

The CiliPadi Family of Lightweight Authenticated Encryption

Version 1.0

Muhammad Reza Z'aba¹, Norziana Jamil², Mohd Saufy Rohmad³,
Hazlin Abdul Rani⁴, and Solahuddin Shamsuddin⁴

¹Faculty of Computer Science and Information Technology, University of Malaya
`reza.zaba@um.edu.my`

²College of Computing and Informatics, Universiti Tenaga Nasional
`norziana@uniten.edu.my`

³Faculty of Electrical Engineering, Universiti Teknologi MARA `saufy@uitm.edu.my`

⁴CyberSecurity Malaysia `hazlin@cybersecurity.my` `solahuddin@cybersecurity.my`

March 29, 2019

Contact Information

For any correspondence, please contact:

CyberSecurity Malaysia

Level 9 Tower 1
Menara Cyber Axis
Jalan Impact
63000 Cyberjaya, Selangor
Malaysia

Tel: +603 8800 7999

Fax: +603 8008 7000

E-mail: cilipadi@cybersecurity.my

Acknowledgement

We thank Teh Je Sen and Yasir Amer Abbas for providing additional security analysis and implementation results.

The CiliPadi Family of Lightweight Authentication Encryption
Version 1.0: March 29, 2019

Contents

1	Introduction	1
2	Preliminaries	1
2.1	Notations	1
2.2	Mode of Operation	1
3	Specification	2
3.1	Parameters	2
3.2	Initialization Phase	4
3.3	Padding	4
3.4	Associated Data Authentication Phase	4
3.5	Message Encryption Phase	5
3.6	Message Decryption Phase	5
3.7	Finalization Phase	5
3.8	The Permutation Function P	5
3.9	The F function	6
3.9.1	AddConstants (AC)	7
3.9.2	SubCells (SC)	7
3.9.3	ShiftRows (SR)	7
3.9.4	MixColumnsSerial (MCS)	8
4	Design Rationale	8
4.1	Key Lengths	8
4.2	Sponge	8
4.3	Permutation	9
5	Performance Analysis	9
5.1	Implementation Results	11
6	Security Analysis	12
6.1	Differential Cryptanalysis	12
6.1.1	Preliminaries	13
6.1.2	P as a Random Permutation	14
6.1.3	Collision-Producing Differentials of CiliPadi	17
6.1.4	Practical Security Bounds	18
6.2	Full Bit Diffusion	19
6.3	Extension to Linear Cryptanalysis	19
7	Strengths and Weaknesses	19
7.1	Strengths	20
7.2	Weaknesses	20
A	Test Vectors	23

1 Introduction

This document describes a family of lightweight authenticated encryption with associated data (AEAD) called CiliPadi¹. There are four flavours of CiliPadi: Mild, Medium, Hot and ExtraHot. The primary member of CiliPadi is CiliPadi-Mild.

2 Preliminaries

This section introduces the notation and other preliminaries as basis to understand subsequent sections.

2.1 Notations

The notation used in this paper is given in Table 1.

	Description
K, N, T	The secret key, nonce, and tag, respectively
$X\ Y$	The concatenation of bit strings X and Y
A	The associated data where $A = A_1\ \dots\ A_s$
M	The message where $M = M_1\ \dots\ M_t$
C	The ciphertext where $C = C_1\ \dots\ C_t\ T$
$ X $	The length of X in bits
S	The internal state where $S = S_r\ S_c$, $r = S_r $ and $c = S_c $
n	The length of the internal state S , i.e. $n = S = r + c$
$0_2^x, 1_2^x$	An all-zeros and all-ones binary string of x bits, respectively
$[X]^i$	The first i bits (or leftmost bits) of X
$(X_1\ \dots\ X_x) \stackrel{y}{\leftarrow} X$	The parsing of the bit string X into x equally-sized y -bit strings

Table 1: Notation

A number written in typewriter font is always treated as a 4-bit hexadecimal value, i.e. $0, 1, \dots, f$. If the value is subscripted with ‘2’, as shown in Table 1 then it is treated as a 1-bit value, which applies only for 0 and 1.

2.2 Mode of Operation

The CiliPadi $[n, r, a, b]$ mode of operation is based on the MonkeyDuplex construction [3, 6] and depicted in Figure 1. It consists of four phases: initialization, associated data authentication, message encryption/decryption, and finalization. The state length is n bits initialized with the value of the secret key K and nonce N . The bitrate is r bits and the capacity is $c = n - r$ bits. The permutation for the initialization and finalization phases has a rounds while the

¹*Cili padi* is a Malay word for the bird’s eye chili. The name is chosen due to the Malay proverb, *kecil-kecil cili padi* which means tiny but powerful. In our context, CiliPadi has a small footprint but secure for lightweight applications.




Figure 1: CiliPadi mode of operation

permutation for the associated data and message encryption/decryption phases has b rounds where $b > a$.

3 Specification

This section formally describes the specification of CiliPadi.

3.1 Parameters

The CiliPadi $[n, r, a, b]$ family of authenticated encryption (AE) scheme consists of configurable parameters. For the purpose of evaluation, we propose four flavours of CiliPadi which is listed in Table 2 according to increasing level of security. They are CiliPadi-Mild, CiliPadi-Medium, CiliPadi-Hot and CiliPadi-ExtraHot. The lengths stated in the table are all in bits.

CiliPadi-	Algorithm $[n, r, a, b]$	Key	Length of			No. of rounds	
			Nonce	Tag	Block	P_n^a	P_n^b
Mild	$[256, 64, 18, 16]$	128	128	64	64	18	16
Medium	$[256, 96, 20, 18]$	128	128	96	96	20	18
Hot	$[384, 96, 18, 16]$	256	128	96	96	18	16
ExtraHot	$[384, 128, 20, 18]$	256	128	128	128	20	18

Table 2: CiliPadi parameters where the primary member is CiliPadi-Mild

Formally, the CiliPadi family of AE accepts a k -bit secret key K and a 128-bit nonce N . These values become the initial value of the n -bit internal state $S = K || N$. The state is then updated by the permutation P_n^a . If the sr -bit associated data $A = A_1 || \dots || A_s$ is non-empty, it will be subsequently processed, along with the internal state, by the associated data authentication phase. Encryption takes the padded message $M = M_1 || \dots || M_t$ and outputs the ciphertext $C = C_1 || \dots || C_t$ and tag T where $|M_i| = |C_i| = T = r$ bits. Decryption takes the ciphertext C and tag T and outputs the original message M if and only if C is authentic, else it outputs \perp . The components and high-level overview of CiliPadi are given in Figures 2 and 3, respectively. The descriptions are provided in the following sections.

<hr/> Proc Init(K, N) <hr/> 1 $S \leftarrow K \ N$ 2 return S <hr/>	<hr/> Proc AD(S, A) <hr/> 1 for $i = 1, \dots, s$ do 2 $S \leftarrow P_n^b((S_r \oplus A_i) \ S_c)$ 3 end 4 $S \leftarrow (S_r \ S_c \oplus (0_2^{c-1} \ 1))$ 5 return S <hr/>
<hr/> Proc Finalization(S) <hr/> 1 $S \leftarrow P_n^a(S)$ 2 $T \leftarrow [S]^k \oplus K$ 3 return T <hr/>	<hr/>
<hr/> Proc MEnc(S, M) <hr/> 1 for $i = 1, \dots, t - 1$ do 2 $C_i \leftarrow S_r \oplus M_i$ 3 $S \leftarrow P_n^b(C_i \ S_c)$ 4 end 5 $C_t \leftarrow S_r \oplus M_t$ 6 $S \leftarrow (C_t \ S_c)$ 7 return (S, C) <hr/>	<hr/> Proc MDec(S, C) <hr/> 1 for $i = 1, \dots, t - 1$ do 2 $M_i \leftarrow S_r \oplus C_i$ 3 $S \leftarrow P_n^b(C_i \ S_c)$ 4 end 5 $M_t \leftarrow S_r \oplus C_t$ 6 $S \leftarrow (C_t \ S_c)$ 7 return (S, M) <hr/>
<hr/> Proc $P_{256}^r(S)$ <hr/> 1 $(X_1 \ \dots \ X_4) \xleftarrow{64} S$ 2 for $i = 1, \dots, r$ do 3 $Y_1 \leftarrow F_1(X_1) \oplus X_2$ 4 $Y_2 \leftarrow X_3$ 5 $Y_3 \leftarrow F_2(X_3) \oplus X_4$ 6 $Y_4 \leftarrow X_1$ 7 $X \leftarrow Y$ 8 end 9 $S \leftarrow X$ 10 return S <hr/>	<hr/> Proc $P_{384}^r(S)$ <hr/> 1 $(X_1 \ \dots \ X_6) \xleftarrow{64} S$ 2 for $i = 1, \dots, r$ do 3 $Y_1 \leftarrow F_1(X_1) \oplus X_2$ 4 $Y_2 \leftarrow X_3$ 5 $Y_3 \leftarrow F_2(X_5) \oplus X_6$ 6 $Y_4 \leftarrow X_1$ 7 $Y_5 \leftarrow F_3(X_3) \oplus X_4$ 8 $Y_6 \leftarrow X_5$ 9 $X \leftarrow Y$ 10 end 11 $S \leftarrow X$ 12 return S <hr/>
<hr/> Proc $F_i^i(X)$ <hr/> 1 $(x_1 \ \dots \ x_4) \xleftarrow{16} X$ 2 $(w_1 \ w_2) \xleftarrow{2} l$ 3 $(z_1 \ z_2) \xleftarrow{3} rc[i]$ 4 $x_1 \leftarrow (0_2^2 \ w_1 \ 0_2 \ z_1 \ 0_2^8)$ 5 $x_2 \leftarrow (0_2^2 \ w_2 \ 0_2 \ z_2 \ 0_2^8)$ 6 $x_3 \leftarrow (2 \ 0_2 \ z_1 \ 0_2^8)$ 7 $x_4 \leftarrow (3 \ 0_2 \ z_2 \ 0_2^8)$ 8 $RC \leftarrow (x_1 \ x_2 \ x_3 \ x_4)$ 9 $X \leftarrow \text{LED1r}(X, RC)$ 10 $X \leftarrow \text{LED1r}(X, 0_2^{64})$ 11 return X <hr/>	<hr/> Proc LED1r(X, RC) <hr/> 1 $X \leftarrow \text{AC}(X, RC)$ 2 $X \leftarrow \text{SC}(X)$ 3 $X \leftarrow \text{SR}(X)$ 4 $X \leftarrow \text{MCS}(X)$ 5 return X <hr/>

Figure 2: Components of CiliPadi

Proc Encrypt(K, N, A, M)	Proc Decrypt(K, N, A, C, T)
<ol style="list-style-type: none"> 1 $S \leftarrow \text{Init}(K, N)$ 2 $S \leftarrow \text{AD}(S, A)$ 3 $(S, C) \leftarrow \text{MEnc}(S, M)$ 4 $T \leftarrow \text{Finalization}(S)$ 5 return (C, T) 	<ol style="list-style-type: none"> 1 $S \leftarrow \text{Init}(K, N)$ 2 $S \leftarrow \text{AD}(S, A)$ 3 $(S, M) \leftarrow \text{MDec}(S, C)$ 4 $T' \leftarrow \text{Finalization}(S)$ 5 if $(T = T')$ then <li style="padding-left: 2em;">6 return M 7 else <li style="padding-left: 2em;">8 return \perp 9 end

Figure 3: High-level overview of the encryption and decryption of CiliPadi

3.2 Initialization Phase

The n -bit state S is initialized with the value of the k -bit key followed by the 128-bit nonce N . The internal state S is initialized as

$$S = K \| N.$$

Note that the nonce must never be repeated to encrypt different messages using the same secret key. The internal state can also be viewed as the concatenation of the r -bit rate S_r and c -bit capacity S_c parts, i.e. $S = S_r \| S_c$. The state is then processed by the n -bit permutation P_n^a , which is described later. The output of this phase is fed to the associated data authentication phase, if the associated data is non-empty.

3.3 Padding

Both the associated data and message blocks are individually padded only if its length is not a multiple of r bits. Padding is performed by adding a bit 1, and then as many zero bits as necessary until the padded data is in multiple of r bits. If the length of the last block is $r - 1$ bits, then only bit 1 is added.

3.4 Associated Data Authentication Phase

If the associated data $A = A_1 \| \dots \| A_s$ is non-empty, then A_1 is XORed with the inner state S_r . The state S is then updated by the permutation P_n^b . This process is repeated for A_i ($i = 1, \dots, s$) until all associated data blocks are processed:

$$S \leftarrow P_n^b((S_r \oplus A_i) \| S_c).$$

After the last associated data is processed, the outer state S_c is XORed with the binary string $0_2^{c-1} \| 1_2$ to denote the completion of the associated data phase:

$$S \leftarrow (S_r \| S_c \oplus (0_2^{c-1} \| 1_2)).$$

The output of this phase is fed to either the message encryption or decryption phase, described next.

3.5 Message Encryption Phase

There are two main inputs for this phase. The first comes from either the initialization (if the associated data is empty) or associated data authentication phase. The second input comes from the padded message $M = M_1 \parallel \dots \parallel M_t$. The current inner state S_r is first XORed with the first message block M_1 to produce the ciphertext block C_1 . If there are more message block available, then this process is repeated until all message blocks are processed, except for the last block where the permutation P_n^b is not applied:

$$S \leftarrow \begin{cases} P_n^b((S_r \oplus M_i) \parallel S_c), & \text{for } i = 1, \dots, t-1, \\ (S_r \oplus M_i) \parallel S_c, & \text{for } i = t. \end{cases}$$

The i -th ciphertext block is extracted from the current state prior to the execution of P_n^b as $C_i = S_r \oplus M_i$. The output of this phase is fed into the finalization phase.

3.6 Message Decryption Phase

The decryption phase is almost identical to encryption. It receives two inputs. The first is the outcome of either the initialization or associated data authentication phase. The second input is the ciphertext $C = C_1 \parallel \dots \parallel C_t$. The first ciphertext block assumes the value of S_r and then the state S is updated by P_n^b . This process is repeated until the second last ciphertext block. The last ciphertext block is not processed by P_n^b .

$$S \leftarrow \begin{cases} P_n^b(C_i \parallel S_c), & \text{for } i = 1, \dots, t-1, \\ C_i \parallel S_c, & \text{for } i = t. \end{cases}$$

The corresponding message block is only released when the tag is verified, i.e. only after successful execution of the finalization phase. The i -th message block is extracted from the current state prior to the execution of the permutation P_n^b , i.e. $M_i = S_r \oplus C_i$.

3.7 Finalization Phase

This phase updates the internal state S to output a single r -bit tag. After the state has been processed by P_n^a , S_r is XORed with the first r bits of the key K . The tag T is the result of this XOR as follows.

$$\begin{aligned} S &\leftarrow P_n^a(S_r \parallel S_c) \\ T &\leftarrow S_r \oplus [K]^r \end{aligned}$$

When decrypting ciphertext, the original message M will only be released if the computed tag above matches the one supplied by the sending party.

3.8 The Permutation Function P

The permutation function P makes use of an unkeyed 2-round² of the lightweight block cipher LED [19] as the round- and line- dependent F -function in a Type-II

²This refers to a full 2 rounds where no operation is omitted in the last (second) round.




Figure 4: Type-II GFN employed in CiliPadi for $d = 4$ and $d = 6$

generalized feistel network (GFN) [24]. We refer a Type-II GFN that accepts d input sub-blocks as a d -line Type-II GFN. For CiliPadi, d is an even number and each line is of length n/d bits. There are $d/2$ F functions in each round that accepts input from odd-numbered lines. Let $X_1 \parallel \dots \parallel X_d$ denote the input lines. They are updated by the F -function in the i -th round as follows.

$$\begin{aligned} X_j &\leftarrow X_j & \text{for } j = 1, 3, \dots, d-1 \\ X_j &\leftarrow X_j \oplus F_{j/2}^i(X_{j-1}) & \text{for } j = 2, 4, \dots, d \end{aligned}$$

After the above transformations, the lines are shuffled by the permutation function π before being used as input to the next round. For instance, the shuffle $\pi = \{2, 3, 4, 1\}$ means that the first input line is mapped to the second output line, the second input line to the third output line, and so forth. The same shuffle π is used in both the message encryption and decryption phases. For CiliPadi, the shuffling³ used are given in Table 3 and depicted in Figure 4 for $d = 4$ and $d = 6$.

Input length n (in bits)	Number of lines d	Shuffle π
128	2	$\{2, 1\}$
256	4	$\{4, 1, 2, 3\}$
384	6	$\{4, 1, 2, 5, 6, 3\}$
512	8	$\{4, 1, 2, 5, 8, 3, 6, 7\}$

Table 3: Shuffling used in the Type-II GFN

3.9 The F function

As mentioned earlier, the round- and line- dependent F function is an unkeyed 2-round of the LED [19] block cipher where no operation is omitted in the last (i.e.

³Note that in [23], the index starts with 0, ours start with 1.




Figure 5: A single round of LED

second) LED round. A single LED round⁴ consists of the following four operations, depicted in Figure 5, applied in sequence to the 64-bit input: AddConstants, SubCells, ShiftRows and MixColumnsSerial. The input to LED is 64 bits partitioned into 16 4-bit cells. Let $x = x_1 \parallel \dots \parallel x_{16}$ denote this input which can be depicted as a 4×4 matrix which is entered row-wise as follows.

$$\begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ x_5 & x_6 & x_7 & x_8 \\ x_9 & x_{10} & x_{11} & x_{12} \\ x_{13} & x_{14} & x_{15} & x_{16} \end{bmatrix}$$

In AES, the input is entered column-wise.

3.9.1 AddConstants (AC)

Initialize a 6-bit round constant $rc_1 \parallel \dots \parallel rc_6$ set to all zeros. In each round of the permutation, the constant is updated by shifting its value 1 bit to the left. The new value for rc_6 is taken as $rc_1 \oplus rc_2 \oplus 1$. We also add the F -function number encoded as a 4-bit value $l = l_1 \parallel \dots \parallel l_4$ as one of the parameters in the constant value. This is to ensure that every F -function has different constant value.

$$\begin{bmatrix} l_1 \parallel l_2 & rc_1 \parallel rc_2 \parallel rc_3 & 0 & 0 \\ l_3 \parallel l_4 & rc_4 \parallel rc_5 \parallel rc_6 & 0 & 0 \\ 2 & rc_1 \parallel rc_2 \parallel rc_3 & 0 & 0 \\ 3 & rc_4 \parallel rc_5 \parallel rc_6 & 0 & 0 \end{bmatrix}$$

The above matrix is XORed to the current value of the state in the first round of LED in the F -function. The constants for the second round of LED are set to all-zeros. The complete values of the round constants are given in Table 4.

3.9.2 SubCells (SC)

This operation substitutes the current value of each 4-bit cell with another 4-bit value, using the s-box of PRESENT [9] given in Table 5.

3.9.3 ShiftRows (SR)

This operation rotates the second, third and fourth row of the state matrix by one, two and three cells to the left.

⁴We are not referring to v2 of LED [20].

Rnd.	Value	Rnd.	Value	Rnd.	Value	Rnd.	Value
1	(0, 1, 0, 1)	6	(7, 6, 7, 6)	11	(3, 6, 3, 6)	16	(1, 6, 1, 6)
2	(0, 3, 0, 3)	7	(7, 5, 7, 5)	12	(7, 4, 7, 4)	17	(3, 5, 3, 5)
3	(0, 7, 0, 7)	8	(7, 3, 7, 3)	13	(7, 1, 7, 1)	18	(7, 2, 7, 2)
4	(1, 7, 1, 7)	9	(6, 7, 6, 7)	14	(6, 3, 6, 3)	19	(6, 5, 6, 5)
5	(3, 7, 3, 7)	10	(5, 7, 5, 7)	15	(2, 7, 2, 7)	20	(5, 3, 5, 3)

Table 4: The values for the second column of the round constant state matrix used in the first LED-round of all F -functions. The round constants for the second LED-round are set to all-zeros. The first column of the round constant matrix is (0, 1, 2, 3) for F_1^i , (0, 2, 2, 3) for F_2^j , and (0, 3, 2, 3) for F_3^j .

x	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
$S[x]$	c	5	6	b	9	0	a	d	3	e	f	8	4	7	1	2

Table 5: The s-box S of PRESENT used in LED

3.9.4 MixColumnsSerial (MCS)

This operation multiplies the current value of the state with the following 4×4 matrix.

$$\begin{bmatrix} 4 & 2 & 1 & 1 \\ 8 & 6 & 5 & 6 \\ b & e & a & 9 \\ 2 & 2 & f & b \end{bmatrix}$$

The result of the multiplication is the new value of the state.

4 Design Rationale

This section outlines the rationale for the design choices made for CiliPadi.

4.1 Key Lengths

We recommend two different key lengths, i.e. 128 and 256 bits. The former is meant for applications where resources such as area are very limited. The latter is proposed for applications where performance can be slightly sacrificed to gain more security.

4.2 Sponge

Our construction is based on the MonkeyDuplex [3, 6] Sponge, which evolves from the original Sponge proposed in 2007 [4]. The versatile construction has been extensively scrutinized and deployed in numerous hash functions and AE proposals. These include Keccak [5, 2] (standardized in SHA-3 and ISO/IEC 10118-3), PHOTON [18] (ISO/IEC 29192-5) and one of CAESAR's⁵ final port-

⁵Competition for Authenticated Encryption: Security, Applicability, and Robustness (CAESAR)

folio Ascon [16]. In particular, our use of different numbers of rounds in the initialization, message processing and finalization phases resembles that of Ascon and FIDES [8].

4.3 Permutation

The permutation function makes use of an unkeyed 2-round of the lightweight block cipher LED [19] as the F -function in a Type-II generalized feistel network (GFN) [24]. This is similar to the Simpira v2 permutation framework introduced by Gueron and Mouha [17]. To maximize diffusion, Simpira v2 utilizes a shuffling introduced by Suzuki and Minematsu [23], instead of the traditional left or right rotation of the input sub-blocks. The optimized shuffling allows us to achieve faster full diffusion of the input sub-blocks compared to the conventional Type-II GFN. A full diffusion means that all output sub-blocks are affected by all input sub-blocks.

Note that Simpira v2 was originally designed to utilize native AES instructions such as Intel’s AES-NI present in many modern processors. The aim is to achieve a high throughput implementation. As there is no hardware-specific instructions for LED, this is not our ultimate aim. We chose to follow Simpira v2 due to its flexibility in extending to larger input lengths and also because it is easy to analyze its security with respect to differential and linear cryptanalysis.

By employing a Type-II GFN, it is trivial to extend the input lengths in multiple of 128 bits. The use of 2-round LED as the F -function allows us to borrow the security analysis done on Simpira v2, which utilizes 2-round AES instead.

The LED block cipher is chosen due to its lightweight construction and its similarity to the AES. In contrast to a 1-round LED that has a minimum of one active s-box, a 2-round LED has a minimum of 5 active s-boxes, which is identical to the AES [15]. This allows us to easily extend the results of Gueron and Mouha [17], whom originally use AES in the Simpira v2 framework, to CiliPadi.

The number of rounds for P_n^b , i.e. $b = 16$, is one round extra than the suggested number of rounds for Simpira v2 (i.e. 15) for $d = 4$ and $d = 6$. Furthermore, P_n^a is two rounds more than P_n^b , which should provide ample protection against tag forgery in the finalization phase.

5 Performance Analysis

Our main goal for CiliPadi is a lightweight AEAD scheme with a very low hardware footprint. CiliPadi mainly consist of the internal operation of LED and XOR operation on its datapath. Three MUXes are the logic that control the overall operation of CiliPadi. There are 5 states of finite state machine (FSM), i.e. idle, initialization, authentication data, message and final phase. We make the state definition intuitive and straightforward. For simplicity, we only describe the implementation of CiliPadi-Mild. Other flavours of CiliPadi follow similarly.

At the heart of the implementation is a 2-round LED block cipher that consists of 4 operations: `AddConstansrs`, `SubCells`, `ShiftRows` and `MixColumnsSerial` as described in Section 3.9. Figure 6 shows the datapath structure of CiliPadi.




Figure 6: CiliPadi Datapath

On the left side is the input and output signal that comes from the control unit and from input output port of the datapath.

- key and nonce are both 128-bit signals.
- M_{in} is a 64-bit signal for the message block.
- A_{in} is a 64-bit signal for associated data.
- C_{out} is a 64-bit signal for the ciphertext block.
- K_{in} is the final 64-bit portion key input.
- T_{out} is the final 64-bit tag output.

The datapath for CiliPadi is controlled by the FSM. Input values to the MUXes are supplied by the state output logic design in the FSM. Figure 7 shows the state diagram of the FSM.

- idle state:
 - $init=1$ will transition to $init_phase$.
- $init_phase$ state:
 - r_sel_mux contains the r -bit (bitrate) part of the key concatenated with the nonce.
 - c_sel_mux contains the c -bit (capacity) part of the key concatenated with the nonce.
- $auth_data_phase$ state:
 - $c_sel_mux = c$ is the capacity part of the internal state.
 - r_sel_mux is the associated data XORed with the bitrate part of the state.




Figure 7: CiliPadi Finite State Machine

- $am_sel_mux = AD_in$ contains the associated data.
- message_phase state:
 - $c_sel_mux = 1xorR(0)$ refers to the XOR operation of the least significant bit of the state with the bit 1.
 - $r_sel_mux = MxorR$ is the XOR of the message with the bitrate part of the internal state.
 - $am_sel_mux = M_in$ is the message block input.
- final_phase state:
 - $c_sel_mux = c$ is the capacity part of the internal state.
 - $r_sel_mux = KxorR$ refers to the XOR of the r -bit part of the key with the bitrate part of the state.
 - $am_sel_mux = K_in$ is the key.

State transition is controlled by the status signal from datapath indicated by dotted lines between the state. In each state, MUXes value will change according to the input that is supposedly go to the registers inside the permutation entity. The datapath is controlled by this MUXes values for the final 64 bit tag output and ciphertext block.

5.1 Implementation Results

We implemented and simulated CiliPadi-Mild using VHDL on Xilinx ISE and synthesized on xc6vlx760-2ff1760 FPGA chip inside the Virtex 6 development board. We use RTL approach and make use of an iterative type design. The

other flavours of CiliPadi are obtained from our estimate based on the implementation of Mild. Table 6 shows a comparison of our implementation with other similarly designed AEs, i.e. Beetle [12] and ASCON [16]. These schemes are chosen due to the use of a sponge-based construction. Furthermore, the latter uses different number of rounds for the initialization and message encryption phases, which is similar to our design. The results for ASCON are obtained from Athena’s database [1].

Scheme	LUTs	Slices	Freq. (MHz)	Gbps	Mbps/ LUT	Mbps/ Slice
Beetle[Light+] [12]	616	252	381.592	1.879	3.050	7.369
Beetle[Secure+] [12]	998	434	256	2.520	2.525	5.806
ASCON-128 [1]	1274	451	341.1	3.118	2.447	6.914
ASCON-128a [1]	1587	547	358.6	5.099	3.213	9.322
CiliPadi-Mild	1052	303	639.959	1.138	1.082	3.756
CiliPadi-Medium	747	363	620	1.353	1.8112	3.727
CiliPadi-Hot	1268	373	631	1.685	1.329	4.517
CiliPadi-ExtraHot	1332	392	605	1.760	1.321	4.490

Table 6: Comparison of FPGA implementation of CiliPadi against other similar designs on Virtex 6

Beetle[Light+] and Beetle[Secure+] respectively use a 144-bit key (with 64 bits of security) and 256-bit key (with 121-bit security). On the other hand, both variants of ASCON employ a 128-bit key. The main differences between ASCON-128 and ASCON-128a are the message block size and numbers of permutation rounds. The former accepts a 64-bit block size while the latter, 128 bits.

The execution of CiliPadi-Mild and CiliPadi-Hot producing one ciphertext block and tag is 72 clock cycles. For CiliPadi-Medium and CiliPadi-ExtraHot, the number of clock cycles is 88. The primary member of CiliPadi, i.e. Mild, occupies 303 slices, which is about 20% higher than Beetle[Light+] but lower than other AEs compared in Table 6. However, our scheme provides 128-bit security as compared to Beetle[Light+]’s 64-bit. In the 256-bit key space, the two flavours of CiliPadi, i.e. Hot and ExtraHot consume fewer numbers of slices compared to Beetle[Secure+]. They even surpassed both variants of the 128-bit key ASCON. The numbers of slices for both Hot and ExtraHot flavours of CiliPadi are below 400 whereas the other compared AEs exceed this number. We believe that the implementation of CiliPadi can be further improved.

6 Security Analysis

6.1 Differential Cryptanalysis

In this section, we analyze the security of CiliPadi against differential cryptanalysis both theoretically and experimentally. We then extend some of these findings to linear cryptanalysis. Note that although there is no LED round subkeys to recover, such analysis is still useful since our analysis falls in the known-key attack model [21]. In our case, the known-key is the round constants.

6.1.1 Preliminaries

P is based on the Simpira v2 framework which is a d -line Type-II GFN. Instead of using AES like Simpira, CiliPadi uses an unkeyed 2-round LED as its round function, F . The design of LED shares same diffusion properties as AES, thus 2 rounds of LED has 5 active s-boxes (AS) [15]. In other words, F has a minimum of 5 active s-boxes given a non-zero input. The maximum differential probability of its s-box (which is essentially PRESENT’s s-box [9]) is 2^{-2} . We evaluate CiliPadi using two approaches, the first of which is based on the notions of P as a random permutation and the second is based on identifying collision-producing differentials. An upper bound of the differential probability for both approaches will be defined based on the number of “active F-functions”. We first provide some brief definitions for these concepts:

P as a random permutation: It has been shown that the security of sponge and duplex constructions rely on their underlying permutations being random [7, 3]. To evaluate P as a random permutation, we take into consideration the entire internal state, S as a whole. The maximum differential probability for P must be 2^{-256} and 2^{-384} for internal states of $n = 256$ and $n = 384$ respectively, to demonstrate some semblance to a random permutation. To obtain a differential path for the entire state S , a related-key/related nonce differential attack can be applied on the initialization phase, whereby an adversary is allowed to inject differences in either the key or nonce (or both of them). We can then define a differential path as $\Delta X \xrightarrow{\hat{p}} \Delta Y$, where \hat{p} is probability that an input difference ΔX leads to an output difference ΔY . We refer to \hat{p} as the differential probability. ΔX and ΔY are defined as

$$\Delta X = (K_1 \oplus K_2) \parallel (N_1 \oplus N_2)$$

and

$$\Delta Y = P_n^a(K_1 \parallel N_1) \oplus P_n^a(K_2 \parallel N_2)$$

respectively.

Thus, if $\Delta X \xrightarrow{\hat{p}} \Delta Y$ holds with $\hat{p} > 2^{-n}$, P is susceptible to a trivial distinguishing attack:

1. Initialize a counter, $c = 0$.
2. Generate 2^n related-key/related-nonce pairs corresponding to ΔX .
3. Encrypt each key-nonce pair to obtain the corresponding output difference.
4. For each key-nonce pair that fulfills, $\Delta X \rightarrow \Delta Y$, increment c .
5. Statistically, the distinguishing attack is successful if $c \approx 2^{\hat{p}}$.

Note that computing the differential path requires that random subkeys be used for each round of the underlying LED cipher to make the s-box inputs independent. These subkeys can be simulated by the addition of round constants. We also note that exhaustively searching through a state space of 256 or 384 bits

exceeds today’s computing capability. Therefore, these results are only of academic interest.

Collision-producing differential: In the AD authentication and message encryption phases, a straightforward application of differential cryptanalysis is difficult because the initial capacity state S_c is unknown to an adversary. However, we can investigate collision-producing differentials for CiliPadi which are differentials that have differences in S_r for both the input and outputs of P , but have zero differences for S_c . Such differentials may be useful in forgery attacks. A collision-producing differential can be defined as

$$\Delta X \parallel 0_2^c \xrightarrow{\hat{p}} \Delta Y \parallel 0_2^c \quad (1)$$

where ΔX can be introduced by injecting a difference in the message block, $\Delta Y \parallel 0^c = P_n^b(\Delta X \parallel 0_2^c)$, and \hat{p} is the probability that the differential holds.

Active F -function analysis: To theoretically estimate the upper bounds of the differential probability, we use the notion of “active F -functions” (AF). An F -function is considered to be active when it receives a non-zero input, similar to the concept of AS. We have implemented a searching algorithm to compute the number of AF for each round which is equivalent to identifying the number of AS for a regular GFN. The algorithm is based on a modified version of Matsui’s branch-and-bound search proposed in [13].

To simplify explanations, we will use the concept of truncated differentials, whereby every 64 bits of the concrete difference is represented as 1 bit in the truncated difference. A non-zero 64-bit block results in a non-zero truncated bit, and vice versa. E.g. for $d = 4$, if a concrete difference consists of four 64-bit differences, $\Delta X = (0_2^{63} \parallel 1_2) \parallel (0_2^{63} \parallel 1_2) \parallel (0_2^{64}) \parallel (0_2^{64})$, the corresponding truncated difference is $\Delta X^T = (1100)$. Thus, the maximum Hamming weight of the truncated difference is equal to d . The odd numbered bits (1,3,6) of the truncated difference will pass through the F -function, thus ”activating” it.

Using the searching algorithm, we identify the number of AF for $d = 4$ and $d = 6$ for up to 30 rounds as shown in Tables 7 and 8 respectively. In the following sections, we use this methodology to first evaluate P as a random permutation before examining its security against the distinguishing attack.

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

Table 7: Active F -function distribution for CiliPadi-Mild/Medium ($d = 4$)

6.1.2 P as a Random Permutation

For P to approximate a random permutation, the differential probability should be at most 2^{-256} (2^{-384}) for $d = 4$ ($d = 6$). As the differential probability of

Round	1	2	3	4	5	6	7	8	9	10
AF	0	1	2	3	4	6	8	10	11	12
Round	11	12	13	14	15	16	17	18	19	20
AF	13	14	15	17	19	21	22	23	24	25
Round	21	22	23	24	25	26	27	28	29	30
AF	26	28	30	32	33	34	35	36	37	39

Table 8: Active F -function distribution for CiliPadi-Hot/ExtraHot ($d = 6$)

the s-box is 2^{-2} , this requires at least $\frac{256}{2} = 128$ ($\frac{384}{2} = 192$) AS. A conservative approach is to assume each AF to contain only 5 AS as was the assumption made for a Type-II GFN that has 2 substitution layers interleaved with a single maximum distance separable (MDS)-based diffusion layer [10] (same as our 2-round LED with the exception that ours have an extra diffusion layer). Based on this approach, therefore, $\frac{256}{2 \times 5} \approx 26$ ($\frac{384}{2 \times 5} \approx 39$) AF is required in order for P to resist differential cryptanalysis. Indeed, 26 (39) AF gives $26 \times 5 = 130$ ($39 \times 5 = 195$) AS. Based on this rough estimate, according to Table 7 (8), P needs to have at least 27 (30) rounds for $d = 4$ ($d = 6$) to avoid any biases from a random permutation.

However, since P makes use of LED, which inherits the wide trails strategy of the AES, we can improve the previous analysis. As illustrated in Figure 8 and proven by the designers of AES, any 4-round differential path provides a minimum of 25 active s-boxes. Suppose that a particular round in LED has one AF where the internal differential paths looks like the first 2 rounds of the 4-round path depicted in the figure. We can observe that MCS causes all 4-bit cells to have nonzero output difference. Then, these 16 nonzero differences become the input to the next subsequent F -function which activate 16 AS in the first LED round and another 4 AS in the second LED round. It is therefore not possible for this second AF to have 5 AS as assumed before. Due to the wide trail strategy, this second AF is guaranteed to contain 20 AS. The AS pattern is $1 \rightarrow 4 \rightarrow 16 \rightarrow 4 \rightarrow 1$.

Table 9 expands the AF distribution given in Tables 7 and 8 by showing examples of truncated differential paths for P_n^a for all flavours of CiliPadi. As given earlier in this section, for $d = 4$, 128 AS are required in order for P to be resistant to differential cryptanalysis. For $d = 6$, the number of AS is 192. Based on Table 9, the truncated differential path for P_{256}^{18} and P_{256}^{20} each contains a minimum of 180 and 185 AS, respectively. On the other hand, the path for P_{384}^{18} and P_{384}^{20} each comprises 265 and 290 AS, respectively. These numbers for CiliPadi are beyond the required number of AS. The upper bounds of the differential probability \hat{p}_u of P_n^a for all variants of CiliPadi are shown in Table 10. For each variant, we also provide the number of truncated paths that correspond to the number of AF. Note that a 6-round iterative truncated differential with 6 AF ($0001 \rightarrow 0001$), and a 16-round iterative truncated differential with 22 AF ($000001 \rightarrow 000001$) exists for $d = 4$ and $d = 6$ respectively.

Practical Confirmation: We now experimentally confirm that our “active F -function” estimation is a conservative lower bound, and that the actual number of AS per AF would be higher than 5 as the number of rounds increases. In




Figure 8: A 4-round differential path for LED that guarantees at least 25 active s-boxes (AS)

other words, the number of AF provides us with an upper bound in terms of differential probability. To perform the differential search, we leverage upon the methodology described in [14]. Here, we focus on only the CiliPadi-Mild as a proof-of-concept and use the truncated differential path shown in Table 9 as a guide. We limit our input difference to have a hamming weight of 1 (only 1 bit out of any 64-bit word will be active at one time). Due to computational limitations, we bound the search based on each round of LED as:

- **LED Round 1:** Based on the input difference, if the number of activated s-boxes is more than 8, we limit the number of branches to 2 for each s-box, whereby we select the two branches with the highest differential probability. Otherwise, we search all branches.
- **LED Round 2:** Based on the input difference, if the number of activated s-boxes is more than 4, we limit the number of branches to 1 for each s-box, whereby we select the branch with the highest differential probability. Otherwise, we search all branches.

Based on this methodology, we found a 4-round concrete path following the truncated differential path $0001 \rightarrow 0111$ with a probability of 2^{-140} :

$$\Delta X = 0_2^{64} \| 0_2^{64} \| 0_2^{64} \| (0_2^{28} \| 2 \| 0_2^{32})$$

$$\Delta Y = 0_2^{64} \| 7f46a6de679866ce \| 8f01218a4117896f \| (0_2^{28} \| 2 \| 0_2^{32})$$

where ΔY is post-shuffle. The breakdown of the concrete differential path, the number of AS per round, and the comparison to our lower bound AS_l , is shown in Table 11. In the second round, there are 5 AF which leads to a probability

Rnd.	Mild (P_{256}^{18})			Medium (P_{256}^{20})			Hot (P_{384}^{18})			ExtraHot (P_{384}^{20})		
	ΔX	AF	AS	ΔX	AF	AS	ΔX	AF	AS	ΔX	AF	AS
1	0001	0	0	0001	0	0	000001	0	0	000011	1	5
2	0010	1	5	0010	1	5	001000	1	5	000001	0	0
3	0110	1	20	0110	1	20	010010	1	20	001000	1	5
4	1110	2	10	1110	2	10	101001	2	10	010010	1	20
5	0111	1	20	0111	1	20	111110	3	60	101001	2	10
6	1100	1	5	1100	1	5	011111	2	10	111110	3	60
7	0001	0	0	0001	0	0	110001	1	20	011111	2	10
8	0010	1	5	0010	1	5	001100	1	5	110001	1	20
9	0110	1	20	0110	1	20	010000	0	0	001100	1	5
10	1110	2	10	1110	2	10	100000	1	5	010000	0	0
11	0111	1	20	0111	1	20	100100	1	20	100000	1	5
12	1100	1	5	1100	1	5	100110	2	10	100100	1	20
13	0001	0	0	0001	0	0	101111	3	60	100110	2	10
14	0010	1	5	0010	1	5	110111	2	10	101111	3	60
15	0110	1	20	0110	1	20	000111	1	20	110111	2	10
16	1110	2	10	1110	2	10	000011	1	5	000111	1	20
17	0111	1	20	0111	1	20	000001	0	0	000011	1	5
18	1100	1	5	1100	1	5	001000	1	5	000001	0	0
19				0001	0	0				001000	1	5
20				0010	1	5				010010	1	20

Table 9: Numbers of AF and AS for truncated differential paths for P_n^a

CiliPadi-	n	a	AF	AS	\hat{p}_u	Truncated Paths
Mild	256	18	18	180	2^{-360}	122
Medium	256	20	19	185	2^{-370}	10
Hot	384	18	23	265	2^{-530}	144
ExtraHot	384	20	25	290	2^{-580}	36

Table 10: Differential probability upper bounds for P_n^a

of $(2^{-2})^5 = 2^{-10}$, verifying the correctness of our implementation. However, although the third round has only 1 AF, the actual number of AS is 28 due to the strong diffusion capability of LED. It confirms our claim that the number of “active F -functions” can be used as a conservative estimate of the security margin, and will lead to a conservative lower bound in terms of security margin (and equivalently, an upper-bound in terms of differential probability).

6.1.3 Collision-Producing Differentials of CiliPadi

The number of AF for a collision-producing truncated differential for CiliPadi-Mild and CiliPadi-ExtraHot can be identified by fixing both the input and output truncated differences to “1000” and “110000” respectively (i.e. $1000 \rightarrow 1000$ and $110000 \rightarrow 110000$ because r is a multiple of 64. For ease of analysis, we use $1000 \rightarrow 1000$ as the truncated differential for CiliPadi-Medium, by setting the remaining $96 - 64 = 32$ bits of the bitrate part to nonzero. For CiliPadi-Hot,

Concrete Differential				AS	AS _l
0_2^{64}	0_2^{64}	0_2^{64}	$(0_2^{28} 2 0_2^{32})$	0	0
0_2^{64}	0_2^{64}	$(0_2^{28} 2 0_2^{32})$	0_2^{64}	5	5
0_2^{64}	$(0_2^{28} 2 0_2^{32})$	c25adcad9fdb44b1	0_2^{64}	28	20
$(0_2^{28} 2 0_2^{32})$	c25adcad9fdb44b1	7f46a6de679866ce	0_2^{64}	36	10
0_2^{64}	7f46a6de679866ce	8f01218a4117896f	$(0_2^{28} 2 0_2^{32})$	-	-

Table 11: Example of a concrete differential path for P_{256}^a

we use $100000 \rightarrow 100000$. We then employ the same searching algorithm to identify the truncated differential path with the lowest number of AF for P_n^b . The results are summarized in Table 12 where \hat{p}_{col} denotes the probability of the collision-inducing path. The truncated paths corresponding to each of CiliPadi’s variants are as shown in Table 13. Note that a 6-round iterative truncated collision-producing differential exists for $d = 4$, where $1000 \rightarrow 1000$ with 6 AF.

CiliPadi-	n	b	AF	AS	\hat{p}_{col}
Mild	256	16	18	180	2^{-360}
Medium	256	18	18	180	2^{-360}
Hot	384	16	22	260	2^{-520}
ExtraHot	384	18	26	310	2^{-620}

Table 12: Collision-producing differential probability upper bounds for P_n^b

Practical Confirmation: We now experimentally confirm that the collision probabilities in Table 12 provides conservative lower bounds. Again, we target CiliPadi-Mild as a proof-of-concept and use the truncated differential path shown in Table 13 as a guide. Based on the same methodology described in Section 6.1.2, we found a 3-round concrete path following the truncated differential path with a probability of 2^{-140} :

$$\Delta X = (0_2^{28} || 2 || 0_2^{32}) || 0_2^{64} || 0_2^{64} || 0_2^{64}$$

$$\Delta Y = 8f01218a4117896f || (0_2^{28} || 2 || 0_2^{32}) || 0_2^{64} || 7f46a6de679866ce$$

where ΔY is post-shuffle. The above differential path contains a total of 69 AS, which is significantly higher than the theoretical lower bound of 35 AS for a 3-round differential, as shown in Table 13.

6.1.4 Practical Security Bounds

In practice, the best cryptanalytic attack requires less computational complexity than an exhaustive search of the secret key. CiliPadi has key sizes of 128 and 256 bits, thus any statistical distinguisher for a successful attack must have a probability higher than 2^{-128} and 2^{-256} respectively. Based on Tables 10 and 12, the theoretical upper bounds of the differential probability indicate that all flavours of CiliPadi are highly resistant to differential cryptanalysis and collision attacks. In reality, the differential probabilities are much lower as depicted in

Rnd.	Mild (P_{256}^{16})			Medium (P_{256}^{18})			Hot (P_{384}^{16})			ExtraHot (P_{384}^{18})		
	ΔX	AF	AS	ΔX	AF	AS	ΔX	AF	AS	ΔX	AF	AS
1	1000	1	5	1000	1	5	100000	1	5	110000	1	5
2	1001	1	20	1001	1	20	100100	1	20	000100	0	0
3	1011	2	10	1011	1	10	100110	2	10	000010	1	5
4	1101	1	20	1101	2	20	101111	3	60	001001	1	20
5	0011	1	5	0011	1	5	110111	2	10	011010	2	10
6	0100	0	0	0100	0	0	000111	1	20	111011	3	60
7	1000	1	5	1000	1	5	000011	1	5	010111	1	5
8	1001	1	20	1001	1	5	000001	0	0	101011	3	60
9	1011	2	10	1011	2	10	001000	1	5	110111	2	10
10	1101	1	20	1101	2	5	010000	1	20	000111	1	20
11	1011	2	10	0011	1	5	100000	2	10	000011	1	5
12	1101	1	20	0100	0	0	100100	3	60	000001	0	0
13	1011	2	10	1000	1	5	100110	2	10	001000	1	5
14	1101	1	20	1001	1	20	101111	1	20	010010	1	20
15	0011	1	5	1011	2	10	110101	1	5	101001	2	10
16	0100	0	0	1101	2	20	001110	0	0	111110	3	45
17				0011	1	5				111101	2	10
18				0100	0	0				011100	1	20

Table 13: Numbers of AF and AS for truncated collusion-producing differential paths for P_n^b

the practical confirmation experiments. Therefore, CiliPadi is expected to thwart any differential type attacks.

6.2 Full Bit Diffusion

With the availability of the full AF distribution, we can determine the minimum number of rounds for P to achieve full bit diffusion. Here, the findings from Simpira v2 are directly applicable because the underlying round functions for both Simpira and CiliPadi have the same diffusion properties. For $d = 4$, full bit diffusion is achieved after $4d - 6 = 16 - 6 = 10$ F -functions [17]. Based on Table 7, 9 rounds of P is sufficient for full bit diffusion. As for $d = 6$, full bit diffusion is achieved after 5 rounds [23]. Thus, the current number of rounds of P for all variants of CiliPadi are sufficient to achieve full bit diffusion.

6.3 Extension to Linear Cryptanalysis

The previous findings on differential cryptanalysis can be trivially extended to linear cryptanalysis due to the duality between linear and differential cryptanalysis [22, 10], and also due to PRESENT's s -box having a linear probability of 2^{-2} . Thus, all the results in the previous subsections are applicable to linear cryptanalysis.

7 Strengths and Weaknesses

The following list the expected strengths and weaknesses of CiliPadi.

7.1 Strengths

CiliPadi has the following advantages:

- It is trivial to expand the length of the permutation in multiple of 128 bits due to the use of a Type-II GFN.
- The bitrate can be adjusted to allow different plaintext and tag lengths.
- The design is based on the sponge construction, which have been extensively analyzed and employed in SHA-3 and one of the CAESAR portfolio ASCON.
- Any AES-like block cipher or permutation can be adopted in the F -function to replace the LED block cipher, if desired.

7.2 Weaknesses

The known limitations of CiliPadi are:

- The processing of the message and ciphertext blocks cannot be parallelized because due to the sequential processing of the input blocks.
- The permutation can only be expanded in multiple of 128 bits. Extending with a smaller granularity, e.g. 32 bits, is not supported. This can be addressed by using a smaller block cipher as the F -function such as KATAN and KTANTAN [11].

References

- [1] *Authenticated Encryption FPGA Ranking*. URL: https://cryptography.gmu.edu/athenadb/fpga_auth_cipher/table_view (visited on Mar. 20, 2019).
- [2] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. *Cryptographic Sponge Functions, Version 0.1*. <http://keccak.noekeon.org>. Jan. 2011.
- [3] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. *Permutation-Based Encryption, Authentication and Authenticated Encryption*. DIAC – Directions in Authenticated Ciphers. July 2012.
- [4] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. *Sponge Functions*. ECRYPT Hash Workshop 2007. 2007.
- [5] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. *The Keccak SHA-3 Submission, Version 3*. SHA-3 Cryptographic Hash Algorithm Competition. <http://keccak.noekeon.org>. Jan. 2011.

- [6] Guido Bertoni, Joan Daemen, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. *CAESAR submission: Ketje v2*. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. <https://competitions.cr.ypt.to/round3/ketjev2.pdf>. 2016.
- [7] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. “Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications”. In: *Selected Areas in Cryptography, SAC 2011*. Ed. by Ali Miri and Serge Vaudenay. Vol. 7118. Lecture Notes in Computer Science. Springer, 2012, pp. 320–337. DOI: 10.1007/978-3-642-28496-0_19.
- [8] Begül Bilgin, Andrey Bogdanov, Miroslav Knežević, Florian Mendel, and Qinju Wang. “Fides: Lightweight Authenticated Cipher with Side-Channel Resistance for Constrained Hardware”. In: *Cryptographic Hardware and Embedded Systems – CHES 2013*. Ed. by Guido Bertoni and Jean-Sébastien Coron. Vol. 8086. Lecture Notes in Computer Science. Springer-Verlag, 2013, pp. 142–158. DOI: 10.1007/978-3-642-40349-1_9.
- [9] Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. “PRESENT: An Ultra-Lightweight Block Cipher”. In: *Cryptographic Hardware and Embedded Systems – CHES 2007*. Ed. by Pascal Paillier and Ingrid Verbauwhede. Vol. 4727. Lecture Notes in Computer Science. Springer-Verlag, 2007, pp. 450–466.
- [10] Andrey Bogdanov and Kyoji Shibutani. “Generalized Feistel Networks Revisited”. In: *Designs, Codes and Cryptography* 66.1–3 (2013), pp. 75–97.
- [11] Christophe De Cannière, Orr Dunkelman, and Miroslav Knežević. “KATAN and KTANTAN – A Family of Small and Efficient Hardware-Oriented Block Ciphers”. In: *Cryptographic Hardware and Embedded Systems, CHES 2009*. Ed. by Christophe Clavier and Kris Gaj. Vol. 5747. Lecture Notes in Computer Science. Springer-Verlag, 2009, pp. 272–288.
- [12] Avik Chakraborti, Nilanjan Datta, Mridul Nandi, and Kan Yasuda. “Beetle Family of Lightweight and Secure Authenticated Encryption Ciphers”. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES)* 2018.2 (2018), pp. 218–241. DOI: 10.13154/tches.v2018.i2.218-241.
- [13] Jiageng Chen, Atsuko Miyaji, Chunhua Su, and Je Sen Teh. “Accurate Estimation of the Full Differential Distribution for General Feistel Structures”. In: *Information Security and Cryptology, Inscrypt 2015*. Ed. by Dongdai Lin, XiaoFeng Wang, and Moti Yung. Vol. 9589. Lecture Notes in Computer Science. Springer-Verlag, 2016, pp. 108–124. DOI: 10.1007/978-3-319-38898-4_7.
- [14] Jiageng Chen, Jesen Teh, Zhe Liu, Chunhua Su, Azman Samsudin, and Yang Xiang. “Towards Accurate Statistical Analysis of Security Margins: New Searching Strategies for Differential Attacks”. In: *IEEE Transactions on Computers* 66.10 (Oct. 2017), pp. 1763–1777. DOI: 10.1109/tc.2017.2699190.
- [15] Joan Daemen and Vincent Rijmen. *The Design of Rijndael, AES – The Advanced Encryption Standard*. Springer-Verlag, 2002.

- [16] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer. *Ascon v1.2: Submission to the CAESAR Competition*. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. <https://competitions.cr.ypt.to/round3/asconv12.pdf>. 2016.
- [17] Shay Gueron and Nicky Mouha. “Simpira v2: A Family of Efficient Permutations Using the AES Round Function”. In: *Advances in Cryptology – ASIACRYPT 2016, Part I*. Ed. by Jung Hee Cheon and Tsuyoshi Takagi. Vol. 10031. Lecture Notes in Computer Science. Springer-Verlag, 2016, pp. 95–125.
- [18] Jian Guo, Thomas Peyrin, and Axel Poschmann. “The PHOTON Family of Lightweight Hash Functions”. In: *Advances in Cryptology – CRYPTO 2011*. Ed. by Phillip Rogaway. Vol. 6841. Lecture Notes in Computer Science. Springer, 2011, pp. 222–239. DOI: 10.1007/978-3-642-22792-9\13.
- [19] Jian Guo, Thomas Peyrin, Axel Poschmann, and Matt Robshaw. “The LED Block Cipher”. In: *Cryptographic Hardware and Embedded Systems – CHES 2011*. Ed. by Bart Preneel and Tsuyoshi Takagi. Vol. 6917. Lecture Notes in Computer Science. Springer-Verlag, 2011, pp. 326–341.
- [20] Jian Guo, Thomas Peyrin, Axel Poschmann, and Matt Robshaw. *The LED Block Cipher*. Cryptology ePrint Archive, Report 2012/600. <http://ia.cr/2012/600>. 2012.
- [21] Lars R. Knudsen and Vincent Rijmen. “Known-Key Distinguishers for Some Block Ciphers”. In: *Advances in Cryptology – ASIACRYPT 2007*. Ed. by Kaoru Kurosawa. Vol. 4833. Lecture Notes in Computer Science. Springer-Verlag, 2007, pp. 315–324.
- [22] Mitsuru Matsui. “On correlation between the order of S-boxes and the strength of DES”. In: *Advances in Cryptology – EUROCRYPT ’94*. Ed. by Alfredo De Santis. Vol. 950. Lecture Notes in Computer Science. Springer-Verlag, 1995, pp. 366–375. DOI: 10.1007/bfb0053451.
- [23] Tomoyasu Suzaki and Kazuhiko Minematsu. “Improving the Generalized Feistel”. In: *Fast Software Encryption, FSE 2010*. Ed. by Seokhie Hong and Tetsu Iwata. Vol. 6147. Lecture Notes in Computer Science. Springer-Verlag, 2010, pp. 19–39.
- [24] Yuliang Zheng, Tsutomu Matsumoto, and Hideki Imai. “On the Construction of Block Ciphers Provably Secure and Not Relying on Any Unproved Hypotheses”. In: *Advances in Cryptology – CRYPTO ’89*. Ed. by Gilles Brassard. Vol. 435. Lecture Notes in Computer Science. Springer-Verlag, 1990, pp. 461–480.

A Test Vectors

The following are the test vectors for the recommended flavours of CiliPadi.

CiliPadi-Mild	
Key	0_2^{128}
Nonce	0_2^{128}
AD	1^{64}
Plaintext	0^{64}
Ciphertext	17a2da6e5f74a0c4
Tag	f85b6cc1321172e8
CiliPadi-Medium	
Key	0_2^{128}
Nonce	0_2^{128}
AD	1_2^{96}
Plaintext	0_2^{96}
Ciphertext	8d34c8413153119ac8b41336
Tag	b349ac1f840bf016c931feed
CiliPadi-Hot	
Key	0_2^{256}
Nonce	0_2^{128}
AD	1_2^{96}
Plaintext	0_2^{96}
Ciphertext	bcb4f3795963a9b0f8b8cc8e
Tag	0f63efca252e1bef83888aab
CiliPadi-ExtraHot	
Key	0_2^{256}
Nonce	0_2^{128}
AD	1_2^{128}
Plaintext	0_2^{128}
Ciphertext	e2998e76e93da0ea5ff746579765272f
Tag	18eb6410b0c04b9aa02fe01ce974469d

Table 14: CiliPadi test vectors