
SP_8(y_8) & = & (s_1(y_8),\ s_1(y_8),\ s_1(y_8),\ 0,\ s_1(y_8),\ s_1(y_8),\ s_1(y_8),\ 0)
\end{array} \tag{3}$$

Next, compute the following equation:

$$(z'_1, z'_2, z'_3, z'_4, z'_5, z'_6, z'_7, z'_8) \leftarrow \bigoplus_{i=1}^{8} SP_i(y_i)$$

This technique requires the following operations.

# of table lookups	8
# of XORs	7
Size of table (KB)	16

If the first cache of the target processor is moderately large, replace a few of the tables defined by Equations (3) with the tables below.

$$\begin{array}{rcl}
SP_\alpha(y) & = & (s_1(y),\ s_1(y),\ s_1(y),\ s_1(y),\ s_1(y),\ s_1(y),\ s_1(y),\ s_1(y)) \\
SP_\beta(y) & = & (s_2(y),\ s_2(y),\ s_2(y),\ s_2(y),\ s_2(y),\ s_2(y),\ s_2(y),\ s_2(y)) \\
SP_\gamma(y) & = & (s_3(y),\ s_3(y),\ s_3(y),\ s_3(y),\ s_3(y),\ s_3(y),\ s_3(y),\ s_3(y)) \\
SP_\delta(y) & = & (s_4(y),\ s_4(y),\ s_4(y),\ s_4(y),\ s_4(y),\ s_4(y),\ s_4(y),\ s_4(y))
\end{array} \tag{4}$$

Then, mask the necessary byte positions. This technique requires the following operations if you use just tables of Equations (4).

# of table lookups	8
# of XORs	7
# of ANDs	8
Size of table (KB)	8

When implementing this technique on Alpha architecture [C98], and if the number of registers is insufficient for storing constants for masking operation, use `zap` or `zapnot` instructions.

If your processor can efficiently copy half bits of a register to the other half, for example, `punpckldq/punpckhdq` or `pshufw` instructions in IA-32 [I99] which are realized after Pentium with MMX technology and Pentium III, respectively, prepare SP_1, SP_2, SP_3, and SP_4 defined in Equations (3). Then, compute the following equation:

$$\begin{aligned}&(z'_1, z'_2, z'_3, z'_4, z'_5, z'_6, z'_7, z'_8)\\ &\leftarrow SP_1(y_1) \oplus SP_2(y_2) \oplus SP_3(y_3) \oplus SP_4(y_4) \oplus \nu(SP_1(y_8) \oplus SP_2(y_5) \oplus SP_3(y_6) \oplus SP_4(y_7)),\end{aligned}$$

where ν denotes the operation that copies the first 4 bytes to the last 4 bytes. This technique requires the following operations.

# of table lookups	8
# of XORs	7
# of νs	1
Size of table (KB)	8

32-bit processor. [AU00] shows efficient implementations of Camellia-type substitution and permutation networks. One of the techniques prepares the following tables defined by Equations (5):

$$\begin{aligned}SP_{1110}(y) &= (s_1(y),\ s_1(y),\ s_1(y),\ 0)\\ SP_{0222}(y) &= (0,\ s_2(y),\ s_2(y),\ s_2(y))\\ SP_{3033}(y) &= (s_3(y),\ 0,\ s_3(y),\ s_3(y))\\ SP_{4404}(y) &= (s_4(y),\ s_4(y),\ 0,\ s_4(y))\end{aligned} \tag{5}$$

Then, compute as follows:

$$\begin{aligned}D &\leftarrow SP_{1110}(y_8) \oplus SP_{0222}(y_5) \oplus SP_{3033}(y_6) \oplus SP_{4404}(y_7)\\ U &\leftarrow SP_{1110}(y_1) \oplus SP_{0222}(y_2) \oplus SP_{3033}(y_3) \oplus SP_{4404}(y_4)\\ (z'_1, z'_2, z'_3, z'_4) &\leftarrow D \oplus U\\ (z'_5, z'_6, z'_7, z'_8) &\leftarrow (z'_1, z'_2, z'_3, z'_4) \oplus (U \ggg_8)\end{aligned}$$

This technique requires the following operations.

# of table lookups	8
# of XORs	8
# of rotations	1
Size of table (KB)	4

[AU00] also shows an implementation that is suitable for a processor in which rotation is very costly. The technique prepares the following tables in addition to tables defined by Equations (5):

$$\begin{aligned}SP_{1001}(y) &= (s_1(y),\ 0,\ 0,\ s_1(y))\\ SP_{2200}(y) &= (s_2(y),\ s_2(y),\ 0,\ 0)\\ SP_{0330}(y) &= (0,\ s_3(y),\ s_3(y),\ 0)\\ SP_{0044}(y) &= (0,\ 0,\ s_4(y),\ s_4(y))\end{aligned}$$

Then, compute as follows:

$$\begin{aligned} D &\leftarrow SP_{1110}(y_8) \oplus SP_{0222}(y_5) \oplus SP_{3033}(y_6) \oplus SP_{4404}(y_7) \\ (z'_1, z'_2, z'_3, z'_4) &\leftarrow D \oplus SP_{1110}(y_1) \oplus SP_{0222}(y_2) \oplus SP_{3033}(y_3) \oplus SP_{4404}(y_4) \\ (z'_5, z'_6, z'_7, z'_8) &\leftarrow D \oplus SP_{1001}(y_1) \oplus SP_{2200}(y_2) \oplus SP_{0330}(y_3) \oplus SP_{0044}(y_4) \end{aligned}$$

This technique requires the following operations.

# of table lookups	12
# of XORs	11
Size of table (KB)	8

4.2.8 Making Indices for *s*-box

You can make an index for s-box by simply using shifts and ANDs. However, several processors have special instructions for making an index, for example, `movzx` in IA-32 [I99] and `extbl` in Alpha [C98].

`movzx` is a fast operation in P6, but it can be used only for the two least significant bytes. A straightforward implementation uses `eax`, `ebx`, `ecx`, and `edx` registers for storing (L_r, R_r), and 2 rotations are used for making indices; 2 rotations are used for recovering byte order in the registers every round. However, you can remove 2 rotations for recovering byte order every round if you prepare rotated tables. Note that the byte order in registers returns to a natural order every 4 rounds.

4.3 General Guidelines

This section describes general guidelines. The guidelines are useful to optimize Camellia as well as other block ciphers. Please refer to the optimization manuals for each processor.

Avoid misaligned data accesses. Almost all processors penalize misaligned data access. Align data to the word boundary.

Avoid partial data accesses. Most processors have a function to access a smaller part than word size. However, this function may cause a penalty. Do not access partial data, even if you do not need full size of word and you have sufficient memory.

Be careful of the size of the cache. If the program or its data exceeds the size of the cache, the speed of the program will significantly decrease. Loop unrolling and table expansion are good techniques to speed up the program, but do not exceed the size of the cache.

Use intrinsic functions. Several compilers support intrinsic functions. For example, when you use Microsoft Visual C++ version 6 compiler on IA-32, and declare "`#pragma intrinsic(_lrotl)`" and use "`_lrotl`", the compiler generates rotation instructions in assembly language. Refer to the manual of the compiler that you use for details.

Measuring precise speeds is difficult. The running time of your code depends on many factors: cache hit misses, OS interrupts, and so on. Furthermore, the cryptographic

properties, for example, the number of blocks to be encrypted, also effect the running time.

A few processors have an instruction to get the time stamp. For example, IA-32 (after Pentium) has `rdtsc` [I99] and Alpha has `rpcc` [C98]. It is a good idea to use the time stamp counter for measuring speeds, but you should not directly apply these instructions to out-of-order architectures such as P6 and EV6.

If you want to measure speed precisely, consult good guidebooks. For example, if you use Pentium family processors, refer to [F00].

5 Hardware Evaluations

In Section 3, we showed evaluation results of hardware implementations (ASIC, FPGA) of Camellia. In this Section, we describe the design policies of the four types of logic evaluated in Section 3. The details of each type are described below.

5.1 Type 1: Fast Implementation-1 (Fully loop unrolled architecture)

In Type 1, we evaluate the hardware implementation (ASIC and FPGA) where the goal is to achieve the fastest encryption and decryption speed with no consideration of logic size. Figure 1 outlines the Type 1 logic. Table 6 shows the basic Type 1 components.

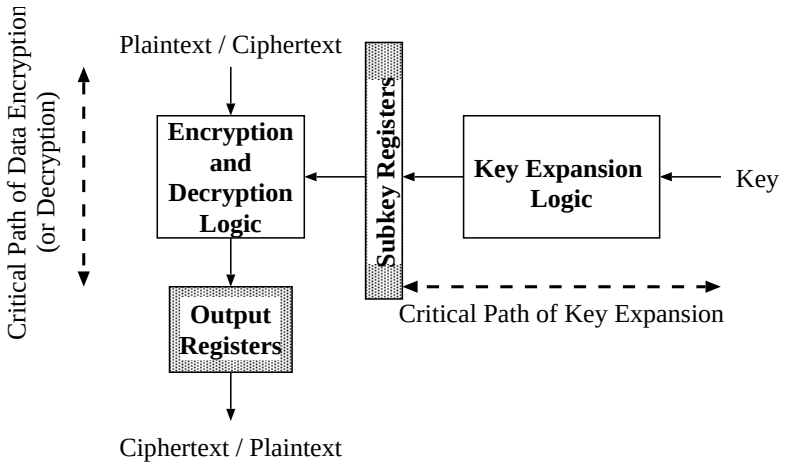

Figure 1: Outline of Type 1 (ASIC, FPGA)

Table 6: The basic Type 1 components

Encryption and decryption logic	Data randomizing logic for encryption and decryption, which consists of combinational logic.
Output register	Register for the encryption (decryption) data.
Key expansion logic	Logic in which subkeys are generated from key, which consists of combinational logic.
Subkey register	Register for the output data of key expansion logic.

The design policies of these basic components are listed below.

1. "Encryption and decryption logic" and "Key expansion logic"
   (a) Loop architecture is not introduced.

(b) Pipeline architecture is not introduced.

(c) Substitution tables (s-boxes) are designed by logic synthesis tool.

2. "Output register" and "Subkey register"

(a) The size of Output register is one block (=128 bits).

(b) The size of Subkey register is the total length of all subkeys in the algorithm.

Under the above design policies, we evaluated Camellia on ASIC and FPGA devices. The results are summarized in Table 3 in Section 3. "Throughput" is defined as follows:

$$\text{Throughput[b/s]} = \frac{\text{Block size(128 [bits])}}{\text{Critical path of data encryption(decryption)[sec]}}.$$

5.2 Type 2: Small Implementation-1 (Loop architecture)

In Type 2, we evaluate the hardware implementations on ASICs and FPGAs with the goal of achieving the smallest logic in encryption (and decryption). Figure 2 outlines the Type 2 logic. Table 7 shows the basic Type 2 components.

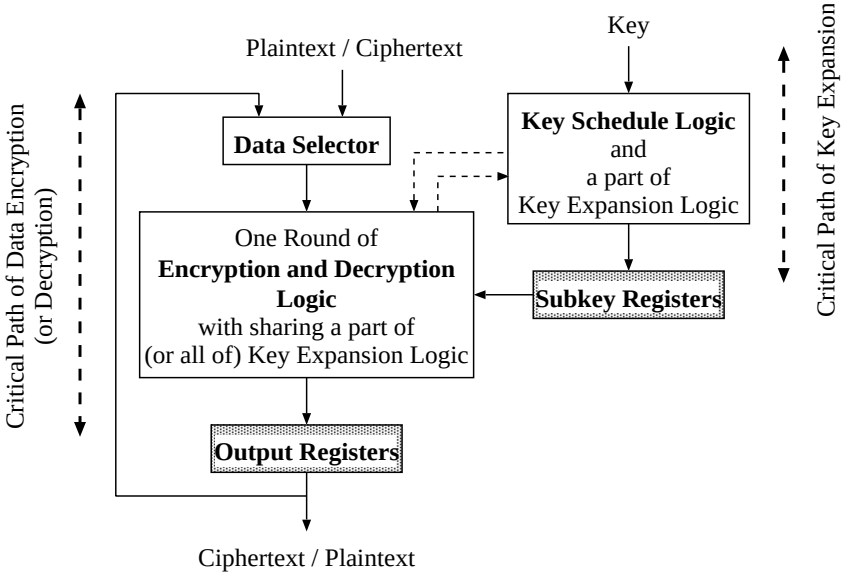

Figure 2: Outline of Type 2 (ASIC, FPGA)

The design policies of these basic components are as follows.

1. "Encryption and decryption logic" and "Key scheduling logic"

(a) Loop architecture is introduced (which consists of one round operation).

(b) Pipeline architecture is not introduced.

Table 7: The basic Type 2 components

Encryption and decryption logic	Data randomizing logic for one round operation of encryption and decryption, which includes (a part of) key expansion logic, and consists of combinational logics.
Output register	Register for the output (and intermediate) data.
Data selector	Selector which selects either encryption/decryption data or output data.
Key scheduling logic	Logic in which subkeys are generated using (a part of) key expansion logic in encryption and decryption logic and consists of combinational logics.
Subkey register	Register for the output data of key scheduling logic.

(c) Substitution tables (s-boxes) are optimized by hand.

(d) Key scheduling logic consists (a part of) key expansion logic and control logic.

2. "Output register", "Subkey register" and "Data selector"

(a) The size of Output register is one block (=128 bits).

(b) The size of Subkey register is that of the subkeys used in Encryption and decryption logic.

(c) Data selector is 2-1 selector, whose size is one block (=128 bits).

Under the above design policies, we evaluated Camellia on ASICs and FPGAs. The results are summarized in Table 3 in Section 3. "Throughput" is defined as follows:

$$\text{Throughput[b/s]} = \frac{\text{Block size(128 [bits])}}{\text{Critical path of data encryption(decryption)[sec]} \times \text{ latency}}.$$

5.3 Type 3: Small Implementation-2 (Special Case for FPGA, Loop architecture)

In Type 3, we evaluated the hardware implementation (FPGA) as a special case of Type 2. In Type 3, we assume that all subkeys are given and are loaded into FPGA internal memory. Figure 3 outlines the Type 3 logic. Table 8 shows the basic Type 3 components.

The design policies of these basic components are as follows.

1. "Encryption and decryption logic"

(a) Loop architecture is introduced (which consists of one round operation).

(b) Pipeline architecture is not introduced.

(c) Substitution tables (s-boxes) are optimized by hand.

2. "Output register", "Subkey memory" and "Data selector"

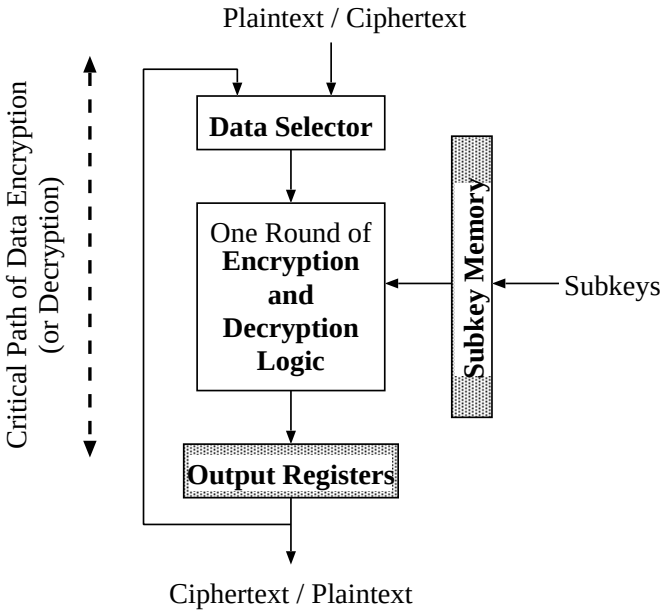

Figure 3: Outline of Type 3 (FPGA)

Table 8: The basic Type 3 components

Encryption and decryption logic	Data randomizing logic for one round operation of encryption and decryption, which includes (a part of) key expansion logic, and consists of combinational logic.
Output register	Register for the output (and intermediate) data.
Data selector	Selector which selects either encryption (decryption) data or output data.
Subkey memory	Memory for the subkeys loaded from outside.

(a) The size of Output register is one block (=128 bits).

(b) The size of Subkey memory is the length of all subkeys in the algorithm.

(c) Data selector is 2-1 selector whose size is one block (=128 bits).

Under the above design policies, we evaluated Camellia on an FPGA. The results are summarized in Table 3. "Throughput" is defined as follows:

$$\text{Throughput[b/s]} = \frac{\text{Block size(128 [bits])}}{\text{Critical path of data encryption(decryption)[sec]} \times \text{ latency}}.$$

5.4 Type 4: Fast Implementation-2 (Pipeline architecture)

In Type 4, we evaluate the hardware implementation (FPGA) where the goal is to achieve the fastest encryption and decryption speed with no consideration of logic size. (The pipeline architecture cannot realize any feedback modes, such as CBC, CFB, and OFB). Figure 4 outlines the Type 4 logic. Table 9 shows the basic Type 4 components.

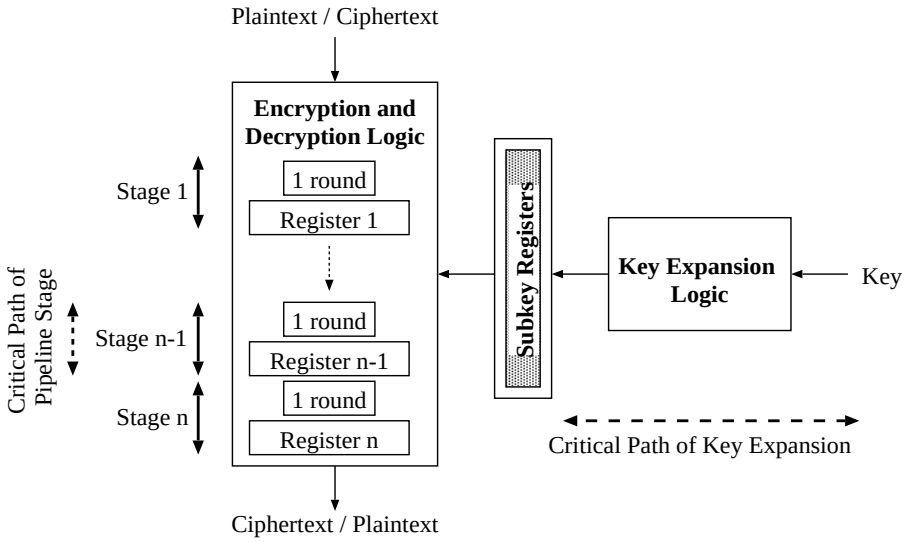

Figure 4: Outline of Type 4 (FPGA)

The design policies of these basic components are as follows.

1. "Encryption and decryption logic" and "Key expansion logic"

   (a) Loop architecture is not introduced.

   (b) Pipeline architecture is introduced.

   (c) Substitution tables (s-boxes) are designed by logic synthesis tool.

   (d) The size of Registers ($1 \sim n$) is one block (= 128 bits).

Table 9: The basic Type 4 components

Encryption and decryption logic	Data randomizing logic for encryption and decryption, which consists of combinational logic, and Registers ($1 \sim n$) for the output and intermediate data.
Key Expansion logic	Combinational logic in which subkeys are generated from the key.
Subkey register	Register for the output data from key expansion logic.

2. "Subkey register"

   (a) The size of Subkey register is the total length of all subkeys in the algorithm.

Under the above design policies, we evaluated Camellia on an FPGA devices. The results are summarized in Table 3. "Throughput" is defined as follows:

$$\text{Throughput[b/s]} = \frac{\text{Block size (128 [bits])}}{\text{Critical path of Pipeline Stage [sec]}}.$$

6 Security

6.1 Differential and Linear Cryptanalysis

The most well-known and powerful approaches to attacking many block ciphers are differential cryptanalysis, proposed by Biham and Shamir [BS93], and linear cryptanalysis, introduced by Matsui [M94]. There are several methods of evaluating security against these attacks, where there is a kind of "duality" relation between them [M95, CV95]: in other words, the security against both attacks can be evaluated in similar ways.

It is known that the upper bounds of differential/linear characteristic probabilities can, for several block ciphers, be estimated using the minimum numbers of differential/linear active s-boxes in some consecutive rounds. Kanda [K00] shows the minimum numbers of differential/linear active s-boxes for Feistel ciphers with conservative SPN (S-P) round function. Hereafter, we assume that linear transformation P is bijective.

Definition 1 The branch number $\mathcal{B}$ of linear transformation P is defined by

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

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

Definition 2 A differential active s-box is defined as an s-box given a non-zero input difference. A linear active s-box is defined as an s-box given a non-zero output mask value.

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

Definition 3 For any given Δx, Δy, Γx, $\Gamma y \in \mathrm{GF}(2^m)$, the differential/linear probabilities of s_i-box: $\mathrm{GF}(2^m) \to \mathrm{GF}(2^m)$ are defined as:

$$\Pr_x[s_i(x) \oplus s_i(x \oplus \Delta x) = \Delta y] = \frac{\#\{x \in \mathrm{GF}(2^m) | s_i(x) \oplus s_i(x \oplus \Delta x) = \Delta y\}}{2^m}$$

$$\Pr_x[x \cdot \Gamma x = s_i(x) \cdot \Gamma y] = \frac{\#\{x \in \mathrm{GF}(2^m) | x \cdot \Gamma x = s_i(x) \cdot \Gamma y\}}{2^m}$$

Definition 4 Let p_s and q_s be the maximum differential/linear probabilities of all s-boxes $\{s_1, s_2, \ldots\}$.

$$p_s = \max_i \max_{\Delta x \neq 0, \Delta y} \Pr_x[s_i(x) \oplus s_i(x \oplus \Delta x) = \Delta y]$$

$$q_s = \max_i \max_{\Gamma y \neq 0, \Gamma x} (2 \Pr_x[x \cdot \Gamma x = s_i(x) \cdot \Gamma y] - 1)^2$$

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

With the above-mentioned techniques, we prove that Camellia offers immunity to these attacks by showing the upper bounds of maximum differential/linear characteristic probabilities, since Camellia is a Feistel cipher whose round function uses the S-P round function.

In the case of Camellia, the maximum differential/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/linear characteristic probabilities of Camellia reduced to 16-round without FL- and FL^{-1}-functions, respectively, we have

$$p \leq p_s^{2(2\mathcal{B}+1)} = (2^{-6})^{22} = 2^{-132} \quad \text{and} \quad q \leq q_s^{2(2\mathcal{B}+1)} = (2^{-6})^{22} = 2^{-132}$$

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 without FL- and FL^{-1}-functions. Since FL- and FL^{-1}-functions are linear for any fixed key, they do not make the average differential/linear probabilities of the cipher higher. Hence, it is proven that Camellia offers enough security against differential and linear attacks.

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 with reduced rounds. The counting algorithm is similar to that described in [M99] except following three items.

- Prepare the table for the number of active s-boxes instead of transition probability table.
- 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/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 with FL- and FL-functions has no differential/linear characteristic with probability higher than 2^{-128} (see Tables 10 and 11).

6.2 Truncated Differential Cryptanalysis

The attacks using truncated differentials were introduced by Knudsen [K95]. 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 bytewise differentials as truncated differentials [MT99].

# of rounds	1	2	3	4	5	6	7	8	9	10	11	12
Estimation based on Th. 1 and 2			2^{-12} (2)	2^{-30} (5)		2^{-42} (7)		2^{-66} (11)				2^{-96} (16)
Camellia	1 (0)	2^{-6} (1)	2^{-12} (2)	2^{-42} (7)	2^{-54} (9)	2^{-66} (11)	2^{-72} (12)	2^{-72} (12)	2^{-78} (13)	2^{-108} (18)	2^{-120} (20)	2^{-132} (22)
without FL/FL^{-1}-functions	1 (0)	2^{-6} (1)	2^{-12} (2)	2^{-42} (7)	2^{-54} (9)	2^{-66} (11)	2^{-78} (13)	2^{-90} (15)	2^{-108} (18)	2^{-126} (21)	2^{-132} (22)	

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

Table 10: Upper bounds of differential characteristic probability of Camellia

# of rounds	1	2	3	4	5	6	7	8	9	10	11	12
Estimation based on Th. 1 and 2			2^{-12} (2)	2^{-30} (5)		2^{-42} (7)		2^{-66} (11)				2^{-96} (16)
Camellia	1 (0)	2^{-6} (1)	2^{-12} (2)	2^{-36} (6)	2^{-54} (9)	2^{-66} (11)	2^{-72} (12)	2^{-72} (12)	2^{-78} (13)	2^{-102} (17)	2^{-120} (20)	2^{-132} (22)
without FL/FL^{-1}-functions	1 (0)	2^{-6} (1)	2^{-12} (2)	2^{-36} (6)	2^{-54} (9)	2^{-66} (11)	2^{-78} (13)	2^{-84} (14)	2^{-108} (18)	2^{-120} (20)	2^{-132} (22)	

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

Table 11: Upper bounds of linear characteristic probability of Camellia

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 [LMM91]. 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 [MT99]. 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 [MSAK00] breaks an 8-round variant of E2 with either IT-Function or FT-Function using 2^{94} chosen plaintexts. In [MSAK00] 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 [MT99, MSAK00]. 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 byte at the XOR of Feistel network. However, the round function of Camellia doesn't have the second s-boxes-layer, and the cancellation sometimes occurs with probability 1. As a result, more than 10-round Camellia is indistinguishable from a random permutation both with/without FL-/FL^{-1}-function layers.

Recently, Sugita et al.'s paper on truncated and impossible differential cryptanalysis of Camellia (without FL-/FL^{-1}-functions) was accepted for ASIACRYPT 2001 [SKI01]. They claim that they found two non-trivial 9-round truncated differentials (with the same input/output differential patterns), which lead to a possible attack of Camellia reduced to 11 rounds without input/output whitenings and FL-/FL^{-1}-functions. However we think it is still open how many rounds of Camellia can be attacked using the truncated differentials.

6.3 Truncated Linear Cryptanalysis

We introduce a new cryptanalysis called truncated linear cryptanalysis.

Due to the duality between differential and linear cryptanalysis, 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, more than 10-round Camellia is indistinguishable from a random permutation without FL-/FL^{-1}-function layers.

6.4 Cryptanalysis with Impossible Differential

The impossible differential means the differential which holds with probability 0, or the differential which never exists. Using such an impossible differential, it is possible to narrow down the candidates of the subkey. It is known that there is at least one 5-round impossible differential in any Feistel network with a bijective round function. Since Camellia has the Feistel network (with FL- and FL^{-1}-functions inserted between every 6 rounds) and the round function is bijective, Camellia has 5-round impossible differentials. Additionally as a recent result, Sugita et al. found a 7-round impossible differential for Camellia (without FL-/FL^{-1}-functions) [SKI01]. We expect FL- and FL^{-1}-functions make attacking Camellia using impossible differentials difficult, since the functions change differential paths depending on key values. In consequent, Camellia with full rounds will not be broken by cryptanalysis using impossible differentials.

6.5 Boomerang Attack

Boomerang attack [W99] requires 2 differentials. Let the probability of the differentials be p_Δ and p_∇. An boomerang attack that is superior than exhaustive key search requires

$$p_\Delta p_\nabla \geq 2^{-64}. \tag{6}$$

Using Table 10, there is no combination that satisfies Inequality (6) for Camellia without FL- and FL^{-1}-functions. The best boomerang probability for Camellia without FL- and FL^{-1}-functions 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 attackable rounds for Camellia without FL- and FL^{-1}-functions is bounded by much shorter than the specification of Camellia, 18, Camellia seems secure against a boomerang attack.

6.6 Higher Order Differential Attack

Higher order differential attack is generally applicable to ciphers that can be represented as Boolean polynomials of low degree. In the higher order differential attack described in [JK97, Theorem 1], the property that if the intermediate bits are represented by Boolean polynomials of degree at lea st d, the $(d+1)$-th order differential of the Boolean polynomial becomes 0 is utilized.

Degrees of Boolean polynomials of the s-boxes The functions affine (over GF(2)) equivalent to the inversion function in GF(2^8) are adopted as the s-boxes. It is known that the degree of the Boolean polynomial of every output bit of the inversion function in GF(2^8) is 7, but the degree for the s-boxes of Camellia is not trivial, since the affine functions are added at the input and output. 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.

Degrees of Boolean polynomials of the entire cipher It is expected that the degree of an intermediate bit in the encryption process increases as the data pass through many s-boxes, whose degree is 7. Therefore, we expect that higher order differential attacks fail against Camellia with full rounds. However, there is still room for further study on higher order differential attacks of Camellia, because there are other approaches for higher order differential attacks. In [KK01] Kawabata et al. shows that Camellia with 10 rounds (without FL- and FL^{-1}-functions) can be attacked faster than exhaustive search when the key size is 256-bit. The attack is applicable to 9 rounds for 192-bit keys and 8 rounds for 128-bit keys. Although the above attack is titled a "higher order differential attack", the used approach is similar to that used for the SQUARE attack.

6.7 SQUARE Attack

The SQUARE attack was proposed as a dedicated attack on SQUARE [DKR97] that exploits its byte-oriented structure. It works well for other byte-oriented ciphers such as Rijndael, Hierocrypt and Camellia. For our cryptanalysis of Camellia using the SQUARE attack, see Sect. 6.8.

The approach of SQUARE attack resembles that of higher order differential attacks: one chooses a certain complete set of plaintexts, and after some rounds of the cipher, predicts a key-independent property with probability one. The higher order differential attack of Camellia by Kawabata et al. [KK01] also takes this approach.

Another SQUARE attack by He and Qing [HQ01] on 6 rounds of Camellia was accepted for ICICS 2001. The claimed attack on 6 rounds of Camellia requires much more complexity than Kawabata and Kaneko's attack (the attack [HQ01] requires 2^{112} encryptions and the attack [KK01] $2^{22}/6$ encryptions), but fewer plaintexts (the attack [HQ01] requires 13×2^8 plaintexts and the attack [KK01] 2^{17} plaintexts).

6.8 Interpolation Attack and Linear Sum Attack

The interpolation attack proposed in [JK97] is typically applicable to attacking ciphers that use simple algebraic functions.

The principle of interpolation attack is that, roughly speaking, if the ciphertext is represented as a polynomial or rational expression of the plaintext whose number of unknown coefficients is N, the polynomial or rational expression can be constructed using N pairs of plaintexts and ciphertexts. Once the attacker constructs the polynomial or rational expression, he can encrypt any plaintext into the corresponding ciphertext or decrypt any ciphertext into the corresponding plaintext for the key without knowing the key. Since N determines the complexity and the number of pairs required for the attack, it is important to make N as large as possible. If N is so large that it is impractical for the attackers to gather N plaintext-ciphertext pairs, the cipher is secure against interpolation attack.

Linear sum attack [A00] is a generalization of the interpolation attack [JK97]. A practical algorithm that evaluates the security against linear sum attack was proposed in [A00]. We searched for linear relations between any plaintext byte and any ciphertext byte over GF(2^8) using the algorithm. Table 12 summarizes the results.

Table 12: 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

Table 12 shows that Camellia is secure against linear sum attack including interpolation attack. It also implies that Camellia is secure against SQUARE attack [DKR97] followed by [A00, Theorem 3].

6.9 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.

6.10 Slide Attack

In [BW99, BW00] the slide attacks were introduced, based on earlier work in [B94, K93]. 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 (See Section 6.11).

6.11 Related-key Attack

We are convinced that the key schedule of Camellia makes related-key attacks [B94, KSW96] very difficult. In these attacks, an attacker must be able to get encryptions using several related keys. If the relation between, say, two keys, is known then if the corresponding relations between the subkeys can be predetermined, it might become possible to predict how the keys would encrypt a pair of different plaintexts. 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.

6.12 Statistical Tests

Most of statistical characteristics depends on the differential attack and other cryptanalytic attacks. For example, it is frequently discussed how many ciphertext bits are complemented when one plaintext bit is complemented. According to the definition and the property of the differential distribution table, the resistance to differential attacks implies that the number of complemented bits is about a half. Of course, we may find a statistical weakness, if we have enough computational resource. However, none in the world has an efficient resource to compute such a statistical measure for a 128-bit block cipher.

Note that the following. It is frequently tested for a round function, because of the limited computational resource. However, we think that it is not significant, because we can construct a cipher that does not show good statistical properties for the round function but shows good statistical properties for a cipher and we can also construct a cipher that shows good statistical properties for the round function but does not show good statistical properties for a cipher.

In the CRYPTREC Report 2000 [C01], it is reported that the avalanche-effect evaluation on Camellia was held and that they found some points which deviate from the expected value in the round function, but no particular characteristics in the data-randomizing part after the 4th round.

6.13 Implementation Attacks

It is well known that a poor implementation can leak information by timing attacks [K96] or power analysis attacks [KJJ99]. Using the classification proposed in [DR99], Camellia is in the group of "favorable" algorithms, since it uses only logical operations and table-lookups and fixed rotations.

On the other hand, Chari et al. [CJRR99] claims that all AES candidates are susceptible to power analysis attacks. As these two papers contradict with each other, how to resist against power analysis attacks is not known, since study on power analysis attacks has just begun. We think that Camellia should be protected by the hardware techniques and should not be evaluated by the security directly derived from the specification, considering the current art. We hope that the study on implementation attacks will be progressed in the near future.

6.14 Brute Force Attacks

Most brute force attacks are applicable to any deterministic block cipher, and the corresponding complexity depends on only the block size or key size*, regardless of its design. Camellia has a block size of 128-bit and allows for the three key sizes of 128-, 192-, and 256-bit. In the discussions below, k denotes the key size in bits.

Exhaustive key search. In exhaustive key search, if an attacker gets one pair of plaintext and ciphertext encrypted in ECB mode, he can find the correct key by encrypting the plaintext with all 2^k possible keys.

A weakness in the key scheduling of the cipher can help improve the efficiency of exhaustive key search attack [K94], but we have not found such a weakness in Camellia. The complexity of the exhaustive key search is estimated to be about 2^{k-1} encryptions on average. Thus, the required complexity for exhaustive key search is 2^{127}, 2^{191}, and 2^{255} encryptions for Camellia with 128-, 192-, and 256-bit keys, respectively. Therefore, Camellia's security against exhaustive key search is adequate.

Time-memory trade-off attack. There are some words that are often used in plaintexts. If an attacker encrypts such a plaintext block using 2^k keys and store them in space for 2^k ciphertexts, then after he gets the corresponding ciphertext, he only has to look it up to find the corresponding key. This attack is called table attack. In this attack, after 2^k encryption is done, the attack complexity is much smaller than is true for exhaustive key search.

Time-memory trade-off attack [H80, KM96] can drastically reduce both time complexity on intercepted ciphertexts of exhaustive key search and space complexity of table attack. However, both attacks require precomputation equivalent to the time complexity of exhaustive key search. The key sizes supported by Camellia are long enough for security against exhaustive key search by today's technology.

Dictionary attack. In dictionary attack, an attacker collects plaintext-ciphertext pairs under the same key and put them in a "dictionary". When the attacker can see only a ciphertext encrypted by the key, he can check if it is in the dictionary. If it is, he has already the plaintext. Since the block size of Camellia is 128 bits, dictionary attack would require the space for 2^{128} different plaintext blocks to allow the attackers to encrypt or decrypt arbitrary messages under an unknown key. The success probability depends on the space for the dictionary, and as the block size is larger, the required space to achieve the same success probability increases exponentially. The 128-bit block cipher Camellia has enough security against this attack.

Matching ciphertext attack. In matching ciphertext attack [K98, Theorem 2], when about the square root of all ciphertexts are available identical ciphertext blocks can be expected with probability more than $\frac{1}{2}$ by the "birthday paradox" for some modes of operations such as ECB, CBC, and CFB modes. Then, valuable information about the plaintexts can be derived. Note

*Strictly speaking, the computation time required for the attack depends on the performance of the block cipher. However, the performance only affects the encryption time and only changes the time complexity by negligible factor.

that this attack is independent of the key size. Since the block size of Camellia is 128 bits, the threat to this attack is small, if encryption of as many as 2^{64} blocks under the same key is not performed.

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.

The performances shown in this paper leave room for further optimizations. The latest performance results will be posted on the Camellia home page: `http://info.isl.ntt.co.jp/camellia/`.

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

[A00] K. Aoki. Practical Evaluation of Security against Generalized Interpolation Attack. *IEICE Transactions Fundamentals of Electronics, Communications and Computer Sciences (Japan)*, Vol. E83-A, No. 1, pp. 33–38, 2000. (A preliminary version was presented at SAC'99).

[ABK98] R. Anderson, E. Biham, and L. Knudsen. Serpent: A Flexible Block Cipher With Maximum Assurance. In *The First AES Candidate Conference*, 1998.

[AIK+00a] K. Aoki, T. Ichikawa, M. Kanda, M. Matsui, S. Moriai, J. Nakajima, and T. Tokita. Implementations of the 128-bit block cipher – *Camellia* –. Technical Report ISEC2000-73, The Institute of Electronics, Information and Communication Engineers, 2000. (in Japanese).

[AIK+00b] K. Aoki, T. Ichikawa, M. Kanda, M. Matsui, S. Moriai, J. Nakajima, and T. Tokita. *Camellia*: A 128-Bit Block Cipher Suitable for Multiple Platforms – Extended Abstract –. In *First NESSIE Workshop*, 2000.

[AU00] K. Aoki and H. Ueda. Optimized Software Implementations of E2. *IEICE Transactions Fundamentals of Electronics, Communications and Computer Sciences (Japan)*, Vol. E83-A, No. 1, pp. 101–105, 2000. (The full paper is available on `http://info.isl.ntt.co.% linebreak[3]jp/e2/RelDocs/`).

[B94] E. Biham. New Types of Cryptanalytic Attacks Using Related Keys. *Journal of Cryptology*, Vol. 7, No. 4, pp. 229–246, 1994. (The extended abstract was appeared at EUROCRYPT'93).

[BS93] E. Biham and A. Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer-Verlag, Berlin, Heidelberg, New York, 1993.

[BW99] A. Biryukov and D. Wagner. Slide Attacks. In L. Knudsen, editor, *Fast Software Encryption — 6th International Workshop, FSE'99*, Volume 1636 of *Lecture Notes in Computer Science*, pp. 245–259, Berlin, Heidelberg, New York, 1999. Springer-Verlag.

[BW00] A. Biryukov and D. Wagner. Advanced Slide Attacks. In S. Vaudenay, editor, *Advances in Cryptology — EUROCRYPT2000*, Volume 1807 of *Lecture Notes in Computer Science*, pp. 589–606, Berlin, Heidelberg, New York, 2000. Springer-Verlag.

[C98] Compaq Computer Corporation. *Alpha Architecture Handbook (Version 4)*, 1998. (You can download the manual from Compaq's technical documentation library: `http://www.support.compaq.com/alpha-tools/documentation/current/chip-docs.html`).

[C01] CRYPTREC. CRYPTREC Report 2000, April 2001.

[CJRR99] S. Chari, C. Jutla, J. R. Rao, and P. Rohatgi. A Cautionary Note Regarding Evaluation of AES Candidates on Smart-Cards. In *Second Advanced Encryption Standard Candidate Conference*, pp. 133–147, Hotel Quirinale, Rome, Italy, 1999. Information Technology Laboratory, National Institute of Standards and Technology.

[CV95] F. Chabaud and S. Vaudenay. Links Between Differential and Linear Cryptanalysis. In A. D. Santis, editor, *Advances in Cryptology — EUROCRYPT'94*, Volume 950 of *Lecture Notes in Computer Science*, pp. 356–365. Springer-Verlag, Berlin, Heidelberg, New York, 1995.

[DKR97] J. Daemen, L. R. Knudsen, and V. Rijmen. The Block Cipher SQUARE. In E. Biham, editor, *Fast Software Encryption — 4th International Workshop, FSE'97*, Volume 1267 of *Lecture Notes in Computer Science*, pp. 54–68, Berlin, Heidelberg, New York, 1997. Springer-Verlag.

[DR98] J. Daemen and V. Rijmen. *AES Proposal: Rijndael*, 1998. (`http://www.esat.kuleuven.ac.be/~rijmen/rijndael/`).

[DR99] J. Daemen and V. Rijmen. Resistance Against Implementation Attacks. A Comparative Study of the AES Proposals. In *The Second AES Candidate Conference*, 1999.

[F00] A. Fog. *How to optimize for the Pentium microprocessors*, 2000. (`http://www.agner.org/assem/`).

[H80] M. Hellman. A Cryptanalytic time-memory trade-off. *IEEE Transactions on Information Theory*, Vol. IT-26, No. 4, pp. 401–406, 1980.

[HQ01] Y. He and S. Qing. Square Attack on Reduced Camellia Cipher. submitted to the 3rd International Conference on Information and Communications Security (ICICS 2001), 2001.

[I99] Intel Corporation. *Intel Architecture Software Developer's Manual (Volume 2: Instruction Set Reference)*, 1999. (You can download the manual from Intel's developer site: `http://developer.intel.com/`).

[ISKM01] T. Ichikawa, T. Sorimachi, T. Kasuya, and M. Matsui. On the criteria of hardware evaluation of block ciphers (1). Technical Report ISEC2001-53, The Institute of Electronics, Information and Communication Engineers, 2001. (in Japanese).

[JK97] T. Jakobsen and L. R. Knudsen. The Interpolation Attack on Block Cipher. In E. Biham, editor, *Fast Software Encryption — 4th International Workshop, FSE'97*, Volume 1267 of *Lecture Notes in Computer Science*, pp. 28–40, Berlin, Heidelberg, New York, 1997. Springer-Verlag.

[K93] L. R. Knudsen. Cryptanalysis of LOKI91. In J. Seberry and Y. Zheng, editors, *Advances in Cryptology — AUSCRYPT'92*, Volume 718 of *Lecture Notes in Computer Science*, pp. 196–208. Springer-Verlag, Berlin, Heidelberg, New York, 1993.

[K94] L. R. Knudsen. Practically secure Feistel ciphers. In R. Anderson, editor, *Fast Software Encryption 1993 — Cambridge Security Workshop (FSE1)*, Volume 809 of *Lecture Notes in Computer Science*, pp. 211–221, Berlin, Heidelberg, New York, 1994. Springer-Verlag.

[K95] L. R. Knudsen. Truncated and Higher Order Differentials. In B. Preneel, editor, *Fast Software Encryption — Second International Workshop*, Volume 1008 of *Lecture Notes in Computer Science*, pp. 196–211. Springer-Verlag, Berlin, Heidelberg, New York, 1995.

[K96] P. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In N. Koblitz, editor, *Advances in Cryptology — CRYPTO'96*, Volume 1109 of *Lecture Notes in Computer Science*, pp. 104–113. Springer-Verlag, Berlin, Heidelberg, New York, 1996.

[K98] L. R. Knudsen. Block Ciphers — A Survey. In B. Preneel and V. Rijmen, editors, *State of the Art in Applied Cryptography*, Volume 1528 of *Lecture Notes in Computer Science*, pp. 18–48, Berlin, Heidelberg, New York, 1998. Springer-Verlag.

[K00] M. Kanda. Practical Security Evaluation against Differential and Linear Attacks for Feistel Ciphers with SPN Round Function. In *SAC2000, Seventh Annual Workshop on Selected Areas in Cryptography, 14-15 August 2000, Workshop Record*, 2000.

[KJJ99] P. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In M. Wiener, editor, *Advances in Cryptology — CRYPTO'99*, Volume 1666 of *Lecture Notes in Computer Science*, pp. 388–397. Springer-Verlag, Berlin, Heidelberg, New York, 1999.

[KK01] T. Kawabata and T. Kaneko. A Study on Higher Order Differential Attack of Camellia. In *Second NESSIE Workshop*, 2001. (This paper is based on T. Kawabata, Y. Ohgaki and T. Kaneko, "A Study on Strength of Camellia against Higher Order Differential Attack," (in Japanese), Technical report of IEICE, ISEC2001-9, pp.55–62, The Institute of Electronics, Information and Communication Engineers, 2001.).

[KM96] K. Kusuda and T. Matsumoto. Optimization of Time-Memory Trade-Off Cryptanalysis and Its Application to DES, FEAL-32, and Skipjack. *IEICE Transactions Fundamentals of Electronics, Communications and Computer Sciences (Japan)*, Vol. E79-A, No. 1, pp. 35–48, 1996.

[KMA+98] M. Kanda, S. Moriai, K. Aoki, H. Ueda, M. Ohkubo, Y. Takashima, K. Ohta, and T. Matsumoto. A New 128-bit Block Cipher ***E2***. Technical Report ISEC98-12, The Institute of Electronics, Information and Communication Engineers, 1998. (in Japanese).

[KSW96] J. Kelsey, B. Schneier, and D. Wagner. Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In N. Koblitz, editor, *Advances in Cryptology — CRYPTO'96*, Volume 1109 of *Lecture Notes in Computer Science*, pp. 237–251. Springer-Verlag, Berlin, Heidelberg, New York, 1996.

[KTM+99] M. Kanda, Y. Takashima, T. Matsumoto, K. Aoki, and K. Ohta. A Strategy for Constructing Fast Round Functions with Practical Security against Differential and Linear Cryptanalysis. In S. Tavares and H. Meijer, editors, *Selected Areas in Cryptography — 5th Annual International Workshop, SAC'98*, Volume 1556 of *Lecture Notes in Computer Science*, pp. 264–279, Berlin, Heidelberg, New York, 1999. Springer-Verlag.

[LMM91] X. Lai, J. L. Massey, and S. Murphy. Markov Ciphers and Differential Cryptanalysis. In D. W. Davies, editor, *Advances in Cryptology — EUROCRYPT'91*, Volume 547 of *Lecture Notes in Computer Science*, pp. 17–38. Springer-Verlag, Berlin, Heidelberg, New York, 1991.

[M94] M. Matsui. Linear Cryptanalysis Method for DES Cipher. In T. Helleseth, editor, *Advances in Cryptology — EUROCRYPT'93*, Volume 765 of *Lecture Notes in Computer Science*, pp. 386–397. Springer-Verlag, Berlin, Heidelberg, New York, 1994. (A preliminary version written in Japanese was presented at SCIS93-3C).

[M95] M. Matsui. On Correlation Between the Order of S-boxes and the Strength of DES. In A. D. Santis, editor, *Advances in Cryptology — EUROCRYPT'94*, Volume 950 of *Lecture Notes in Computer Science*, pp. 366–375. Springer-Verlag, Berlin, Heidelberg, New York, 1995.

[M97] M. Matsui. New Block Encryption Algorithm MISTY. In E. Biham, editor, *Fast Software Encryption — 4th International Workshop, FSE'97*, Volume 1267 of *Lecture Notes in Computer Science*, pp. 54–68, Berlin, Heidelberg, New York, 1997. Springer-Verlag. (A preliminary version written in Japanese was presented at ISEC96-11).

[M99] M. Matsui. Differential Path Search of the Block Cipher E2. Technical Report ISEC99-19, The Institute of Electronics, Information and Communication Engineers, 1999. (in Japanese).

[MIYY88] M. Matsui, T. Inoue, A. Yamagishi, and H. Yoshida. A note on calculation circuits over GF(2^{2n}). Technical Report IT88-14, The Institute of Electronics, Information and Communication Engineers, 1988. (in Japanese).

[MSAK00] S. Moriai, M. Sugita, K. Aoki, and M. Kanda. Security of E2 against Truncated Differential Cryptanalysis. In H. Heys and C. Adams, editors, *Selected Areas in Cryptography — 6th Annual International Workshop, SAC'99*, Volume 1758 of *Lecture Notes in Computer Science*, pp. 106–117, Berlin, Heidelberg, New York, 2000. Springer-Verlag.

[MT99] M. Matsui and T. Tokita. Cryptanalysis of a Reduced Version of the Block Cipher E2. In L. Knudsen, editor, *Fast Software Encryption — 6th International Workshop, FSE'99*, Volume 1636 of *Lecture Notes in Computer Science*, pp. 71–80, Berlin, Heidelberg, New York, 1999. Springer-Verlag. (Japanese version was presented at SCIS99.).

[RDP+96] V. Rijmen, J. Daemen, B. Preneel, A. Bosselaers, and E. De Win. The Cipher SHARK. In D. Gollmann, editor, *Fast Software Encryption — Third International Workshop*, Volume 1039 of *Lecture Notes in Computer Science*, pp. 99–111. Springer-Verlag, Berlin, Heidelberg, New York, 1996.

[SKI01] M. Sugita, K. Kobara, and H. Imai. Security of Reduced Version of the Block Cipher Camellia against Truncated and Impossible Differential Cryptanalysis. submitted to ASIACRYPT 2001, 2001.

[W99] D. Wagner. The Boomerang Attack. In L. R. Knudsen, editor, *Fast Software Encryption — 6th International Workshop, FSE'99*, Volume 1636 of *Lecture Notes in Computer Science*, pp. 156–170, Berlin, Heidelberg, New York, 1999. Springer-Verlag.

[Y01a] C.-H. Yang. Performance Evaluation of AES/DES/Camellia on the 6805 and H8/300 CPUs. In *Proceedings of the 2001 Symposium on Cryptography and Information Security*, Volume II of *SCIS2001*, pp. 727–730, Oiso, Japan, 2001. Technical Group on Information Security (IEICE).

[Y01b] C.-H. Yang. Supplementary information for C.H. Yang SCIS 2001 paper. http://www.geocities.com/chyang00/SCIS2001, 2001.

A History

Ver 2.0 (September 26, 2001)

- Abstract was renewed with the latest performance figures.
- Section 1, the paragraph of "Future developments" was renewed based on the current status. The title was also changed into "Standardization activities".
- Section 3 was renewed with the latest performance figures.
- In Section 4.2.7, the equation to calculate Eq.(3) using only four tables, SP_1, SP_2, SP_3, SP_4, was corrected.
- Section 5 was renewed by adding the latest information on hardware evaluations.
- In Section 6.1 (Differential and Linear Cryptanalysis), an erratum in Table10 "Upper bounds of differential characteristic probability of Camellia" (in the row of "without FL/FL^{-1}-functions") was fixed.
- Section 6.2 (Truncated Differential Cryptanalysis) was renewed by adding the recent result.
- Section 6.4 (Cryptanalysis with Impossible Differential) was renewed by adding the recent result. An erratum was also fixed: "more than 6 rounds" $\rightarrow$ "more than 5 rounds"
- Section 6.6 (Higher Order Differential Attack) was renewed based on the recent result.
- Section 6.7 (SQUARE Attack) was added.
- Section 6.12 (Statistical Tests) was renewed by adding more information.