Differential Cryptanalysis of WARP

Je Sen Teh* and Alex Biryukov

***Abstract*—The proliferation of resource-constrained Internet-of-Things (IoT) devices that transmit sensitive data on a daily basis has led to the need for lightweight ciphers with minimal computational requirements. WARP is an energy-efficient lightweight block cipher that is currently the smallest 128-bit block cipher in terms of hardware. It was proposed by Banik et al. in SAC 2020 as a lightweight replacement for AES-128 without changing the mode of operation. This paper proposes key-recovery attacks on WARP based on differential cryptanalysis in single and related-key settings. We searched for differential trails for up to 20 rounds of WARP, with the first 19 having optimal differential probabilities. We also found that the cipher has a strong differential effect, whereby 16 to 20-round differentials have substantially higher probabilities than their corresponding individual trails. A 23-round key-recovery attack was then realized using an 18-round differential distinguisher. Next, we formulated an automatic boomerang search using SMT that relies on the Feistel Boomerang Connectivity Table to identify valid switches. We designed the search as an add-on to the CryptoSMT tool, making it applicable to other Feistel-like ciphers such as TWINE and LBlock-s. For WARP, we found a 21-round boomerang distinguisher which was used in a 24-round rectangle attack. In the related-key setting, we describe a family of 2-round iterative differential trails, which we used in a practical related-key attack on the full 41-round WARP.**

***Index Terms*—Constrained Devices, IoT, Symmetric-key, Block ciphers, Differential cryptanalysis, Boomerang distinguisher, Rectangle attack, Related-key, WARP, GFN**

I. Introduction

Lightweight cryptography is currently one of the most heavily researched areas in recent years [1]. This is due in part to the widespread use of resource-constrained devices such as smart or IoT devices which transmit sensitive information on a daily basis. Compared to other symmetric-key primitives, lightweight block ciphers have received the most attention in terms of development and cryptanalytic efforts. Although most lightweight block ciphers have block sizes of 64 bits there were a number of 128-bit block ciphers with lower area and/or power requirements than AES such as MIDORI [2] and GIFT-128 [3]. These 128-bit lightweight block ciphers are usually based on the Substitution-Permutation Network (SPN) design paradigm which generally takes up more hardware space due to the inversion of their confusion and diffusion layers.

To overcome this hurdle, Banik et al. adopted the Type-2 Generalized Feistel Network [4] in their 128-bit block cipher called WARP which was proposed in SAC 2020 [5]. The design team consisted of the minds behind multiple well-known lightweight block ciphers such as GIFT, MIDORI and TWINE [6]. The motivation behind designing WARP as a 128-bit cipher with a 128-bit key was to realize a direct replacement for AES-128 without having to change the underlying mode of operation. By adopting MIDORI's S-box (for reduced latency and area) and a simple alternating key schedule, the designers found that WARP only requires 763 Gate Equivalents (GE) for a bit-serial encryption-only circuit and has better energy consumption than MIDORI, which is widely considered the current state-of-the-art in terms of 128-bit low-energy ciphers.

Related Work. To the best of our knowledge, the only prior third-party cryptanalysis result for WARP was an attack described in an IACR ePrint [7]. In its current form, it seems to be invalid due to a possible error in the permutation pattern. As

Manuscript received xx, 2021; revised xx, 2021.

Je Sen Teh and Alex Biryukov are both with the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg.

Je Sen Teh is also with the School of Computer Sciences, Universiti Sains Malaysia (e-mail:jesen_teh@usm.my)

*Corresponding Author

such, the only cryptanalysis results to date are the ones provided by WARP's designers. In terms of differential and linear cryptanalysis, WARP has more than 64 active S-boxes after 19 rounds. The designers also found a 21-round impossible differential distinguisher and a 20-round integral distinguisher for the cipher. A meet-in-the-middle attack is expected to be feasible for at most 32 rounds of WARP. Although no concrete attacks were described, 41 rounds of WARP is expected to be secure against these attacks.

Our Contributions. In this paper, we cryptanalyze WARP using differential cryptanalysis. By using an SMT-aided differential search, we found differential trails for up to 20 rounds of WARP, with the first 19 guaranteed to be optimal. These differential trails confirm that the lower bounds provided by the designers cannot be improved. We then performed a differential cluster search for each of these trails and found that WARP has a strong differential effect from round 13 onward.

Next, we implemented an automatic search for boomerang (or more specifically, rectangle) distinguishers which includes the Feistel Boomerang Connectivity Table (FBCT) [8]. The boomerang search was written as a new module for the CryptoSMT tool [9] rather than one that was specifically catered to WARP[1]. We showcase its flexibility by also applying it to TWINE and LBlock-s [10]. Using our tool, we were able to find a 21-round boomerang distinguisher for WARP with a differential probability, $\mathsf{DP} = 2^{-121.11}$.

We also performed a search for related-key differential trails for WARP. As a result, we found that WARP has a family of 2-round iterative related-key differential trails with low weight. These iterative trails can be concatenated to form distinguishers for the full 41-round WARP with $\mathsf{DP} = 2^{-40}$. These trails exist due to the interaction between the cipher's nibble-wise permutation, simple alternating key schedule and subkey XOR operation performed *after* the S-box. The interaction between these design elements also led to another interesting observation whereby knowledge of the input difference for a Feistel-subround can be propagated to the next round without having to guess its corresponding subkey. This property was leveraged in all of our key recovery attacks to target specific subkeys.

[1]The implementation is publicly available under an open-source license at https://github.com/jesenteh/cryptosmt-boomerang, along with other supplementary codes for this paper.

R	Method	Time	Data	Mem	Sect.
23	SK Diff.	$2^{106.68}$	$2^{106.62}$	$2^{106.62}$	IV-A
24	SK Rect.	$2^{125.18}$	$2^{126.06}$	$2^{127.06}$	IV-B
41	RK Diff.	2^{37*}	2^{37}	$2^{9.59}$	V-B

*Time complexity to recover 60 bits of the key

TABLE I
SUMMARY OF KEY-RECOVERY ATTACKS ON WARP (SK/RK DENOTES SINGLE-KEY/RELATED-KEY)

Finally, we proposed key-recovery attacks on WARP based on the differential distinguishers that were found. In the single-key setting, we have a 23-round differential attack using an 18-round differential distinguisher that has time T, data D and memory M complexities of $(T, D, M) = (2^{106.68}, 2^{106.62}, 2^{106.62})$, followed by a 24-round rectangle attack using a 21-round boomerang distinguisher with $(T, D, M) = (2^{125.18}, 2^{125.06}, 2^{126.06})$. In the related-key setting, we formulated a 25-round attack using a 19-round related-key differential distinguisher for the purpose of computational verification. 16 bits of the secret key were recovered within 2.5 minutes. We extended the same key-recovery framework and introduced a practical attack on all 41 rounds of WARP using a 35-round related-key distinguisher with $(T, D, M) = (2^{37}, 2^{37}, 2^{9.59})$. Our cryptanalytic results are summarized in Table I.

II. Preliminaries

Notations and abbreviations used in this paper are summarized in Table II. The rightmost (least significant) bits or nibbles have an index of 0.

A. Differential Cryptanalysis

A block cipher maps a set of plaintexts to a set of ciphertexts using a key-dependent round function,

Symbol	Meaning
n	Block size in bits
k	Key size in bits
ΔP	Plaintext XOR difference
ΔC	Ciphertext XOR difference
$\alpha, \beta, \delta, \gamma$	n-bit input and output
α_i^j	The i-th nibble of an n-bit XOR difference, α in round j
X_i^j	The i-th nibble of an n-bit binary variable, X in round j
$\#AS$	Number of active S-boxes
$\oplus$	Binary exclusive-OR (XOR)
$\|\|$	Binary concatenation
DP	Differential probability
R	Number of rounds
$DDT(x,y)$/ $FBCT(x,y)$	An entry in the DDT/FBCT for an input x and output y

TABLE II
SYMBOLS AND NOTATION

f_j, where $j \in \mathsf{R}$. The goal of differential cryptanalysis is to find pairs of plaintexts (P_1, P_2) and ciphertexts (C_1, C_2) with a strong correlation between their differences $\alpha = P_1 \oplus P_2$ and $\beta = C_1 \oplus C_2$. The propagation pattern of an input difference α to an output difference β is known as a differential characterisitic or trail. A differential trail consists of a sequence of differences,

$$\alpha \xrightarrow{f_1} \alpha^1 \xrightarrow{f_2} \ldots \xrightarrow{f_{\mathsf{R}-2}} \alpha^{\mathsf{R}-2} \xrightarrow{f_{\mathsf{R}-1}} \beta. \quad (1)$$

An adversary must find a differential trail with sufficiently high differential probability,

$$\mathsf{DP} = \mathsf{Pr}(\alpha \xrightarrow{f_1} \ldots \xrightarrow{f_{\mathsf{R}-1}} \beta). \quad (2)$$

Based on the Markov assumption [11] which allows treating a cipher's rounds independently, the differential probability can be computed as

$$\mathsf{DP} \approx \prod_{j=1}^{\mathsf{R\text{-}1}} \mathsf{Pr}(\alpha^{j-1} \xrightarrow{f_j} \alpha^j), \quad (3)$$

where $\alpha_0 = \alpha$ and $\alpha_{\mathsf{R\text{-}1}} = \beta$. A better estimate of the differential probability can be obtained by collecting differential trails that share the same input and output differences,

$$\mathsf{DP} = \mathsf{Pr}(\alpha \rightarrow \beta) = \sum_{\alpha^1 \ldots \alpha^{\mathsf{R\text{-}2}}} (\alpha \xrightarrow{f_1} \ldots \xrightarrow{f_{\mathsf{R}-1}} \beta). \quad (4)$$

When cryptanalyzing a block cipher, an adversary maximizes the probability of the differential by enumerating as many differential trails as possible, which can be automated using methods such as Matsui's algorithm [12], MILP [13], boolean satisfiability problem (SAT) and satisfiability modulo theory (SMT) solvers [14] .

In the related-key setting, an adversary is allowed to also have a difference in the encryption key, and not only in the plaintext. However, the adversary cannot specify the value of the key itself and the attack must be valid for any pair of keys with the given difference. In the past, there have been practical attacks that rely on the related-key property [15]. Ciphers that are vulnerable to related-key attacks are not recommended for use in protocols where key integrity is not guaranteed [16].

B. Boomerang and Rectangle Attacks

The boomerang attack is a variant of differential cryptanalysis that concatenates two shorter differentials to form a longer distinguisher. The classical boomerang attack involves decomposing a target cipher, E into two subciphers, $E = E_1 \circ E_0$. The input and output differences of E_0 are denoted as α and β while for E_1, they are denoted as γ and δ. The probability that $\alpha \xrightarrow{E_0} \beta$ is p and $\gamma \xrightarrow{E_1} \delta$ is q. The boomerang attack was later reformulated

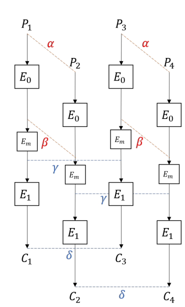

Fig. 1. Sandwich attack

as a chosen plaintext attack called the rectangle attack [17], [18] by encrypting many pairs with the input difference α and searching for a quartet which satisfies $C_1 \oplus C_3 = C_2 \oplus C_4 = \delta$ when $P_1 \oplus P_2 = P_3 \oplus P_4 = \alpha$. Although the probability of a quartet to be a right quartet is reduced to $2^{-n}p^2q^2$, counting over all possible β's and δ's as long as $\beta \neq \delta$ improves the probability to $2^{-n}\hat{p}^2\hat{q}^2$, where $\hat{p} = (\sum_i \mathsf{Pr}^2(\alpha \xrightarrow{E_0} \beta_i))^{\frac{1}{2}}$ and $\hat{q} = (\sum_j \mathsf{Pr}^2(\gamma \xrightarrow{E_1} \delta_j))^{\frac{1}{2}}$.

With the introduction of the sandwich attack [19], [20], the boomerang connectivity table (BCT) [21] and its Feistel counterpart [8], we can systematically enumerate these trails while guaranteeing their compatibility. The sandwich attack (Figure 1) decomposes the cipher into 3 components, $E = E_1 \circ E_m \circ E_0$, where E_m is the middle transition round with a switching probability, r that can be calculated using BCT or FBCT. The connectivity tables already the various switches that have been used in the past to improve the probability of boomerang distinguishers such as the ladder, S-box and Feistel switches [19]. The probability of obtaining a right quartet is

$$\hat{p}^2\hat{q}^2 = \sum_{i,j}(\hat{p_i}^2\hat{q_j}^2 r_{i,j}), \tag{5}$$

where $\hat{p}_i = \mathsf{Pr}(\alpha \xrightarrow{E_0} \beta_i)$, $\hat{q}_j = \mathsf{Pr}(\gamma_j \xrightarrow{E_1} \delta)$ and $r_{i,j} = \mathsf{Pr}(\beta_i \xrightarrow{E_m} \gamma_j)$.

C. Specification of WARP

The block cipher WARP is a 41-round, 128-bit block cipher with a 128-bit key designed based on a 32-nibble Type-2 GFN. The i-th round's state is divided into 32 nibbles, $X^i = X^i_{31}||X^i_{30}||...||X^i_1||X^i_0$, where $X^i_j \in \{0,1\}^4$. It has a simple key schedule that first divides the secret key into two 64-bit round keys, $K = K^1||K^0$, then alternates between them (starting from K^0). Each 64-bit round key is divided into 16 nibbles, $K^i = K^i_{15}||K^i_{14}||...||K^i_1||K^i_0$, where $K^i_j \in \{0,1\}^4$, $i \in \{0,1\}$. The round function is illustrated in Figure 2 while the S-box and permutation pattern, π are shown in Tables III and IV respectively. Apart from using the inverse permutation, π^{-1}, the decryption algorithm is the same.

x	0	1	2	3	4	5	6	7
$S(x)$	C	A	D	3	E	B	F	7
x	8	9	A	B	C	D	E	F
$S(x)$	8	9	1	5	0	2	4	6

TABLE III
WARP 4-BIT S-BOX

To avoid the complement property of Feistel-type ciphers [22], the designers of WARP opted for the key XOR operation to be after the S-box. However, this design decision leads to the following property:

Property 1 (Subround Filters). *Since XOR with the key is done after the S-box in the Feistel-subround which works on two nibbles, it allows to partially decrypt and propagate the knowledge of the difference to the next round. This can be done for both the top and bottom rounds.*

Based on Figure 3, we can see that partially encrypting P_1 and P_2 that correspond to the input difference, α allows to immediately check if the given pair is valid if the left nibble of the output difference, β_L is known. We can do this without having to guess the corresponding key nibble, K^j_i because the output difference of the S-box, which we denote as γ, can be directly computed from the

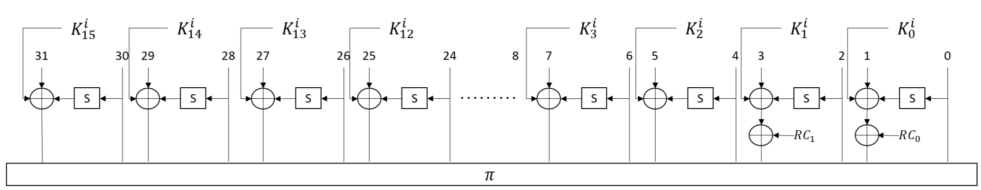

Fig. 2. Round Function of WARP

x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$\pi(x)$	31	6	29	14	1	12	21	8	27	2	3	0	25	4	23	10
$\pi^{-1}(x)$	11	4	9	10	13	22	1	30	7	28	15	24	5	18	3	16
x	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
$\pi(x)$	15	22	13	30	17	28	5	24	11	18	19	16	9	20	7	26
$\pi^{-1}(x)$	27	20	25	26	29	6	17	14	23	12	31	8	21	2	19	0

TABLE IV
WARP PERMUTATION

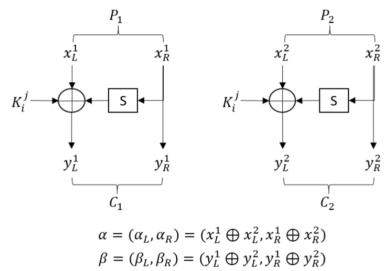

Fig. 3. Difference propagation for a pair of nibbles

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	2	4	0	2	2	2	0	2	0	0	0	0	0	2	0
2	0	4	0	0	4	0	0	0	0	4	0	0	4	0	0	0
3	0	0	0	0	2	0	4	2	2	2	0	0	0	2	0	2
4	0	2	4	2	2	2	0	0	2	0	0	2	0	0	0	0
5	0	2	0	0	2	0	0	4	0	2	4	0	2	0	0	0
6	0	2	0	4	0	0	0	2	2	0	0	0	2	2	0	2
7	0	0	0	2	0	4	2	0	0	0	0	2	0	4	2	0
8	0	2	0	2	2	0	2	0	0	2	0	2	2	0	2	0
9	0	0	4	2	0	2	0	0	2	2	0	2	2	0	0	0
A	0	0	0	0	0	4	0	0	0	0	4	0	0	4	0	4
B	0	0	0	0	2	0	0	2	2	2	0	4	0	2	0	2
C	0	0	4	0	0	2	2	0	2	2	0	0	2	0	2	0
D	0	0	0	2	0	0	2	4	0	0	4	2	0	0	2	0
E	0	2	0	0	0	0	0	2	2	0	0	0	2	2	4	2
F	0	0	0	2	0	0	2	0	0	0	4	2	0	0	2	4

TABLE V
WARP'S DIFFERENTIAL DISTRIBUTION TABLE

known values of x_R^1 and x_R^2. Thus, we can check if $\alpha_L \oplus \gamma = \beta_L$ because the effect of the round key has been negated by the XOR operation. The same property exists for the bottom rounds whereby partially decrypting known values of C_1 and C_2 when α_L is a known difference allows to check if $\beta_L \oplus \gamma = \alpha_L$. We will use these subround filters in our key recovery attacks.

III. SEARCHING FOR WARP DISTINGUISHERS

A. *Differential Distinguishers*

We use CryptoSMT [9] to search for both differential trails and differentials for WARP. First, a script was written to generate the SMT model that describes its differential propagation. Then, we enumerate the optimal differential trails for each round and perform differential clustering. Our findings are summarized in Table VI where $\#AS$ refers to the number of active S-boxes and the weight of a differential trail is calculated as $W = -\log_2 \mathsf{DP}$. For up to 19 rounds, we verified that the minimum number of active S-boxes mentioned in WARP's design specification was indeed the lower bound and also found the optimal differential trails for each of these rounds. The time required to find differential trails increased sharply with the number of rounds, compounded with the fact that we are dealing with a 128-bit block size. Finding a trail for Round 17 onward would take up to half a day, longer if a trail did not exist for a particular weight. We managed to find a differential trail for 20 rounds of WARP with a weight of 140 but could not verify its optimality.

Next, we clustered these differentials with a time limit of 24 hours. The results in Table VI show that for the first 12 rounds, all differentials either had 1 or very few trails each. From Round 13 onward, however, there was a sharp increase in the number of trails. We managed to find all trails for the 13-round to 15-round differentials, which

R	$\#AS$	Trail			Differential	
		W_{opt}	α	β	W_{diff}	#Trails
1	0	0	0000 0000 0000 0000 0000 0000 0000 0010	0000 0000 0000 0000 0000 0000 0100 0000	0	1
2	1	2	0000 0000 0000 0000 0000 0000 4000 0000	0000 4000 0000 0000 0000 0000 0000 0200	2	1
3	2	4	0000 2000 0021 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0200 0000 0000	4	1
4	3	6	0000 0000 002C 0000 0000 0000 0000 0000	0000 0000 0000 0400 000C 2000 0000 0000	6	1
5	4	8	0000 A000 00AA 0000 0000 0000 0000 0000	F000 0000 0000 0000 0A00 0000 0F00 0000	8	1
6	6	12	0092 0000 0000 0000 0000 0000 9000 0042	0000 0000 0000 9002 0000 0002 0000 0000	12	1
7	8	16	0000 7DA0 FF00 0000 0000 0000 0000 0000	0000 0000 0000 A005 000A 000F 0000 00F0	16	1
8	11	22	0000 00FA 5A00 0000 00A0 0000 0000 0000	0000 0000 0000 A705 000A 500A 0000 A005	22	1
9	14	28	0000 0000 0000 1000 0000 C2C0 4200 0012	2900 0020 0000 0000 0120 0000 0104 0001	28	1
10	17	34	E000 00EE EE00 0000 00E0 00EE 0000 0000	E000 0000 0E0E 0000 0E0E 00E0 0E00 E00E	33.19	7
11	22	44	0012 0000 1000 1290 1212 0000 1000 0042	2000 0000 0101 0000 0101 0020 0100 2004	43.19	7
12	28	56	1212 0000 4000 0042 0012 0000 4000 4240	0200 0202 0212 0200 1002 0212 4040 0010	55.42	5
13	34	68	0020 2000 0024 2000 0000 0020 2121 0021	0010 0202 1000 0000 1200 0240 4000 1202	62.37	1600
14	40	80	0000 0010 1292 0012 0010 1000 0042 C000	0000 1002 0200 4202 40C0 0002 C002 4202	72.14	21528
15	47	94	0000 00A0 5A5A 005A 00A0 5000 0057 5000	A500 A005 000A 0700 0AA5 55A5 057A 0AA0	85.54	497248
16	52	104	A000 5AAA 0000 0000 0000 A05A 005A 0000	0A00 000A 000A 0000 0057 0A50 005A 500A	90.52	800152
17	57	114	0000 A000 0000 0075 0000 A500 0000 7000	000A 5000 0550 0000 AA00 000A 0000 0A00	95.66	734494
18	91	122	0000 A0AF 005A 0000 A000 AA75 0000 0000	000A 5000 0AA0 0000 5A00 000A 0000 0A00	104.62	626723
19	66	132	5000 A55A 0000 0000 0000 70AA 00A5 0000	0500 0050 00A0 0A00 00A5 A00A 5007 000A	118.07	594111
20	70*	140*	0000 50AA 0057 0000 F000 5AAF 0000 0000	0A00 A000 0000 500A 0000 050A 0000 F50A	122.71	545045

*Number of active S-boxes and/or differential probability not confirmed to be optimal

TABLE VI
WARP DIFFERENTIALS FOR ROUNDS 1 TO 20

ranged from 1600 to just under 500000 trails. The remaining cluster sizes were bounded by the time limit. The results show that WARP has a significant differential effect at higher rounds, whereby rounds 16 to 20 have an improvement to their differential probabilities by a factor of at least $2^{13.48}$.

B. Boomerang Distinguishers

To find boomerang distinguishers for WARP, we formulated an automatic boomerang search based on CryptoSMT's differential search functionality. The overall goal of the automatic search is to maximize $\hat{p}^2\hat{q}^2 = \sum_{i,j}(\hat{p_i}^2\hat{q_j}^2 r_{i,j})$ by finding as many E_0 and E_1 trails that are compatible. The compatibility of the upper and lower trails is determined using the FBCT[2]. WARP's FBCT shown in Table VII. The proposed boomerang search procedure is as follows:

1) Search for an E_0 trail with R_{E_0} rounds for up to a weight limit of W_{upper}.
2) Search for an E_1 trail with R_{E_1} rounds for up to a weight limit of W_{lower}. Limit the search to only compatible trails by propagating β from E_0 through E_m, then including blocking constraints in the SMT model for each of its S-boxes based on entries in the FBCT. If a valid E_1 trail is found then:
   a) If this is the first iteration, fix the input and output differences of the boomerang distinguisher to α and δ for all future iterations.
   b) Calculate the switching probability, $r_{i,j}$ based on β, γ, the linear layer, π and FBCT as
   $$r_{i,j} = \prod_{k=\{2,4,\ldots,28,30\}} \frac{FBCT(\beta_k, \pi^{-1}(\gamma_k) - 1)}{16}. \quad (6)$$
   c) For the clustering process, limit the search to $W_{init} + \frac{n}{l}$ where W_{init} is the weight of the initial trail and l controls the upper limit of the search, e.g. for $l = 64$, the upper weight limit of the clustering process is $W_{init} + 2$. Set individual limits for E_0 and E_1.
   d) Perform differential clustering for E_0 if it has not yet been done. Denote the resulting differential probability as $\hat{p}_i$.
   e) Perform differential clustering for E_1. Denote the resulting differential proba-

[2]For more information about the FBCT, readers can refer to the work by Boukerrou et al. [8]

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
1	16	16	4	4	0	0	0	0	0	0	0	0	0	0	0	0
2	16	4	16	4	4	0	4	0	0	4	0	4	4	0	4	0
3	16	4	4	16	0	0	0	0	0	0	0	0	0	0	0	0
4	16	0	4	0	16	0	4	0	0	0	0	0	0	0	0	0
5	16	0	0	0	0	16	0	0	0	0	8	0	0	0	0	8
6	16	0	4	0	4	0	16	0	0	0	0	0	0	0	0	0
7	16	0	0	0	0	0	0	16	0	0	8	0	0	8	0	0
8	16	0	0	0	0	0	0	0	16	0	0	0	0	0	0	0
9	16	0	4	0	0	0	0	0	0	16	0	4	0	0	0	0
A	16	0	0	0	0	8	0	8	0	0	16	0	0	8	0	8
B	16	0	4	0	0	0	0	0	0	4	0	16	0	0	0	0
C	16	0	4	0	0	0	0	0	0	0	0	0	16	0	4	0
D	16	0	0	0	0	0	0	8	0	0	8	0	0	16	0	0
E	16	0	4	0	0	0	0	0	0	0	0	0	4	0	16	0
F	16	0	0	0	0	8	0	0	0	0	8	0	0	0	0	16

TABLE VII
WARP'S FEISTEL BOOMERANG CONNECTIVITY TABLE

bility as $\hat{q}_j$

f) Update the current boomerang probability with $\hat{p_i}^2\hat{q_j}^2 r_{i,j}$.

g) Add blocking constraints to the SMT model to prevent the current E_1 trail from being found again, then repeat Step 2.

3) If no more valid E_1 trails can be found, clear all blocking constraints for E_1, add constraints to the SMT model to block the current E_0 trail from being found again, then repeat Step 1.

The search itself is generic to Feistel-like ciphers and can be adapted to other design paradigms such as SPN. Several examples of boomerang distinguishers for TWINE and LBlock-s found using the automated search are shown in Table VIII[3]. The best boomerang distinguishers that found for WARP are summarized in Table IX.

C. Related-key Differential Distinguishers

Although its designers do not claim any security in the related-key setting, WARP could possibly be used to design other primitives such as hash functions or used in certain applications for which resilience against related-key attacks are important[4]. We found that WARP has a family of 2-round iterative related-key differential trails:

Property 2 (2-round Related-key Trails). *Let i be an odd-numbered index (1,3,...,29,31) of a nonzero nibble in the input difference and x be the nibble's difference. The input difference α consists of all zero nibbles except $\alpha_i = x$. When $K^1_{\frac{\pi^{-1}(i)}{2}} = y$, $K^0_{\frac{(\pi^{-1})^2(i)-1}{2}} = x$ and $K^0_{\frac{i-1}{2}} = x$, we have a 2-round related-key differential trail from $\alpha \rightarrow \alpha$ with $\mathsf{DP} = \frac{DDT(x,y)}{16}$.*

Depending on the DDT (Table V), these trails can either have a differential probability of $\frac{4}{16} = 2^{-2}$ or $\frac{2}{16} = 2^{-3}$. Figure 4 illustrates two examples of trails described in Property 2. For a more concrete example, we set $i = 3$, $x = 1$ and $y = 2$ and have the following differential propagation that follows the red trail in Figure 4:

$$\texttt{0000...0000000000001000} \xrightarrow[2^{-2}]{2r} \texttt{0000...0000000000001000},$$

where the key difference is $\Delta K = \{\Delta K_1 = \texttt{0000000000200000}, \Delta K_0 = \texttt{0000000010000010}\}$. The trail's differential probability is $\frac{DDT(1,2)}{16} = 2^{-2}$. We can then concatenate this iterative related-key differential trail 20.5 times to construct a 41-round distinguisher $\mathsf{DP} = 2^{-40}$.

IV. DIFFERENTIAL ATTACKS ON WARP

We denote an R-round cipher, E as $E = E_f \circ E' \circ E_b$, where E' is our differential distinguisher. The

[3]These distinguishers only serve to showcase the flexibility of the proposed boomerang search, and may not be the best boomerang distinguishers found for these ciphers.

[4]Other block ciphers such as GIFT have also been extensively cryptanalyzed using related-key attacks despite not claiming any security in this setting [23], [24].

Cipher	R ($\mathsf{R}_{E_0} + \mathsf{R}_{E_m} + \mathsf{R}_{E_1}$)	α	δ	$\sum_{i,j}(\hat{p_i}^2\hat{q_j}^2 r_{i,j})$
TWINE	15 (7+1+7)	3890 0000 0097 0000	0DB0 0010 0D00 0C00	$2^{-58.92}$
TWINE	16 (8+1+7)	A250 0000 0056 0000	A000 0702 0050 0002	$2^{-61.62}$
LBlock-s	15 (7+1+7)	0420 0004 0600 0004	6600 0000 4020 0004	$2^{-58.64}$

TABLE VIII
BOOMERANG DISTINGUISHERS FOR OTHER CIPHERS

R ($R_{E_0} + \mathsf{R}_{E_m} + \mathsf{R}_{E_1}$)	α	δ	p^2q^2r	$\sum_{i,j}(\hat{p_i}^2\hat{q_j}^2 r_{i,j})$
20 (9+1+10)	0000 0000 0000 1000 0000 C2C0 4200 0012	0202 0040 0200 1002 4000 0000 0202 0000	2^{-124}	$2^{-114.24}$
21 (10+1+10)	E000 00EE EE00 0000 00E0 00EE 0000 0000	2000 0000 0104 0000 0404 0020 0100 2004	2^{-142}	$2^{-121.11}$

TABLE IX
BOOMERANG DISTINGUISHERS FOR WARP

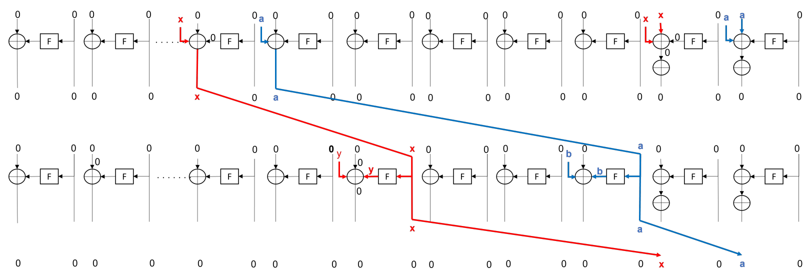

Fig. 4. Two examples of the 2-round iterative related-key differential trails for WARP where $i = 3$ and $i = 1$ are represented by the red and blue trails respectively (Property 2)

R_b-round E_b and R_f-round E_f are rounds added before and after the distinguisher, respectively. The input difference of E_b and the output difference of E_f are denoted as ΔP and ΔC. We denote the number of active or unknown bits of ΔP as r_b while the $n-r_b$ inactive or fixed bits are denoted as $\hat{r}_b$ and $\bar{r}_b$ for 0s and 1s, respectively. Analogously, these bits are denoted as r_b, $\hat{r}_b$ and $\bar{r}_b$ for ΔC. We adopt a targeted approach for the key counting procedure by strategically guessing and filtering m bits of keys involved in subround filters described in Property 1.

A. 23-round Attack using 18-round Differential

We use the 18-round differential from Table VI with $\mathsf{DP} = 2^{-104.62}$ to mount an attack on 23-round WARP by adding 2 rounds at the beginning and 3 rounds at the end. The 23-round key recovery model is depicted in Table X, where we have ($r_b = 56$, $\hat{r}_b = 56$, $\bar{r}_b = 16$) and ($r_f = 72$, $\hat{r}_f = 46$, $\bar{r}_f = 10$). We guess a total of $m = 56$ subkey

R		
	Input Difference (ΔP)	???? A000 A0?? ???A ?F?? 0000 A0A0 ??50
1	After S-box (ΔX^1)	0?0? A000 A07? 0?0A 0F0? 0000 A0A0 0?50
	After π (ΔY^1)	00?7 00?A F000 00?A A0?A 00?A ?500 0000
2	After S-box (ΔX^2)	00A7 000A F000 000A A00A 000A 5500 0000
	After π (ΔY^2)	0000 A0AF 005A 0000 A000 AA75 0000 0000
	Differential distinguisher,	0000 A0AF 005A 0000 A000 AA75 0000 0000
20	$\alpha \to \beta$	000A 5000 0AA0 0000 5A00 000A 0000 0A00
21	After S-box (ΔX^{21})	00?A 5000 ?AA0 0000 ?A00 00?A 0000 ?A00
	After π (ΔY^{21})	00AA A00? A00? 0005 0?00 0?A0 00A0 0?00
22	After S-box (ΔX^{22})	00?A A0?? A0?? 00?5 ??00 ??A0 00A0 ??00
	After π (ΔY^{22})	00?? 000A ??0? 0??A 5?0A ??A0 0000 ?A0?
23	After S-box (ΔX^{23})/ Output Difference (ΔC)	00?? 00?A ???? ???A ???A ??A0 0000 ?A??

Additional Notes: ? denotes an undetermined nibble. Red text denotes subround filters based on Property 1.

TABLE X
THE 23-ROUND KEY RECOVERY MODEL FOR WARP USING AN 18-ROUND DIFFERENTIAL

bits, corresponding to K_i^0 and K_j^1, where $i = \{0, 1, 4, 5, 7, 8, 9, 11, 14\}$ and $j = \{2, 4, 11, 13, 14\}$.

Data Preparation. We let $s = 2$ and collect

$y = 2 \cdot 2^{-56} \cdot \frac{2}{2^{-104.62}} \approx 2^{50.62}$ structures of 2^{56} plaintexts each. The plaintexts traverse all possible values for the active r_b bits while the $\hat{r}_b$ and $\bar{r}_b$ bits are assigned suitable constants. Notably, half of the plaintexts should have the $\bar{r}_b$ bits set to 0 while the other half has these set to 1. We encrypt all 2^{56} plaintexts to obtain 2^{56} corresponding ciphertexts of all structures that are stored in a hash table H, according to the 46 $\hat{r}_f$ bits set to 0. For each pair of structures, we have 2^{111} pairs at the beginning.

Key Recovery. We initialize a list of 2^{56} counters then:

1) We filter wrong pairs using inactive bits of ΔC, leaving $2^{111-56} = 2^{55}$ pairs per structure or $2^{50.62+55} = 2^{105.62}$ pairs in total.
2) The number of filters (Property 1) in the first and final rounds are $v_b = 8$ and $v_f = 6$, respectively as indicated in red in Table X. Thus, the number of valid pairs would be reduced to $2^{105.62-4\times 8-4\times 6} = 2^{49.62}$.
3) **In Round 23:** We can propagate knowledge of the output difference to Round 22 without having to guess any keys (Property 1).
4) **In Round 22:** In this round, guess and filter subkeys for each remaining pair based on Property 1. For example, guess K^0_{14} to partially decrypt ΔX^{23}_{29} to calculate $\Delta X^{22}_2 = \Delta Y^{22}_{29}$, . Directly compute $\Delta X^{22}_3 = \Delta Y^{22}_{14} = \Delta X^{23}_{14}$ from the ciphertext pairs. Check if $\Delta X^{22}_3 \oplus S(\Delta X^{22}_2) = \Delta Y^{21}_3 = 0$. If the equality holds, keep the guessed value of K^0_{14} and the pair, otherwise discard them. There will be around $2^{49.62} \cdot 2^4 \cdot 2^{-4} = 2^{49.62}$ combinations of the remaining pairs associated with the guessed K^0_{14} values. We have 6 more of these filters in Round 22, for which we can guess and filter K^0_1, K^0_4, K^0_5, K^0_7, K^0_8 and K^0_{11} candidates. We expect to have $2^{49.62}$ combinations of the remaining pairs associated with 28-bit key candidates.
5) **In Round 21:** Guess K^1_{14}, K^1_{11} and K^1_4 to calculate ΔX^{21}_2, ΔX^{21}_{14} and ΔX^{21}_{28} respectively, while the remaining differences can be calculated based on the previous key guesses. After going through these filters, there will be $2^{49.62}$ combinations of the remaining pairs associated with 40-bit key candidates. For the remaining two filters, guess (K^0_9,K^1_{13}) to calculate ΔX^{21}_8 and (K^0_0,K^1_2) to calculate ΔX^{21}_{22}. Since there are 2^8 possible subkey candidates involved in each of these 4-bit filters, this will increase the number of combinations to $2^{49.62} \cdot 2^{4\times 2} = 2^{57.62}$ pairs associated with 56-bit keys.
6) **In Round 1:** For all the remaining pairs, we propagate knowledge of the input difference to Round 2.
7) **In Round 2:** Differences ΔY^1_6 and ΔY^1_{24} can be calculated using K^0_0 and K^0_{11} candidates already associated with each remaining pair. ΔY^1_7 and ΔY^1_{25} can be calculated from the plaintext pairs. We then discard combinations of pairs and keys based on the known differences, $\Delta X^2_7 = 5$ and $\Delta X^2_{25} = 0$ due to Property 1. This reduces the number of possible combinations to $2^{57.62-4\times 2} = 2^{49.62}$.
8) Increment the key counters based on the $2^{49.62}$ remaining combinations of pairs associated with the 56 bits of guessed keys. We expect 2 pairs to vote for the right key while the remaining pairs will vote for a random key with a probability of $2^{49.62-56} = 2^{-6.38}$.
9) We select the top $2^{m-a} = 2^{56-52} = 16$ hits in the counter to be candidates that deliver an a-bit or higher advantage [25], then brute-force the 72 remaining bits of the secret key.

Complexity Estimation. The data and memory complexities are $N = 2^{50.62} \cdot 2^{56} \approx 2^{106.62}$ plaintexts and $2^{106.62} + 2^{56} \cdot \frac{56}{128} \approx 2^{106.62}$ 128-bit blocks, respectively. The time complexity of the key recovery is dominated by the final round filtering in Step 2, in which the $2^{105.62}$ pairs need to be partially decrypted. This requires $2^{105.62} \cdot \frac{2}{23} \approx 2^{102.09}$ 23-round WARP encryptions. The brute force complexity is $2^{128-a} = 2^{76}$ Therefore, the time complexity of the 23-round differential attack, including data preparation, is about $2^{106.62} + 2^{102.09} + 2^{76} \approx 2^{106.68}$ 23-round WARP encryptions when $a = 52$.

Success Probability. We calculate the probability of success, $\mathsf{Pr_S}$ of our attack based on the method

R	Input Difference (ΔP)	00?E 0000 E000 ?EE0 ?E?E 0000 E000 00?E
1	After S-box (ΔX^1) After π (ΔY^1)	000E 0000 E000 0EE0 0E0E 0000 E000 000E E000 00EE EE00 0000 00E0 00EE 0000 0000
22	Boomerang distinguisher, $\alpha \rightarrow \delta$	E000 00EE EE00 0000 00E0 00EE 0000 0000 2000 0000 0104 0000 0404 0020 0100 2004
23	After S-box (ΔX^{23}) After π (ΔY^{23})	2000 0000 ?1?4 0000 ?4?4 0020 ?100 20?4 400? 024? 4010 0040 0200 0?0? 0?1? 0200
24	After S-box (ΔX^{24})/ Output Difference (ΔC)	40?? ???? 4010 0040 ?200 ???? ???? ?200

? denotes an undetermined nibble. Red text denotes subround filters based on Property 1.

TABLE XI
24-ROUND KEY RECOVERY MODEL OF THE RECTANGLE ATTACK USING A 21-ROUND (10+1+10) DISTINGUISHER

proposed by Selçuk [25]:

$$\mathsf{Pr_S} = \Phi\left(\frac{\sqrt{s \cdot S_N} - \Phi^{-1}(1 - 2^{-a})}{\sqrt{S_N + 1}}\right), \qquad (7)$$

where the signal-to-noise ratio is calculated as $S_N = \frac{\mathsf{DP}}{2^{-n}}$. With $a = 52$, the probability that the attack succeeds is 92.09%.

B. 24-round Rectangle Attack using 21-round Boomerang Distinguisher

We use the 21-round boomerang distinguisher from Table IX where $\mathsf{R}_{E_0} = 10$, $\mathsf{R}_{E_1} = 10$ and $\sum_{i,j} \hat{p_i}^2 \hat{q_j}^2 r_{i,j}) = 2^{-121.11}$ to mount a rectangle attack on 24-round WARP by appending 1 round at the beginning and 2 rounds at the end. The 24-round key recovery model is depicted in Table XI, which has ($r_b = 20$, $\hat{r}_b = 84$, $\bar{r}_b = 24$) and ($r_f = 60$, $\hat{r}_f = 61$, $\bar{r}_f = 7$). The number of subkey bits that will be guessed are $m_f = 16$, corresponding to K_j^1 where $j = \{2, 8, 10, 15\}$. Details of our rectangle attack are as follows:

Data Preparation. For $s = 2$ right quartets, collect $y = \frac{\sqrt{2} \cdot 2^{64-20}}{\sqrt{2^{-121.11}}} = 2^{105.06}$ structures of 2^{20} plaintexts each. The plaintexts are assigned all possible combinations of the r_b active bits while the other bits are assigned suitable constants. Encrypt 2^{20} plaintexts of each structure to obtain 2^{20} corresponding ciphertexts, which are stored in a hash table, H_1 indexed by the r_b bits of the plaintext.

Key Recovery. Initialize a list of 2^{16} counters then:

1) Construct a set $S = \{(P_1, C_1, P_2, C_2) : E_b(P_1) \oplus E_b(P_2) = \alpha\}$ without having to guess any keys in E_b as follows:
   a) For every plaintext P_1 in a structure, determine the known $\hat{r_b} + \bar{r_b}$ bits in P_2 by calculating $P_2 = P_1 \oplus \Delta P$.
   b) The unknown nibbles in P_2, which are all left input nibbles to Feistel-subrounds, can be calculated from their corresponding right input nibbles. Let the pairs of nibbles for P_1 and P_2 be denoted as (x_L^1, x_R^1) and (x_L^2, x_R^2) respectively (see Figure 3). We already know the values for the right halves (x_R^1, x_R^2) after Step 1(a) and we also know the value of x_L^1 from P_1. We can then calculate the remaining unknown value as
   $$x_L^2 = x_L^1 \oplus S(x_R^1) \oplus S(x_R^2).$$
   c) After calculating all the unknown bits of P_2, check H_1 to find the corresponding plaintext-ciphertext pair indexed by the r_b bits of P_2. Since $v_b = 5$, we expect $2^{20 \times 2 - 1} \cdot 2^{-4 \times v_b} = 2^{19}$ pairs in S.
2) The size of S is $N = 2^{105.06} \cdot 2^{(19+1) \times 2^{-1}} = 2^{115.06}$ chosen plaintexts. Insert S into a hash table H_2 indexed by the 61 $\hat{r}_f$ bits of C_1 and C_2. For each element of S, check H_2 to find (P_1, C_1, P_2, C_2) where (C_1, C_3) and (C_2, C_4) collide in the $\hat{r}_f + \bar{r}_f = 68$ known bits. There will be $(2^{115.06})^2 \cdot 2^{-2 \times 68} = 2^{94.12}$ quartets remaining.
3) Since the number of subround filters is $v_f = 9$, the number of valid quartets would be reduced to $2^{94.12 - 8 \times 9} = 2^{22.12}$. The filtering effect due to Property 1 is twofold since it is applicable to both pairs in a quartet.
4) **In Round 24:** Propagate the knowledge of the output difference to Round 23 (Property 1).
5) **In Round 23:** Perform the guess-and-filter procedure for subkeys in Round 23. For example, ΔX_0^{23} can be directly computed from the ciphertext pairs.Guess K_{15}^1 and partially decrypt (C_1, C_3) and (C_2, C_3) to obtain ΔX_1^{23} for each pair. Check if each pair in the

quartet fulfills $\Delta X_1^{23} \oplus S(\Delta X_0^{23}) = \delta_1 = 0$. If the equality holds for both pairs in the quartet, keep the guessed key and the quartet, otherwise, discard them. There will be around $2^{22.12} \cdot 2^4 \cdot 2^{-8} = 2^{18.12}$ combinations of the remaining quartets associated with the guessed K_{15}^1 values. Guess and filter 3 more subkeys, K_2^1, K_8^1 and K_{10}^1, which leaves $2^{18.12} \cdot 2^{-4\times 3} = 2^{6.12}$ combinations of the remaining quartets associated with the guessed keys.

6) Increment the key counters based on the $2^{6.12}$ remaining combinations of quartets associated with the 16 bits of guessed keys. On average, 2 quartets will vote for the right key while the remaining quartets will vote for a random key with a probability of $2^{6.12-16} = 2^{-7.88}$.
7) Select the top $2^{16-12} = 16$ hits in the counter and brute force the 112 remaining bits of the secret key.

Complexity Estimation. The data complexity of the attack is $N = 2^{105.06} \cdot 2^{20} \approx 2^{125.06}$ chosen plaintexts. The memory complexity includes the space required to store the hash tables and the key counters, which is $2 \cdot 2^{125.06} + 2^{24} \cdot \frac{24}{128} \approx 2^{126.06}$ 128-bit blocks. To prepare the quartets, we require around $2^{125.06}$ 24-round encryptions and $2N$ memory accesses, for which we make a conservative assumption is equivalent to 1 encryption round. The time complexity of the key recovery is dominated by the final round filtering in Step 3, which is approximately $\theta = 2^{94.12} \cdot \frac{4}{24} \approx 2^{91.54}$. The overall time complexity of the 24-round attack is $2^{125.06} + 2 \cdot 2^{125.06} \cdot \frac{1}{24} + 2^{93.54} + 2^{116} \approx 2^{125.18}$ 24-round WARP encryptions when $a = 12$.

Success Probability. When $a = 12$ and $S_N = \frac{\sum_{i,j}(\hat{p_i}^2 \hat{q_j}^2 r_{i,j})}{2^{-n}} = 2^{-121.11}$, the probability that the attack succeeds is 86.2%.

V. Related-key Differential Attacks on WARP

A. 25-round Related-key Differential Attack

We concatenate 9 instances of the 2-round iterative related-key differential trail described in subsection III-C and append 1 more round to form a 19-round distinguisher with DP$= 2^{-18}$. After appending 6 rounds to the end of this distinguisher, we have a 25-round key-recovery model depicted in Table XII where $r = 19$. We guess a total of 16 subkey bits, corresponding to K_i^0 where $i = \{4, 7, 10, 14\}$. Although it may be possible to guess more key bits to reduce the computational complexity of the final brute force step, we stick with 16 bits so we can computationally verify the attack efficiently.

R (ΔK)		
$0 - r$	RK distinguisher, $\alpha \rightarrow \beta$	0000 0000 0000 0000 0000 0000 0000 1000 0000 0000 0000 0000 0000 0100 0000 0000
$r+1$ (ΔK_1)	After S-box (ΔX^{r+1}) After π (ΔY^{r+1})	0000 0000 0000 0000 0000 ?100 0000 0000 0000 0000 0000 0000 0000 0000 0000 100?
$r+2$ (ΔK_0)	After S-box (ΔX^{r+2}) After π (ΔY^{r+2})	0000 0000 0000 0000 1000 0000 0000 00?? ?000 0000 0000 0000 0000 0100 0?00 0000
$r+3$ (ΔK_1)	After S-box (ΔX^{r+3}) After π (ΔY^{r+3})	?000 0000 0000 0000 0000 ?100 ??00 0000 0000 0?00 00?0 0000 0000 000? 0000 100?
$r+4$ (ΔK_0)	After S-box (ΔX^{r+4}) After π (ΔY^{r+4})	0000 ??00 00?0 0000 1000 00?? 0000 00?? ?00? ?000 0000 ?00? 0000 0100 0?00 0?00
$r+5$ (ΔK_1)	After S-box (ΔX^{r+5}) After π (ΔY^{r+5})	?0?? ?000 0000 ?0?? 0000 ?100 ??00 ??00 0??0 0?00 0??? 000? ??00 00?? 0000 100?
$r+6$ (ΔK_0)	After S-box (ΔX^{r+6}) Output Difference (ΔC)	???0 ??00 ???? 00?? ??00 00?? 0000 00??

? denotes an undetermined nibble. Red text denotes subround filters based on Property 1.

TABLE XII
Key recovery model using an r-round related key distinguisher where $r \in \{3, 5, ..., 31\}$

Data Preparation. Encrypt 2^{19} pairs of plaintexts, (P_1, P_2) using a pair of related keys, $(K, K \oplus \Delta K)$. We expect $s = 2$ right pairs. There is a strong filtering effect at the output difference ΔC, which has 60 inactive bits and 5 subround filters (Property 1). The probability of a wrong pair surviving is $2^{20} \cdot 2^{-60} \cdot 2^{-4\times 5} = 2^{-60}$, which implies that only the right pairs remain.

Key Recovery. For all the remaining (right) pairs:

1) **In Round 25:** Propagate knowledge of the output difference to Round 24 without having to guess any keys (Property 1).
2) **In Round 24:** Guess K_{14}^0 to calculate 2_2^{24}, then derive 2_3^{24} from the ciphertext pairs. Since all pairs are valid, each pair will be associated with at least one possible 4-bit key

candidate. Guess K^0_{10}, K^0_7 and K^0_4 associated with the remaining subround filters.

Complexity Estimation. The data complexity of the attack is $N = 2 \cdot 2^{19} = 2^{20}$ chosen plaintexts and it finds correctly 16-bits of the key, which is sufficient for verification of the attack correctness. The memory requirement of the attack is negligible.

Computational Verification. We first calculated the average probability of the 19-round distinguisher. Using 10 randomly selected keys and plaintext pairs, the average differential probability was $2^{18.1}$. We then execute the 25-round attack 10 times on a PC with an Intel Core i7-9700K 3.60GHz processor and 32GB of RAM. The correct subkey always has the highest count of s. The correct 16-bit key will be among the top 2 candidates 70% of the time. On average, the attack completes in under 2.5 minutes using an unoptimized implementation.

B. 41-round (Full) Related-key Differential Attack

The key recovery model for a 41-attack using a 35-round distinguisher with $\mathsf{DP} = 2^{-34}$ is depicted in Table XII where $r = 35$. We generate 2^{35} pairs and expect 2 right pairs. From Round 40 to Round 36, we guess a total of 60 subkey bits (16 in Round 40, 12 in Round 39, 16 bits in 38, 12 in Round 37 and 4 in Round 36). There are subround filters in Rounds 37 and 38 that require guessing 12 key bits, key counters that can accommodate 2^{12} possibilities are required. The memory requirement is $(6 \cdot 2^4 \cdot \frac{4}{128} + 2 \cdot 2^{12} \cdot \frac{12}{128}) \approx 2^{9.59}$ 128-bit blocks. The time complexity for the guess-and-determine procedure is negligible, therefore recovering 60 bits of the key comes mainly from encrypting the 2^{36} chosen plaintexts. We can either brute force the remaining 68 bits, which would then dominate the time complexity or use faster auxiliary techniques to find the rest of the key.

VI. Conclusion

This paper described cryptanalytic attacks on the lightweight block cipher WARP. We show that its first 20 rounds have high-probability differentials due to a strong differential effect. Then, by using an automatic search for boomerang distinguishers, boomerang distinguishers for up to 21 rounds were found. We also described a family of 2-round iterative related-key trails which can be concatenated to form a full 41-round distinguisher for WARP. Key-recovery attacks were then demonstrated using the identified distinguishers. In the single-key setting, we attacked 23 and 24 rounds of WARP using an 18-round differential and 21-round boomerang distinguisher, respectively. Next, we computationally verified a 25-round related-key attack on WARP using a 19-round distinguisher, where 16 subkey bits were recovered in 2.5 minutes. Using the same framework, a practical related-key attack on the full WARP was introduced. All attack complexities were summarized in Table I. To the best of our knowledge, these are the first (valid) 3rd party cryptanalysis results for WARP.

Acknowledgements

Je Sen Teh was supported by the Luxembourgish Fonds National de la Recherche (FNR) and the German Research Foundation (DFG) project APLICA (C19/IS/13641232).

References

[1] A. Fotovvat, G. M. E. Rahman, S. S. Vedaei, and K. A. Wahid, “Comparative performance analysis of lightweight cryptography algorithms for iot sensor nodes,” *IEEE Internet Things J.*, vol. 8, no. 10, pp. 8279–8290, 2021. [Online]. Available: https://doi.org/10.1109/JIOT.2020.3044526

[2] S. Banik, A. Bogdanov, T. Isobe, K. Shibutani, H. Hiwatari, T. Akishita, and F. Regazzoni, “Midori: A block cipher for low energy,” in *ASIACRYPT 2015, Part II*, ser. LNCS, T. Iwata and J. H. Cheon, Eds., vol. 9453. Springer, Heidelberg, Nov. / Dec. 2015, pp. 411–436.

[3] S. Banik, S. K. Pandey, T. Peyrin, Y. Sasaki, S. M. Sim, and Y. Todo, “GIFT: A small present - towards reaching the limit of lightweight encryption,” in *CHES 2017*, ser. LNCS, W. Fischer and N. Homma, Eds., vol. 10529. Springer, Heidelberg, Sep. 2017, pp. 321–345.

[4] Y. Zheng, T. Matsumoto, and H. Imai, “On the construction of block ciphers provably secure and not relying on any unproved hypotheses,” in *CRYPTO’89*, ser. LNCS, G. Brassard, Ed., vol. 435. Springer, Heidelberg, Aug. 1990, pp. 461–480.

[5] S. Banik, Z. Bao, T. Isobe, H. Kubo, F. Liu, K. Minematsu, K. Sakamoto, N. Shibata, and M. Shigeri, “WARP : Revisiting GFN for lightweight 128-bit block cipher,” in *Selected Areas in Cryptography - SAC 2020 - 27th International Conference, Halifax, NS, Canada, October 21-23, 2020*, ser. Lecture Notes in Computer Science,

O. Dunkelman, M. J. J. Jr., and C. O'Flynn, Eds., vol. 12804. Springer, 2020, pp. 535–564. [Online]. Available: https://doi.org/10.1007/978-3-030-81652-0_21

[6] T. Suzaki, K. Minematsu, S. Morioka, and E. Kobayashi, "TWINE : A lightweight block cipher for multiple platforms," in *SAC 2012*, ser. LNCS, L. R. Knudsen and H. Wu, Eds., vol. 7707. Springer, Heidelberg, Aug. 2013, pp. 339–354.

[7] M. Kumar and T. Yadav, "MILP based differential attack on round reduced WARP," Cryptology ePrint Archive, Report 2020/1598, Version 20210802:090929, 2020, https://ia.cr/2020/1598.

[8] H. Boukerrou, P. Huynh, V. Lallemand, B. Mandal, and M. Minier, "On the feistel counterpart of the boomerang connectivity table introduction and analysis of the FBCT," *IACR Trans. Symmetric Cryptol.*, vol. 2020, no. 1, pp. 331–362, 2020. [Online]. Available: https://doi.org/10.13154/tosc.v2020.i1.331-362

[9] Stefan Kölbl, "CryptoSMT: An easy to use tool for cryptanalysis of symmetric primitives," https://github.com/kste/cryptosmt.

[10] L. Zhang, W. Wu, Y. Wang, S. Wu, and J. Zhang, "LAC," Tech. Rep., 2014, available at https://competitions.cr.yp.to/round1/lacv1.pdf.

[11] X. Lai, J. L. Massey, and S. Murphy, "Markov ciphers and differential cryptanalysis," in *EUROCRYPT'91*, ser. LNCS, D. W. Davies, Ed., vol. 547. Springer, Heidelberg, Apr. 1991, pp. 17–38.

[12] M. Matsui, "On correlation between the order of S-boxes and the strength of DES," in *EUROCRYPT'94*, ser. LNCS, A. D. Santis, Ed., vol. 950. Springer, Heidelberg, May 1995, pp. 366–375.

[13] N. Mouha, Q. Wang, D. Gu, and B. Preneel, "Differential and linear cryptanalysis using mixed-integer linear programming," in *Information Security and Cryptology - 7th International Conference, Inscrypt 2011, Beijing, China, November 30 - December 3, 2011.*, ser. Lecture Notes in Computer Science, C. Wu, M. Yung, and D. Lin, Eds., vol. 7537. Springer, 2011, pp. 57–76. [Online]. Available: https://doi.org/10.1007/978-3-642-34704-7_5

[14] N. Mouha and B. Preneel, "Towards finding optimal differential characteristics for ARX: Application to Salsa20," Cryptology ePrint Archive, Report 2013/328, 2013, https://eprint.iacr.org/2013/328.

[15] G. Tsudik and E. Van Herreweghen, "On simple and secure key distribution," in *ACM CCS 93*, D. E. Denning, R. Pyle, R. Ganesan, R. S. Sandhu, and V. Ashby, Eds. ACM Press, Nov. 1993, pp. 49–57.

[16] E. Biham, O. Dunkelman, and N. Keller, "A unified approach to related-key attacks," in *FSE 2008*, ser. LNCS, K. Nyberg, Ed., vol. 5086. Springer, Heidelberg, Feb. 2008, pp. 73–96.

[17] J. Kelsey, T. Kohno, and B. Schneier, "Amplified boomerang attacks against reduced-round MARS and Serpent," in *FSE 2000*, ser. LNCS, B. Schneier, Ed., vol. 1978. Springer, Heidelberg, Apr. 2001, pp. 75–93.

[18] E. Biham, O. Dunkelman, and N. Keller, "New results on boomerang and rectangle attacks," in *FSE 2002*, ser. LNCS, J. Daemen and V. Rijmen, Eds., vol. 2365. Springer, Heidelberg, Feb. 2002, pp. 1–16.

[19] A. Biryukov and D. Khovratovich, "Related-key cryptanalysis of the full AES-192 and AES-256," in *ASIACRYPT 2009*, ser. LNCS, M. Matsui, Ed., vol. 5912. Springer, Heidelberg, Dec. 2009, pp. 1–18.

[20] O. Dunkelman, N. Keller, and A. Shamir, "A practical-time related-key attack on the KASUMI cryptosystem used in GSM and 3G telephony," in *CRYPTO 2010*, ser. LNCS, T. Rabin, Ed., vol. 6223. Springer, Heidelberg, Aug. 2010, pp. 393–410.

[21] C. Cid, T. Huang, T. Peyrin, Y. Sasaki, and L. Song, "Boomerang connectivity table: A new cryptanalysis tool," in *EUROCRYPT 2018, Part II*, ser. LNCS, J. B. Nielsen and V. Rijmen, Eds., vol. 10821. Springer, Heidelberg, Apr. / May 2018, pp. 683–714.

[22] A. Biryukov and I. Nikolic, "Complementing Feistel ciphers," in *FSE 2013*, ser. LNCS, S. Moriai, Ed., vol. 8424. Springer, Heidelberg, Mar. 2014, pp. 3–18.

[23] B. Zhao, X. Dong, W. Meier, K. Jia, and G. Wang, "Generalized related-key rectangle attacks on block ciphers with linear key schedule: applications to SKINNY and GIFT," *Des. Codes Cryptogr.*, vol. 88, no. 6, pp. 1103–1126, 2020. [Online]. Available: https://doi.org/10.1007/s10623-020-00730-1

[24] R. Zong, X. Dong, H. Chen, Y. Luo, S. Wang, and Z. Li, "Towards key-recovery-attack friendly distinguishers: Application to GIFT-128," *IACR Trans. Symmetric Cryptol.*, vol. 2021, no. 1, pp. 156–184, 2021. [Online]. Available: https://doi.org/10.46586/tosc.v2021.i1.156-184

[25] A. A. Selçuk, "On probability of success in linear and differential cryptanalysis," *Journal of Cryptology*, vol. 21, no. 1, pp. 131–147, Jan. 2008.