Improved Differential Meet-in-the-Middle Cryptanalysis

Zahra Ahmadian[1(✉)], Akram Khalesi[1], Dounia M'Foukh[2], Hossein Moghimi[1], and María Naya-Plasencia[2]

[1] Shahid Beheshti University, Tehran, Iran
{z_ahmadian,a_khalesi}@sbu.ac.ir, h.moghimi@mail.sbu.ac.ir
[2] Inria, Paris, France
{dounia.mfoukh,maria.naya_plasencia}@inria.fr

Abstract. In this paper, we extend the applicability of differential meet-in-the-middle attacks, proposed at Crypto 2023, to truncated differentials, and in addition, we introduce three new ideas to improve this type of attack: we show how to add longer structures than the original paper, we show how to improve the key recovery steps by introducing some probability in them, and we combine this type of attacks with the state-test technique, that was introduced in the context of impossible differential attacks. Furthermore, we have developed a MILP-based tool to automate the search for a truncated differential-MITM attack with optimized overall complexity, incorporating some of the proposed improvements. Thanks to this, we can build the best known attacks on the cipher `CRAFT`, reaching 23 rounds against 21 previously; we provide a new attack on 23-round `SKINNY-64-192`, and we improve the best attacks on `SKINNY-128-384`.

Keywords: differential meet-in-the-middle cryptanalysis · truncated cryptanalysis · parallel partitioning · state-test technique · tool · `SKINNY` · `CRAFT`

1 Introduction

Symmetric cryptanalysis is crucial for trusting symmetric primitives: the more we cryptanalyze a primitive without success, the more confidence we will have in it. It is essential for determining the security margin of the cryptographic constructions and being able to anticipate problems. Many different cryptanalysis families exist, like differential attacks [9], linear attacks [25], meet-in-the-middle attacks [13], invariant sub-space attacks [23], integral attacks [21]. Often new techniques and variants are proposed to augment these attacks, but proposing new families is less common.

In [11] a new cryptanalysis was proposed: the differential meet-in-the-middle (MITM) attack. It combines ideas from differential and MITM attacks, allowing to build the best known attack on the cipher `SKINNY-128-384` in the single-tweak setting [11]. As described by the authors, these attacks could be seen in

part as a new way of performing the key recovery in differential attacks, or as MITM ones where instead of looking for a partial collision at some middle state, we look for states that verify, with certain probability, a differential relation. The authors proposed in the SKINNY-128-384 scenario, a way of extending the attack using parallel partitioning to gain one round, techniques usually applied in MITM attacks [2,3,10], but not applicable during classical differential attacks. Some questions were left unsolved in the original paper, such as whether these attacks could be seen as just a new way of performing the key recovery part, but were in essence differential attacks. Another interesting question was if they could be combined with truncated differential attacks – as the probability in both directions of a truncated path is often not symmetric, it seemed at first counter-intuitive to apply it. We have considered and answered these two questions, and in addition proposed two additional improvements to the technique: allowing some probability in the key-guessing part, and combining it with the state-test technique, introduced in [12] in the context of impossible differential attacks.

We have applied our new techniques to CRAFT [8], SKINNY-128-384 [7] and SKINNY-64-192 [7], providing the best known attacks in the two first cases, and an attack reaching the highest number of rounds as the best attack for the third. These attacks can be seen in Table 1.

This paper is organized as follows: Sect. 2 presents the previous framework of differential MITM attacks. Section 3 describes our proposition of combining this attack with truncated differentials and Sect. 4 the newly proposed improvements. Section 5 presents our new tool that finds the distinguishers providing the best overall attacks, considering in addition most of the new improvements on the external rounds, and Sect. 6 and Sect. 7 describe the new applications. The paper is ended with a conclusion.

2 Preliminaries: Differential Meet-in-the-Middle

The differential meet-in-the-middle (MITM) technique, introduced in [11], represents a novel approach for the cryptanalysis of symmetric primitives. This attack combines two significant families of symmetric cryptanalysis attacks: the meet-in-the-middle attack and the differential attack. In [11], it is described as both an extension of classical MITM attacks and as a new key recovery method to apply in differential cryptanalysis.

This new technique was successfully applied to SKINNY-128-384, the 128-bit block cipher variant with a 384-bit tweakey, breaking 25 out of the 56 rounds in the single-tweakey setting [11]. This application highlights the attack's efficiency by surpassing the best known attack on this cipher by two rounds. Furthermore, another instance provided in [11], involved AES-256, where this technique managed to break 12 rounds of the cipher in the related-key model. This attack requires only 2 related keys while the previous attacks with the same number of related keys achieve a maximum of 10 rounds.

In this section, we will provide an overview of how differential MITM attacks are constructed, along with the original improvements introduced in [11]. One

Table 1. Summary of the best known cryptanalysis on `CRAFT`, `SKINNY-64-192` and `SKINNY-128-384` in the single tweak setting (related tweak the strongest setting).

Cipher	Rounds	Time	Data	Memory	Attack	Setting/Model	Ref
`CRAFT`	18	$2^{101.7}$	$2^{60.92}$	2^{84}	Rectangle	STK	[17]
	19	$2^{114.68}$	2^{56}	2^{109}	DS-MITM	STK,CP	[24]
	19	$2^{112.61}$	$2^{60.92}$	2^{72}	Rectangle	SK	[30]
	20	$2^{126.96}$	2^{56}	2^{109}	DS-MITM	STK,CP	[24]
	21	$2^{106.53}$	$2^{60.99}$	2^{100}	ID	STK,CP	[18]
	21	$\mathbf{2^{116}}$	$\mathbf{2^{56}}$	$\mathbf{2^{68}}$	**Tr-Diff-MITM**	**STK**	**Sect. 6**
	22	$\mathbf{2^{125}}$	$\mathbf{2^{58}}$	$\mathbf{2^{72}}$	**Tr-Diff-MITM**	**STK**	**Sect. 6**
	23	$\mathbf{2^{125}}$	$\mathbf{2^{60}}$	$\mathbf{2^{68}}$	**Tr-Diff-MITM**	**STK**	**Sect. 6**
`SKINNY-64-192`	21	$2^{180.01}$	2^{44}	$2^{191.55}$	DS-MITM	STK	[29]
	21	$2^{174.42}$	$2^{62.43}$	2^{168}	ID	STK/CP	[18]
	22	$2^{183.97}$	$2^{47.84}$	$2^{74.84}$	ID	STK	[31]
	23	2^{188}	2^{52}	2^{4}	MITM	STK	[14]
	23	2^{184}	2^{60}	2^{8}	MITM	STK	[6]
	23	2^{188}	2^{28}	2^{4}	MITM	STK	[6]
	23	$\mathbf{2^{188}}$	$\mathbf{2^{56}}$	$\mathbf{2^{104}}$	**Tr-Diff-MITM**	**STK**	**Sect. 7.1**
`SKINNY-128-384`	23	2^{376}	2^{104}	2^{8}	MITM	STK	[14]
	23	2^{372}	2^{96}	$2^{352.46}$	DS-MITM	STK	[29]
	23	$2^{361.9}$	2^{117}	$2^{118.5}$	Diff-MITM	STK	[11]
	24	$2^{361.9}$	2^{117}	2^{183}	Diff-MITM	STK	[11]
	24	$2^{372.5}$	$2^{122.3}$	$2^{123.8}$	Diff-MITM	STK	[11]
	25	$2^{372.5}$	$2^{122.3}$	$2^{188.3}$	Diff-MITM	STK	[11]
	25	$\mathbf{2^{378.9}}$	$\mathbf{2^{117}}$	$\mathbf{2^{165}}$	**Diff-MITM**	**STK**	**Sect. 7.2**
	25	$\mathbf{2^{366}}$	$\mathbf{2^{122.3}}$	$\mathbf{2^{188.3}}$	**Diff-MITM**	**STK**	**Sect. 7.2**

MITM: Meet In The Middle	ID: Impossible Differential
CP: Chosen Plaintext	DS-MITM: Demirci-Selcuk
Diff-MITM: Differential MITM	Tr-Diff-MITM: Truncated Differential MITM
STK: Single-Tweak/Tweakey	SK: Single Key

improvement, achieved through a parallel treatment of data partitions, extends the attack for one round more, mostly at no cost, particularly when the key was exclusively added to half of the internal state. Another one is a technique designed to reduce data complexity in the originally full code-book attack scenario.

2.1   Framework of the Differential MITM Attack

Consider an n-bit cipher E decomposed into three sub-ciphers $E_{out} \circ E_m \circ E_{in}$ of r_{in}, r_m and r_{out} rounds respectively, as depicted in Fig. 1. Finally, suppose that the efficient differential $(\alpha \rightarrow \beta)$, covering the r_m middle rounds, with probability 2^{-p} is a distinguisher for E_m. The core idea of the attack involves employing the MITM approach. In other words, given a pair of plaintext/ciphertext (P, C), we independently compute the candidate plaintexts and ciphertexts $\tilde{P}$ and $\tilde{C}$ such that they follow the differential characteristic as summarized on Fig. 1. $\tilde{P}$ is computed from the plaintext P and the difference α for each possible value of the associated key k_{in}, and $\tilde{C}$ from the ciphertext C and the difference β for each possible value of the associated key k_{out}.

Fig. 1. Framework of the differential meet-in-the-middle attack.

Thus the aim of the following procedure is to find a pair of plaintext/ciphertext (P, C) and $(\tilde{P}, \tilde{C})$ such that:

$$\begin{cases} E_{in}(P) \oplus E_{in}(\tilde{P}) &= \alpha \\ E_{out}^{-1}(C) \oplus E_{out}^{-1}(\tilde{C}) &= \beta. \end{cases} \tag{1}$$

Procedure. First, we randomly pick a plaintext P and ask the encryption oracle for its ciphertext C.

– *Upper Part:* From P, α and some key information, we aim to compute $\tilde{P}$ such that $E_{in}(P) \oplus E_{in}(\tilde{P}) = \alpha$, if the key guess is correct. Thus, we want to minimize the amount of key information, denoted by k_{in}, needed. For each possible value i of k_{in}, we have a different candidate $\tilde{P}^i$. Thus, we have $2^{|k_{in}|}$ candidates at the end of this step. For each $\tilde{P}^i$, we ask the oracle for its encryption $\hat{C}^i$, and store them in a hash table.
– *Lower Part:* Similarly, given C, β and a minimized amount of key information, denoted by k_{out}, we can compute $\tilde{C}^j$ such that $E_{out}^{-1}(C) \oplus E_{out}^{-1}(\tilde{C}^j) = \beta$, if the key guess is correct. We have $2^{|k_{out}|}$ candidates for $\tilde{C}^j$.

Actually, during the procedure before the upper and lower parts, we can first guess the subkey bits of k_{in} and k_{out} in common thanks to the possible linear relations between k_{in} and k_{out} given by the key schedule.

Furthermore, given that $P(\alpha \rightarrow \beta) = 2^{-p}$, we have to repeat the procedure 2^p times using 2^p different plaintext/ciphertexts pairs (P_l, C_l) to have a good pair $(P_l, \tilde{P}_l^i)$ and $(C_l, \tilde{C}_l^j)$, satisfying the distinguisher. When this is the case, we will get a collision between a $\hat{C}^i = E(\tilde{P}^i)$ of the upper part and a $\tilde{C}^j$ of the lower part. Each collision (i, j) yields a candidate key.

For each P_l, we initially have $2^{|k_{in}|+|k_{out}|}$ candidate pairs $(\hat{C}_l^i, i)$ and $(\tilde{C}_l^j, j)$ in search of a collision. After matching through the relation $\hat{C}_l^i = \tilde{C}_l^j$, we are left with $2^{|k_{in}|+|k_{out}|-n}$ candidates. Thus, in the end, for each P_l, the number of expected collisions would be $2^{|k_{in}|+|k_{out}|-n-|k_{in}\cap k_{out}|}$.

Complexity. The time complexity for the computations in the upper and lower parts of the procedure is $2^{|k_{in}|+p}$ and $2^{|k_{out}|+p}$, respectively. And as explained above, the number of expected candidate keys is $2^{|k_{in}|+|k_{out}|-n-|k_{in}\cap k_{out}|+p}$.

Thus, the time complexity is:

$$2^p \times 2^{|k_{in} \cap k_{out}|} (2^{|k_{in}|-|k_{in} \cap k_{out}|} + 2^{|k_{out}|-|k_{in} \cap k_{out}|} + 2^{|k_{in}|+|k_{out}|-n-2|k_{in} \cap k_{out}|}).$$

With this, we recover $k_{in} \cup k_{out}$. So, if we expect fewer key candidates than the whole key space $\mathcal{K}$ of size 2^k, then we can guess the remaining bits of the master key and test the guess with additional pairs. Thus, we recover the whole key with a complexity smaller than the cost of an exhaustive key search, and the additional cost of $2^{k-|k_{in} \cup k_{out}|} \times \max(1, 2^{|k_{in}|+|k_{out}|-n-|k_{in} \cap k_{out}|+p})$ to be added to the time complexity $\mathcal{T}$. In the case that we need to guess the remaining bits of the master key, specifically if $|k_{in}| + |k_{out}| - n - |k_{in} \cap k_{out}| + p > 0$, the total time complexity would be:

$$\mathcal{T} = 2^{p+|k_{in} \cap k_{out}|} (2^{|k_{in}|-|k_{in} \cap k_{out}|} + 2^{|k_{out}|-|k_{in} \cap k_{out}|} + 2^{|k_{in}|+|k_{out}|-n-2|k_{in} \cap k_{out}|})$$
$$+ 2^{k-n+p}. \tag{2}$$

The data complexity is $\mathcal{D} = \min(2^n, 2^{p+\min(|k_{in}|,|k_{out}|)})$, and the memory complexity is $\mathcal{M} = 2^{\min(|k_{in}|-|k_{in} \cap k_{out}|, |k_{out}|-|k_{in} \cap k_{out}|)}$.

2.2   Improvement: Parallel Partitions for Layers with Partial Subkeys

In [11], two methods for improving the original differential MITM attack are proposed. The first method, elaborated upon below, focuses on adding an extra round at the beginning or the end of the attack, in some specific cases.

In the cases where the round key addition only affects a part of the state, as is the case with SKINNY [7] or GIFT [5] block ciphers, the differential MITM attack can be extended by one round. Suppose $m < n$ bits of the state are affected by the round key addition. The framework of the improvement is given in Fig. 2, where one round is added at the end of the attack covering $r-1$ rounds of the cipher and we have eliminated all the transformations after the round key addition of round r, if any. Let A and B denote the states before and after adding K_r, respectively. The main idea of the improvement is to only keep the ciphertexts that satisfy the following condition: we fix the $n - m$ bits that are not affected by the key addition in A and B, and compute all the 2^m possible values for A and B. Now, we will repeat the procedure 2^{p-m} times. So, we can apply this improvement without increasing the time complexity, if $p > m$.

Then, from all the 2^m possible values for A, we can proceed with the lower part of the differential MITM attack. We get $2^{|k_{out}|+m}$ possible candidates $(A, \tilde{A}, j)$. In parallel, we do the same for each of the 2^m possible values for B, and we proceed with the upper part of differential MITM attack and we get $2^{|k_{in}|+m}$ possible candidates $(B, \tilde{B}, i)$; hence $2^{|k_{in}|+|k_{out}|+2m}$ total candidates which is a factor 2^m higher than before, as we are comparing this to performing the attack without structures 2^m times. Nevertheless, note that we have to match A and B and their associated pairs $\tilde{A}$ and $\tilde{B}$ through the relation $A \oplus B = \tilde{A} \oplus \tilde{B}$. This also yields the value of K_r, already determined by k_{in} and k_{out}, usually.

Fig. 2. Framework of the parallel partitioning improvement.

So, the number of expected collisions does not increase: in total, we need to collide on $n - m$ bits of the key-free parts of $\tilde{A}$ and $\tilde{B}$; and we have to collide for both pairs in the m bits once the key is added, providing a total filtering of $(n - m) + 2m = n + m$. As we now have 2^m additional potential candidates, but a 2^{m+n} filter, we expect a total proportion of remaining candidates of 2^{-n}, as in the attack without parallel partitions. In [11], they applied this technique to reach one more round of SKINNY-128-384 without increasing the time complexity and thus mounting the best known attack on this cipher in the single tweak setting.

2.3 Reducing Data Needed with Imposed Conditions

The idea is to fix x bits of the plaintext P and of the associated plaintext $\tilde{P}$ to a specific value, thereby restricting them to a set of 2^{n-x} plaintexts instead of the whole codebook. We can choose the plaintext P as desired but the probability to get a $\tilde{P}$ with the fixed x bits is 2^{-x}. Then, the probability of the attack decreases by 2^x as we have to repeat the procedure 2^x times to get a pair that satisfies the differential characteristic along with this new condition. On the other hand, it means that during the upper and lower part of the attack, we will discard a proportion of 2^x tuples of $(P, \tilde{P}, i)$ and $(C, \tilde{C}, j)$, the ones that do not satisfy the conditions. Thus the data and memory complexities will decrease by 2^x.

Furthermore, considering the attributes of the differential MITM attack, we can derive the following two bounds on the number of fixed bits x:

$$p + x \leq n - x \text{ and } 2^{p+x}(2^{|k_{in}|} + 2^{|k_{out}|}) < 2^k.$$

This technique particularly applies when the whole codebook would be needed, as it is in the differential MITM attack on SKINNY-128-384 presented in [11].

3 Truncated Differential Meet-in-the-Middle Attack

Truncated differential cryptanalysis was introduced at FSE in 1994 [20] by Knudsen as an extension of differential cryptanalysis and has proven its efficacy by

successfully attacking several ciphers which seemed to be secure against differential attacks. It is the case with the KLEIN block cipher, for which its security against bit-wise differential attack had been proved but was broken by some truncated differential attacks [22,27]. Thus the main extension of the differential MITM attack that we will explain in this section is the truncated differential MITM attack. A challenge of building this extension is that since the probability of a truncated differential characteristic is not the same in both directions, then it was not clear how to properly take this propriety into account in the truncated differential MITM extension.

The main idea of truncated differential-MITM attacks is to use, instead of a differential path, a truncated differential path as the underlying distinguisher of the attack. As stated in Appendix A of the longer version of this paper [1], a truncated differential operates based on sets of input and output differences rather than concrete ones. It considers whether a word (typically of the S-box size in the cipher) has a non-zero difference or not, regardless of its concrete value. One advantage of the truncated differential attack is that during the key recovery step, we do not need to know the concrete values of the states just before and after the distinguisher. Consequently, we may need to guess fewer subkey words, potentially allowing us to reach more rounds. Additionally, for certain ciphers, truncated differential distinguishers can reach more rounds than concrete differentials (as in [26]). Finally, the search space for truncated differentials is much smaller than that of concrete differentials, making it easier to deal with an automated method [26].

3.1   Framework of the Truncated Differential MITM Attack

Similar to the differential MITM attack, consider the n-bit cipher E decomposed into three sub-ciphers: $E_{out} \circ E_m \circ E_{in}$, of r_{in}, r_m and r_{out} rounds respectively, as depicted in Fig. 3. Finally, suppose that $(\Delta_{in} \xrightarrow{E_m} \Delta_{out})$ is a truncated distinguisher for E_m with the probability of 2^{-p}, where $|\Delta_{in}| = 2^{\delta_{in}}$ and $|\Delta_{out}| = 2^{\delta_{out}}$. So $(\Delta_{in} \xrightarrow{E_m} \Delta_{out})$ is an efficient differential if $p < n - \delta_{out}$.

We randomly pick a pair of plaintext/ciphertext pair (P, C), and try to generate appropriate candidates for $(\tilde{P}, \tilde{C})$ such that the difference of these two data ensures the truncated difference Δ_{in} at round r_{in}, i.e. $E_{in}(P) \oplus E_{in}(\tilde{P}) \in \Delta_{in}$, and the truncated difference Δ_{out} at round $r_{in} + r_m$, i.e. $E_{out}^{-1}(C) \oplus E_{out}^{-1}(\tilde{C}) \in \Delta_{out}$. The procedures of the upper and lower parts of the attack are as follows.

Upper Part: To generate $\tilde{P}$, we guess some key material, denoted by k_{in}. For each P, and each guess i of k_{in}, there are $2^{\delta_{in}}$ distinct candidates $\tilde{P}_l^i$, each corresponding to an input difference $l \in \Delta_{in}$. All the $2^{|k_{in}|+\delta_{in}}$ ciphertexts $\hat{C}_l^i = E_K(\tilde{P}_l^i)$, along with its associated key material i are stored in a hash table.

Fig. 3. Framework of the truncated differential meet-in-the-middle attack. The probability of the distinguisher in the forward (backward) direction is 2^{-p} ($2^{-p'}$).

Lower Part: In the ciphertext side, given C, for each guess j of k_{out}, we compute all the $2^{\delta_{out}}$ ciphertexts corresponding to the $2^{\delta_{out}}$ possible differences $m \in \Delta_{out}$. So, there would be in total $2^{|k_{out}|} \times 2^{\delta_{out}}$ values for $\tilde{C}_m^j$.

Number of Pairs: For the correct key guess, the transition $(\Delta_{in} \xrightarrow{E_m} \Delta_{out})$ will happen with probability 2^{-p}. So, a number of 2^p pairs of data is expected to be tested to find the correct key. For each P and a fixed key material i, we can provide $2^{\delta_{in}}$ differential pairs $(P, \tilde{P})$, each of which corresponds to a specific value of Δ_{in}. So, it is required to repeat the upper and lower parts of the attack for a number of $2^{p-\delta_{in}}$ plaintexts P. On the other hand, despite the concrete differential-MITM attack, the output difference Δ_{out} does not have a specific single value, but it belongs to a set of size $2^{\delta_{out}}$, whole of which should be checked to certainly determine if the event $(\Delta_{in} \xrightarrow{E_m} \Delta_{out})$ has occurred or not.

3.2 Attack Complexities

As done in the differential MITM attack, we first guess the possible linear relations between k_{in} and k_{out}, i.e. $k_{in} \cap k_{out}$. During the upper part of the attack, for each guess of $k_{in} - k_{in} \cap k_{out}$, we get $2^{\delta_{in}}$ candidates $\tilde{P}_l^i$ hence $2^{|k_{in}|+\delta_{in}-|k_{in}\cap k_{out}|}$ candidates for $(P, \tilde{P}, i)$. Similarly, in the lower part, we have $2^{|k_{out}|+\delta_{out}-|k_{in}\cap k_{out}|}$ candidates triplets for $(C, \tilde{C}, j)$. So, there is $2^{|k_{in}|+\delta_{in}+|k_{out}|+\delta_{out}-2|k_{in}\cap k_{out}|}$ candidates for (i, j). Let's denote $\hat{C} = E(\tilde{P})$. The matching of $\hat{C}_l^i$ with $\tilde{C}_m^j$ leaves us with $2^{|k_{in}|+\delta_{in}+|k_{out}|+\delta_{out}-2|k_{in}\cap k_{out}|-n}$ candidates. Moreover, similar to the original differential-MITM attack, we can guess the remaining bits of the master key and test the guess with additional pairs. Thus, the time complexity of the attack is:

$$\mathcal{T} = 2^{p-\delta_{in}} \times 2^{|k_{in}\cap k_{out}|}(2^{|k_{in}|+\delta_{in}-|k_{in}\cap k_{out}|} + 2^{|k_{out}|+\delta_{out}-|k_{in}\cap k_{out}|})$$
$$+ 2^{p-\delta_{in}} \times 2^{|k_{in}\cap k_{out}|}(2^{|k_{in}|+\delta_{in}+|k_{out}|+\delta_{out}-2|k_{in}\cap k_{out}|-n}) \qquad (3)$$

The data and memory complexities are similar to the differential MITM ones.

$$\mathcal{D} = \min(2^n, 2^{p-\delta_{in}+\min(|k_{in}|+\delta_{in}, |k_{out}|+\delta_{out})}), \qquad (4)$$
$$\mathcal{M} = \min(2^{|k_{in}|+\delta_{in}-|k_{in}\cap k_{out}|}, 2^{|k_{out}|+\delta_{out}-|k_{in}\cap k_{out}|}). \qquad (5)$$

The second line of (3) refers to the number of candidates for $k_{in} \cup k_{out}$ to be tested, which should be less than the whole set $k_{in} \cup k_{out}$. This holds if $p + |k_{in}| + |k_{out}| - |k_{in} \cap k_{out}| - n + \delta_{out} < |k_{in} \cup k_{out}|$, implying that $p < n - \delta_{out}$, which is ensured by an efficient distinguisher.

Remark 1. It holds for the reverse transition that $P(\Delta_{out} \xrightarrow{E_m^{-1}} \Delta_{in}) = 2^{-p'}$, where $p' = p + \delta_{out} - \delta_{in}$. Using this equality, one can infer that all the complexities of the reverse attack (the chosen ciphertext scenario) are equivalent to that of the forward attack: Eqs. (3), (4), (5).

4   New Improvements to Differential MITM Attacks

We present in this section three new improvements that can be incorporated into either the truncated or the original variants of differential MITM attacks. These improvements include an extension of the parallel partitioning technique, the state-test technique, and the probabilistic key recovery technique. The extended parallel partitioning technique, built upon the concept initially proposed in [11] (also reviewed in Sect. 2.2 of this paper), expands its range of applicability. The other two techniques focus on minimizing the information needed to be guessed during the key guessing step, which has a direct impact on the complexity. Specifically, the state-test technique, adopted from the impossible differential attacks [12], guesses a word of the state instead of a larger-size key material, thereby decreasing the total information that needs to be guessed. The probabilistic key recovery technique introduces a probability into the key-guessing step to reduce the number of active words in the key recovery parts of the attack, and consequently, the keybits involved.

4.1   Improving the Parallel Partitioning

As explained in Sect. 2.2, the original parallel partitioning method proposed in [11] effectively extends the attack for one round, in situations where the round key addition impacts only a portion of the cipher's state. This extension has no additional cost in time or data complexities if certain specific conditions are met. In this section, inspired by the structures commonly employed in MITM attacks like initial structures [4] and bicliques [19], we explain how the applicability range can be extended, specifically in two directions: One round extension, in the case of SPN ciphers with a whole-state key addition (applied to `CRAFT` in Sect. 6 of this paper). Additionally, more than one round extension, in the case of SPN ciphers with a partial-state key addition (applied in Sect. 7 of this paper). *The General Idea.* Without loss of generality, we explain the procedure for the round(s) extensions at the end of the cipher. A general view is shown in Fig. 4. Suppose that the cipher state is formed of W words of size s bits ($n = Ws$). Let A be the final state of the (truncated) differential-MITM attack, and B be the ciphertext, typically one or two rounds after A. Without any extra conditions, there are 2^{Ws} possibilities for A and B values, each. A set of F words,

Fig. 4. High-level representation of the proposed structures when added at the end of the attack. They could also be added at the input without loss of generality.

representing independent conditions, enforced at any point within the internal states of the added rounds, would reduce the number of possibilities for each of A and B to $2^{(W-F)s}$. The whole set of these possible values for A and B is called a pair of initial structures of size $2^{(W-F)s}$. These conditions typically involve forcing some words within the internal state into fixed values or imposing linear relationships on specific internal state words. This set of conditions of size Fs bits are selected so that having them, along with k_{in} (resp. k_{out}), allows one to uniquely determine an equivalent of Fs information bits of B (resp. A). Therefore, it makes sense to consider the Fs-bit internal state condition as the starting point for the structures.

As in the previous parallel partitioning method, we next perform the upper and lower procedures for both structures of size $2^{(W-F)s}$ in parallel, generating lists of $(B, \tilde{B}, i)$ and $(A, \tilde{A}, j)$. Therefore, we need to repeat the modified attack $2^{p-(W-F)s}$ times, while the first two terms of time complexity in (2) and (3), would be increased by a factor of $2^{(W-F)s}$, and the third term by a factor of $2^{2(W-F)s}$. So it is required to have efficient sieving on the candidates when merging the two partial lists of solutions, to retain a number well below the exhaustive search. The sieving over the candidate solutions is done in two ways: Firstly, a 2^{-Fs} sieving over $\tilde{B}$ and $\tilde{A}$, in the words involved in the starting point. These words computed from the both sides, (i.e. from the Fs information bits of $\tilde{A}$ and $\tilde{B}$ possibly with the aid of the associated k_{out} and k_{in}, respectively) should return the same values. Secondly, an L-bit sieving by new linear relations between A and B (similarly, $\tilde{A}$ and $\tilde{B}$) independent from those Fs-bit matching over the starting point. To profit from the whole sieving potential, we would need to discover $2(W-F)s$ linear relations between the pairs A and B, as well as between $\tilde{A}$ and $\tilde{B}$. The accurate value of L, upper bounded by $2(W-F)s$, depends on the particular cases. Finally, as far as $(W-F)s \leq p$, this technique imposes no additional costs to the time complexity of the original attack, besides the possible increase in the size of k_{in} and k_{out} if more bits are needed to check the linear relations.

The time complexity of the truncated differential-MITM is then:

$$\mathcal{T} = 2^{p-(W-F)s-\delta_{in}+|k_{in}\cap k_{out}|} \times (2^{(W-F)s} \times 2^{|k_{in}|+\delta_{in}-|k_{in}\cap k_{out}|}$$
$$+2^{(W-F)s} \times 2^{|k_{out}|+\delta_{out}-|k_{in}\cap k_{out}|} + 2^{|k_{in}|+\delta_{in}+|k_{out}|+\delta_{out}+2(W-F)s-Fs-L-2|k_{in}\cap k_{out}|}).$$

A positive side effect of this improvement is that thanks to the linear relations checked in the second sieving, we recover most of the time the whole values of the last subkey bits, which increases the size of the number of bits recovered, as we will see in the applications.

In the following, we provide a brief description of two specific examples illustrating how this technique is applicable for more than one round (for SKINNY), or even when the key is added to the entire state (for CRAFT).

The Case of SKINNY. In Sect. 7.1, we add two rounds at the end of our truncated differential MITM attack on SKINNY-64-192 as shown in Fig. 10. For the cipher SKINNY, knowing the first key row of the first added key allows checking of all the linear equations needed to profit from the whole sieving potential. In our attack, we construct structures of size 2^{40}, thus to profit from the whole sieving potential, we want to find 10 linear relations between the pairs A and B, as well as between $\tilde{A}$ and $\tilde{B}$. In our case, we guessed 3 subkey bits of the first key row of the first key which gives us 8 linear equations out of the 10 linear equations we want. And we guess the 2 subkey bits of the third column of the second key to find the 2 remaining linear relations. A similar idea is applied in our improvements of the SKINNY-128-384 attacks of Sect. 7.2.

The Case of CRAFT. In our attack of CRAFT in Sect. 6, we can check less than all the linear equations as we have a bigger margin - we do not need to use all the potentially available sieving. Thus we get our linear equation by checking the relation $\mathtt{MC}(A) \oplus B = \mathtt{MC}(\tilde{A}) \oplus \tilde{B}$ for all the non-fixed words, which erases the unknown key bits. Since we fixed 5 words in A and B, we get 11 linear relations on 4 bits each.

4.2 Probabilistic Key Recovery Technique

Usually, in the key recovery step of the differential attacks, we extend with probability one the differential characteristic for some rounds in both sides. Our second idea for improvement is to force one or more transitions to have zero difference, paying some probability. Thus the number of active words will decrease and fewer key bits will need to be guessed.

Suppose we are in the case of a differential attack. In the classical case, we propagate Δ_{in} and Δ_{out} with probability 1 for r_{in} rounds backward and r_{out} rounds forward respectively, determining the truncated differences that pairs of plaintexts/ciphertexts should verify. Then we will test the pairs verifying this truncated differential and the possible keys that would lead to this differential.

Here instead of extending Δ_{in} and Δ_{out} for r_{in} and r_{out} rounds with probability 1, we allow these transitions to happen with probability $2^{-p_{in}}$ and $2^{-p_{out}}$. Thus the overall probability for a random pair to follow the differential path is now $2^{-p-p_{in}-p_{out}}$. Therefore we have to repeat the attack $2^{p_{in}+p_{out}}$ more times. However, the number of pairs we keep for each side decreases by $2^{p_{in}}$ for the upper part and of $2^{p_{out}}$ for the lower part, so this is often compensated in the final time complexity. On the other hand, this technique restricts the large diffusion of active nibbles in the upper and lower sides, resulting in smaller sets of

k_{in} and k_{out}. We will show an application of this improvement in our attack on `CRAFT` in Sect. 6.

Comparison with Differential Attacks. In differential attacks, the complexity of the key recovery step is usually quite low, thanks to early abort techniques, and therefore the improvement of using probabilistic key recovery thechnique might be less advantageous, though it could still be applied. Here the number of involved keybits is directly associated to the final complexity.

4.3 Applying the State-Test Technique

The state-test technique was introduced in [12,15] in the context of impossible differential and MITM attacks respectively, to reduce the amounts of bits guessed in the key guessing step. The main idea is to test a part of the internal state that will uniquely define a partition of the involved key bits instead of guessing these keybits, reducing, therefore, the complexity of the guess.

During the key guessing part, some internal state words needed for computing $\tilde{P}$ or $\tilde{C}$ are smaller than the key materials that affect them. Thus guessing the state words instead of the key bits involved decreases the time complexity. Indeed, if l key or subkey bits are only needed to compute s bits of internal state with no differences but required to compute some internal state with differences. In such cases, we can guess s bits instead of the l key bits as P (or C) is fixed and this will define a disjoint partition of the involved keybits. Thus the time complexity of this part decreases of 2^{l-s}. We use this technique in both applications of Sect. 6 and Sect. 7.

This idea can be easily applied to differential-MITM attacks, unlike to classical differential attacks, as the plaintext for which we guess the keys is fixed, therefore defining a disjoint partition of the involved keybits.

Analysis of the Secret Information Recovered. As instead of recovering direct bits from the key we can recover non-linear equations, we have to be careful when computing the overall complexity that the number of recovered keybits of the actual key is bigger than the number of candidates for the triplets $(P, \tilde{P}, |k_{in} \cup k_{out}|)$ so that we do not have to deal with counters. We will see how to deal with particular cases in the applications.

5   MILP Modeling of the Truncated Differential-MITM Attack

In this section, we describe the automatic MILP-aided method for searching the truncated differential-MITM attack. We first introduce the modeling of the basic attack. Then, we explain how the state-test and probabilistic key recovery techniques can be incorporated into the model. We leave the inclusion of the parallel partitioning method as an open problem for future work. All source codes are available at https://github.com/CraftSkinny.

5.1   MILP Model of the Basic Attack

The set of constraints used in our model can be divided into three parts: constraints describing the distinguisher, constraints associated to k_{in} and k_{out}, and constraints describing the objective function.

Constraints Associated with the Distinguisher. This set of constraints is derived according to the method given in [26]. Once the model is solved, the approximated values of transition probabilities will be replaced by the accurate ones, given in the Branching Property Tables (BPT) described in Appendix D of [1]. To be more conservative, we can examine the accurate method given in [16], which computes the exact value of the probability for a given path. Moreover, we develop a distinguisher-only model with an accurate DBT, to compute the clustering effect on the differential probability, by summing up the probabilities of all the paths with $(\Delta_{in}, \Delta_{out})$ fixed to that of the optimum solution.

Constraints Determining k_{in} and k_{out}. We explain the method for the identification of the set k_{out}. A similar scenario holds for k_{in}, as well. The set k_{out} is determined by two factors: First, all the subkeys involved in the *differential propagation* of Δ_{out} to the ciphertext, and second, all the subkeys involved in the *value determination* of active words of Δ_{out} from the ciphertext. These two concepts have been previously used and modeled in other works such as automation of MITM attacks [28]. In Appendix B of [1], the two given theorems unify the description of the MILP modeling of the differential propagation and value determination through the linear layer of a given matrix **M** with input and output **a** and **b**, respectively.

These two concepts have been denoted in Fig. 5, where the differential propagation and value determination of active words are indicated by $\{D_i\}$ and $\{V_i\}$ chains, respectively. The active words in round r, i.e. D_r only depend on the active words in D_{r-1}, however, the words whose values should be determined in round r, i.e. V_r, depend on both D_{r-1} and V_{r-1}, i.e. on $D_{r-1} \vee V_{r-1}$.

There is a nuanced aspect at the starting point of chain $\{V_i\}$. Note that since the truncated differential attack is not dependent on the concrete values of differences, the attacker avoids the need to guess the subkeys required for value determination just before and after the distinguisher. Moreover, this property may partially overspread to the next rounds. For instance, consider a differential output of a distinguisher for `Skinny-64-192`, denoted by $\Delta_{out} = (a, b, 0, c)$, propagating to $(a + c, a, b, a)$ after the `MixColumn` operation. It's important to note that we do not have to determine the values of all four active words; rather, we only need to determine the values of the second and fourth words, while the other two can remain undetermined. This is because the differences in the second and fourth words are independent of the others, allowing them to have any non-zero value, thanks to the truncated nature of the attack.

Finally, we have specified the active words as well as the words whose values need to be determined over all states of the first r_{in} and the last r_{out} rounds. With this information, determining k_{in} and k_{out} becomes straightforward, as it

Fig. 5. Differential propagation and value determination in the lower part, where $r_1 = r_{in} + r_m$ and $r_2 = r_{in} + r_m + r_{out}$. The solid arrows show the differential propagation, the dashed ones show the value determination trace and the dotted arrows show the update of V_i by $V_i \vee D_i$, for $r_1 + 2 \le i \le r_2$. D_{r_1} corresponds to Δ_{out}. We refer to [1] for the Theorems and Propositions used.

equates to $D \vee V$ at the internal states where the round keys are introduced. All details of the MILP parameters for the two cases studied in this paper, i.e. CRAFT and Skinny-64-192, are given in Tables 2 and 3, where D_i is the input differential state of round $r_1 \le i \le r_2$, $D_{r_{in}} = \Delta_{in}$, $D_{r_{in}+r_m} = \Delta_{out}$, $V_{r_{in}} = V_{r_{in}+r_m} = 0$, and $T_i = D_i \vee V_i$. Finally, $T_i[\mathtt{R}_j]$ denotes the j^{th} row of state T_i.

Constraints Associated with the Objective Function. The model should minimize the time complexity given in (3) while preserving the efficient distinguisher constraint. Let the time complexity be dominated by the integer-valued variable u. The set of constraints associated to these parameters would be defined as follows.

$$
\begin{aligned}
\min \ & u \\
s.t. \ & p < n - \delta_{out} \\
& p + |k_{in}| < u \\
& p + |k_{out}| + \delta_{out} - \delta_{in} < u \\
& |k_{in} \cup k_{out}| + p + \delta_{out} - n < u
\end{aligned}
\tag{6}
$$

5.2 MILP Model of the Improved Attack

State-Test Enhanced Attack. In this section, we propose a general method for MILP modeling of the state-test technique introduced in Sect. 4.3, for reducing the guessed key material on each side. This technique specifically influences the constraints related to the value determination, in such a way that if a word of the internal state is preferred to be guessed, the value-determination trace corresponding only to this specific word should be aborted, then. To MILP model this event, for each state word whose value is supposed to be determined, we define a new binary variable s, indicating whether the corresponding word should undergo the state-test ($s = 1$) or not ($s = 0$). Then, the value-determination

Table 2. `Skinny-64-192` MILP parameters.

	Lower Part		Upper Part	
	Parameters	Description	Parameters	Description
Rounds	r_1	$r_{in} + r_m$	r_1	0
$(r_1 \leq i \leq r_2)$	r_2	$r_{in} + r_m + r_{out} - 1$	r_2	r_{in}
Differential	**M**	MC	**M**	Inverse MC
Propagation	**a**	$\mathrm{SR}(D_i)$	**a**	D_i
	b	D_{i+1}	**b**	$\mathrm{SR}(D_{i-1})$
Value	**M**	Inverse MC	**M**	MC
Determination	**a**	V_{i+1}	**a**	$\mathrm{SR}(V_{i-1})$
	b	$\mathrm{SR}(T_i)$	**b**	T_i
Involved Subkeys	$|k_{out}|$	$\sum\limits_{i=r_1+1}^{r_2} T_i[\mathrm{R_0, R_1}]$	$|k_{in}|$	$\sum\limits_{i=r_1}^{r_2-2} T_i[2] + (\bigvee\limits_{j \neq 2} T_i[\mathrm{R}_j])$

Table 3. `CRAFT` MILP parameters.

	Lower Part		Upper Part	
	Parameters	Description	Parameters	Description
Rounds	r_1	$r_{in} + r_m$	r_1	0
$(r_1 \leq i \leq r_2)$	r_2	$r_{in} + r_m + r_{out} - 1$	r_2	r_{in}
Differential	**M**	MC	**M**	Inverse MC
Propagation	**a**	D_i	**a**	$\mathrm{P}^{-1}(D_i)$
	b	$\mathrm{P}^{-1}(D_{i+1})$	**b**	D_{i-1}
Value	**M**	Inverse MC	**M**	MC
Determination	**a**	$\mathrm{P}^{-1}(V_{i+1})$	**a**	V_{i-1}
	b	T_i	**b**	$\mathrm{P}^{-1}(T_i)$
Involved subkeys	$|k_{out}|$	$\sum\limits_{\substack{i=r_1+1 \\ i \ even}}^{r_2} (\bigvee \mathrm{P}^{-1}(T_i)) + \\ \sum\limits_{\substack{i=r_1+1 \\ i \ odd}}^{r_2} (\bigvee \mathrm{P}^{-1}(T_i))$	$|k_{in}|$	$\sum\limits_{\substack{i=r_1 \\ i \ even}}^{r_2-2} (\bigvee (T_i)) + \\ \sum\limits_{\substack{i=r_1 \\ i \ odd}}^{r_2-2} (\bigvee (T_i))$

constraints would be as regular if $s = 0$ or aborted if $s = 1$. Theorem 3 from [1] identifies the required constraints corresponding to the value-determination for state-test enhanced attack.

Finally, in the objective function constraints (6), $|k_{in}|$ should be replaced by $|k_{in}| + s_{in}$ where $s_{in} = \sum_{0 \leq i \leq r_{in}-2} Hw(\mathbf{s}_i)$, and $|k_{out}|$ should be replaced by $|k_{out}| + s_{out}$ where $s_{out} = \sum_{r_{in}+r_m+1 \leq i \leq r_{in}+r_m+r_{out}-1} Hw(\mathbf{s}_i)$.

Probabilistic Key Recovery Ehanced Attack. In order to incorporate this technique into the model, it suffices to replace the differential propagation constraints generated according to Theorem 1 from [1] by the MILP model of the probabilistic truncated differential propagation given in [26]. In the upper part, this model is used for the inverse of `MixColumn` matrix as $\mathbf{M}$, and in the lower part, with `MixColumn` matrix. Then, the constraint $p + |k_{in}| < u$ of (6) should be modified

as $p + p_{out} + |k_{in}| < u$, and $p + |k_{out}| + \delta_{out} - \delta_{in} < u$ should be modified to $p + p_{in} + |k_{out}| + \delta_{out} - \delta_{in} < u$.

6   Application on 23-Round CRAFT

CRAFT is a lightweight, tweakable block cipher designed with the goal of protecting implementations against differential fault analysis while also providing strong security guarantees within the related-tweak model. The specification of this cipher is provided in the long version of this paper [1].

Security Claim. In [8], the authors presented optimum differentials for 13 and 14 rounds of CRAFT and claimed that using those differentials the attacker cannot have a successful single-tweak differential attack on 22 rounds. The best previous known attack on CRAFT [18] is a 21-round impossible differential attack with time complexity of $2^{106.53}$, data complexity of $2^{60.99}$ and memory complexity of 2^{100}.

Using truncated differential-MITM, enhanced with parallel partitioning for whole-state key addition ciphers, probabilistic key guessing, and the state-test techniques we managed to provide the best attack on CRAFT [8], improving by 2 rounds the best known attack, as detailed in Table 1.

6.1   An Attack on 23 Rounds of CRAFT

The truncated differential-MITM attack proposed in this section is composed of a 22-round core attack followed by one additional round using the parallel partitioning method. We conducted an automatic search for the optimum 22-round core attack on CRAFT, enhanced with the state-test and probabilistic key guessing techniques, based on the MILP method proposed in Sect. 5. We set $r_m = 11$, $r_{in} = 6$, and $r_{out} = 5$. The 11-round distinguisher and the core 22-round attack are represented in Figs. 6 and 7 respectively, with the following parameters:

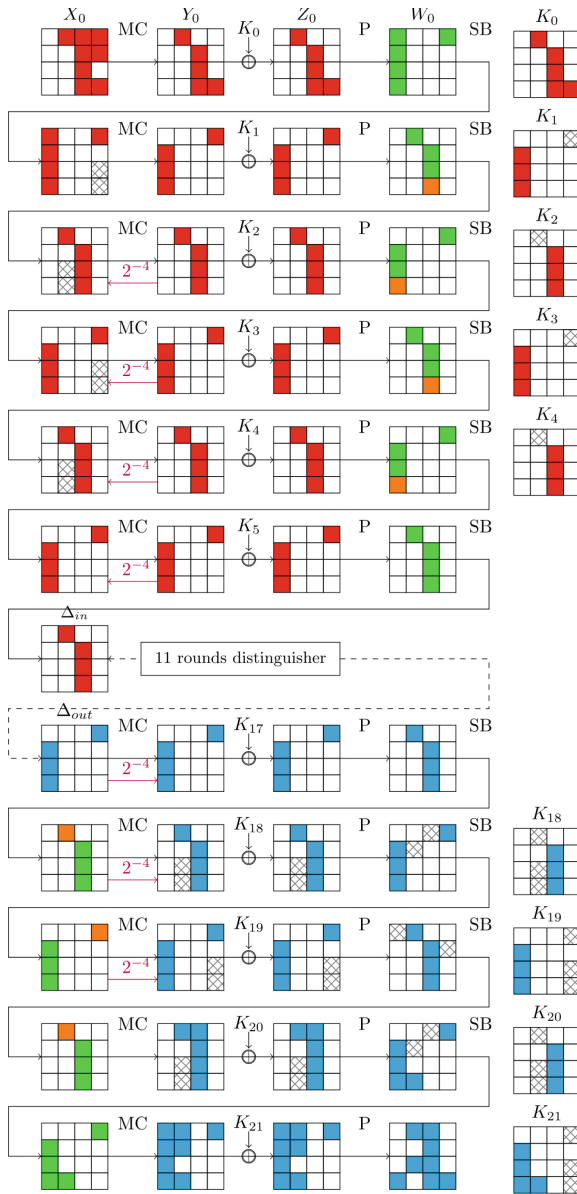| p	p_{in}	p_{out}	s_{in}	s_{out}	δ_{in}	δ_{out}	$|k_{in}|$	$|k_{out}|$
44	16	12	16	12	16	16	32	32

Fig. 6. CRAFT 11-round truncated differential distinguisher with $p = 44$

where s_{in} and s_{out} is the number of state-test bits to guess. The clustering effect was also examined which has a negligible impact. Since the matrix M used in

Fig. 7. The 22-round core part of the 23-round attack on `CRAFT` not including the structures. Differential propagation in the upper (lower) part has been shown in red (blue). The state-test nibble is shown in orange. The gray-striped nibbles are whose values are no longer required thanks to the state-tests, except K_{21} that the XOR of gray-striped nibbles are required. (Color figure online)

Table 4. Subkey nibbles guessed during the 23-round attack of CRAFT.

Subkey	k_{in}	k_{out}	$\lvert k_{in} \cap k_{out} \rvert$
Even	$K_0[1, 6, 10, 14, 15]$	$K_0[6, 10, 14]$	3 nibbles
Odd	$K_1[4, 8, 12]$	$K_1[4, 8, 12, 13, 3 \oplus 11 \oplus 15]$	3 nibbles

Fig. 8. The last round of the 23-round attack on CRAFT, using parallel partitioning technique. The red nibbles are known or can be computed from the fixed values and k_{in} in the upper part of the attack and the blue nibbles are known or can be computed from the fixed values in the lower part of the attack as $Y_{22} = \mathtt{MC}(\mathtt{SB}(W_{21}))$. The purple subkey nibbles are known in both the upper and lower parts.

the MixColumn operation is involutory, and the input and output differences of the distinguisher are the same, the propagation of active nibbles in the upper and lower parts is very similar. Thus, as shown in Table 4, the subkey words needed on both parts have many words in common thanks to the key schedule and the position of the odd and even subkeys, which allows us to efficiently apply the parallel partitioning technique. The parallel key guessing steps of this attack benefit from the state-test technique from Sect. 4.3 and the probabilistic key recovery technique of Sect. 4.2.

Probabilistic Key Recovery Technique. It is imposed that the difference during the MixColumn transition on nibbles $X_2[2], X_3[0], X_4[2], X_5[0], Y_{17}[0], Y_{18}[2]$ and $Y_{19}[0]$ is null. This decreases the overall probability by a factor of $2^{p_{in}+p_{out}} = 2^{16+12}$, but the number of active words, and consequently $\lvert k_{in} \rvert$ and $\lvert k_{out} \rvert$, also decreases dramatically. This attack is easy to adapt to 21 rounds, as we will show in Sec 6.2.

State-Test Technique. We use the state-test technique to guess four and three nibbles of the internal state in the upper and lower parts respectively: $W_1[14]$, $W_2[12], W_3[14], W_4[12], X_{20}[1], X_{19}[3], X_{18}[1]$ (detailed equations in [1]).

Parallel Guessing Step. From a pair (P, Y_{22}), we need to know the values of the active nibbles before the S-boxes, in green in Fig. 7, to be able to compute the

associated value of $\tilde{P}$ for each k_{in} and of $\tilde{Y}_{22}$ for each k_{out}. We will first describe the procedure over the 22 first rounds, and then show how to deal with the last round with parallel partitions.

1. We first guess the $6 \times 4 = 24$ subkey bits of $(k_{in} \cap k_{out})$ given in Table 4. We then fix $F = 5$ nibbles $Y_{22}[1, 6, 10, 14, 15]$, the words denoted by F in Fig. 8 and since we know the subkey nibbles $K_{22}[6, 10, 14]$ for both the upper and lower parts, and of $K_{22}[1, 15]$ for the upper part, then we also have the fixed values of the ciphertext $C[1, 6, 10, 14, 15]$.
2. As 5 nibbles from C and Y_{22} are fixed, we can build a pair of structures of size $2^{4(16-5)} = 2^{44}$ for Y_{22} and for the ciphertext C. We decrypt the 2^{44} ciphertexts C_s and we obtain the associated plaintexts P_s.
3. For each of the $2^{|k_{in} - k_{in} \cap k_{out}|} \times 2^{s_{in}} = 2^{8+16} = 2^{24}$ possible values i of k_{in} and the state-test relations, and each of the 2^{44} P_s from the structure defined at the previous step, compute all possible tuples $(P_s, \tilde{P}_s, i)$ such that they generate Δ_{in} after 6 rounds and they follow the probabilistic differential transitions in the first rounds. There are 2^{24} possible values for k_{in}, and it needs to be done for each $2^{\delta_{in}} = 2^{16}$ possible differences in Δ_{in}, but we expect only a proportion of $2^{-p_{in}} = 2^{-16}$ to verify the upward path, so the expected number of total solutions per P_s will be $2^{24+16-16}$, and if we use rebound-like techniques as shown in the original paper [11] the cost of this step would be equal to the number of solutions: $2^{44+24} = 2^{68}$. We store them in a hash table.
4. Similarly, for each of the $2^{|k_{out} - k_{in} \cap k_{out}|} \times 2^{s_{out}} = 2^{8+12} = 2^{20}$ possible values j of k_{out} and the state-test relations, and each of the 2^{44} states $Y_{22,s}$, compute all possible tuples $(Y_{22}, \tilde{Y}_{22}, j)$ such that there is the Δ_{out} difference on the state before the 17th S-box layer. It needs to be done for each $2^{\delta_{out}} = 2^{16}$ possible difference in Δ_{out} but we expect only a proportion of $2^{-p_{out}} = 2^{-12}$ to verify the downward path, so the expected number of total solutions per $Y_{22,s}$ will be $2^{20+16-12}$. As the previous step, we expect a cost and a number of solutions of $2^{44+24} = 2^{68}$. For each solution, check for possible matches on the hash table. The match is performed on two quantities:
   - The values of $\tilde{Y}_{22}[1, 6, 10, 14, 15]$ can be fully computed from $\tilde{C}[1, 6, 10, 14, 15]$ and the guessed k_{in}: 20-bit filter;
   - The linear relations between Y_{22} and C and between $\tilde{Y}_{22}$ and $\tilde{C}$, i.e. $Y_{22} \oplus C = \tilde{Y}_{22} \oplus \tilde{C}$, excluding the nibbles we could fully compute as they were already used in the 20 bit filter: 44-bit filter (11 equations on 4 bits each);
5. Repeat from Step 1 until the right key is retrieved: as the structures check 2^{44} plaintexts, we need to repeat all this procedure $2^{44-16+16+12-44} = 2^{12}$, to expect having a pair of data that verifies the middle distinguisher and the external transitions 2^{-56}.

The data complexity of the attack is 2^{64} as we ask the whole codebook to the oracle. But, applying the improved data technique from [11], we can fix $\frac{64-56}{2} = 4$ bits of the ciphertexts to reduce the data needed without increasing the time complexity, hence the data complexity will become $\mathcal{D} = 2^{60}$.

The memory complexity is determined by Step 3 where 2^{68} words of $64 \times 2 + 16 + 8 = 152$ bits each are stored in the hash table, so $\mathcal{M} = 2^{68}$.

The time complexity so far is

$$2^{12}2^{24}\left(2^{44}2^{24}2^{16-16} + 2^{44}2^{20}2^{16-12} + 2^{68+68-20-44}\right) = 2^{108}.$$

But the attack is not yet finished. Indeed, we have recovered $5 \times 4 = 20$ bits of K_1, as well as 64 bits of K_0, as the unguessed bits of $K_0 = K_{22}$ would be revealed as a side effect of the second sieving, explained in Step. 4. In addition, we recover $7 \times 4 = 28$ bits of information in key bits due to the 7 nibbles of state-tests, so $64 + 20 + 28 = 112$ information-bits of the key in total, which is bigger than the number of candidates, i.e. 108. The big question now is how to determine the whole key from each candidate because of the complex form of the state test equations.

How to Recover the Whole Key. The whole key K_0 is known. Considering this and rewriting the state-test equations given in Appendix E of [1], we recover the following values and relations. From the first equation we can derive the value of $K_1[3]$, giving us $K_1[11] \oplus K_1[15]$ from the guesses, and we choose to write $K_1[15]$ as a function of $K_1[11]$. Then we rewrite the equations given on rounds 4 and 18 as a function of some variables $x_1, \ldots, x_{24}$ which depend only on the plaintext, the ciphertext, the guessed values of the state tests, the nibbles of K_0 and those nibbles of K_1 from k_{in} and k_{out}. We obtain the following equations:

Equation 4 : $SB(K_1[2] \oplus SB(SB(K_1[0] \oplus x_1) \oplus SB(K_1[14] \oplus x_2) \oplus SB(K_1[7]$
$\oplus x_3) \oplus x_4) \oplus SB(SB(K_1[1] \oplus x_5) \oplus SB(K_1[10] \oplus x_6) \oplus x_7)$
$\oplus SB(SB(K_1[2] \oplus x_8) \oplus x_9)) \oplus SB(K_1[5] \oplus SB(SB(K_1[5] \oplus x_{10})$
$\oplus x_{11}) \oplus x_{12}) \oplus x_{13} = 0,$

Equation 18 : $SB(K_1[2] \oplus K_1[10] \oplus K_1[14] \oplus SB(SB(K_1[7] \oplus K_1[11] \oplus x_{14})$
$\oplus SB(SB(K_1[0] \oplus x_{15}) \oplus SB(K_1[14] \oplus x_{16}) \oplus x_{17})) \oplus SB(SB(K_1[1] \oplus K_1[9]$
$\oplus x_{18}) \oplus SB(K_1[10] \oplus x_{19}) \oplus x_{20}) \oplus SB(SB(K_1[2] \oplus K_1[10] \oplus K_1[14]$
$\oplus x_{21}) \oplus x_9) \oplus SB(K_1[5] \oplus SB(SB(K_1[5] \oplus x_{22}) \oplus x_{11}) \oplus x_{23}) \oplus x_{24} = 0.$

During the attack procedure, we stock the candidates we find after the matching in a table of size 2^s with $s \geq 100$, and sort this table based on $x_1, \ldots, x_{24}$. We will have 2^{96} groups of candidates of size $2^x = 2^{s-96}$ with the same variables. Since $x_1, \cdots, x_{24}$ define Eqs. 4 and 18, each candidate in one group will have the same set of solutions for those two equations. Thus, for each group, we can calculate and store the list of the 2^{20} solutions for Eq. 4 and 18. And in parallel, for each of the 2^x candidates in the group, we calculate the 2^{16} possible solutions of Eqs. 3 and 19. For each of these 2^{16} solutions, we get one match with the stored list of 2^{20} solutions from Eq. 4 and 18 with respect to the nibbles $K_1[0, 1, 9, 11, 14]$. Thus, for each element in the group, we can find 2^{16} possibilities for $K_1[0, 1, 2, 5, 6, 7, 9, 10, 11, 14, 15]$ giving us the whole key. Finally, the overall time cost of this will be:

$$\mathcal{T} = 2^{108-s} 2^{96} \left(2^{20} + 2^x 2^{16} \right),$$

which equals to 2^{124} if $x > 4$ and if $x = 4$ it equals to 2^{125} of small computations for recovering the whole key, which is lower than exhaustive search. And the memory complexity is 2^s as we stock the candidates in a table of size 2^s. For $s = 101$, we get a time complexity of $2^{124.58}$, and many trade-offs are possible.

6.2   Other Attacks on `CRAFT` and Conclusion

The attack described above can be applied to fewer than 23 rounds straightforwardly by subsequently removing one round from each side and adapting the structures to the known key nibbles. These results can be seen in Table 1. In this case, the data will be smaller, as we can apply the data reduction idea from [11].

It is worth pointing out that the authors did not expect differential attacks to reach 22 rounds with the best paths they found. Given that these paths reached 13 rounds, and the distinguisher used in our attack reaches 11 (two less rounds), it makes us deduce that truncated MITM attacks seem to be much more performant on `CRAFT` than differential attacks, and the most performant attack, to the best of our knowledge. We believe this is the case because of its alternated key schedule and the existence of iterative truncated paths, both with period two. The number of rounds reached is still far from the full version, but we expect further attacks with these techniques and better dealing with the state-test relations to reach more rounds.

7   Applications: `SKINNY-64-192` and `SKINNY-128-384`

In this section, we provide two applications on two variants of `SKINNY`. Using a truncated path, state-test technique and the enhanced parallel partitioning method over two rounds, we provide a new attack on 23-round `SKINNY-64-192`, with slightly better time and lower data than the best known attack. In order to illustrate the importance of our improved parallel partitioning method, we have improved the previous differential MITM attack on `SKINNY-128-384`, that provided the best attacks on the single tweakey setting, and we have managed to slightly reduce their time or data complexity thanks to the structures, providing the best current attack on `SKINNY-128-384` in the single tweakey setting. The specification of this cipher is given in the long version of this paper [1]. In the proposed attacks, we use the modified round key addition in the upper part where $U_r = \text{MC}(\text{SR}(K_r))$ is Xored to the output state of the `MixColumn` operation. This shows the fact that (truncated) differential MITM attacks work well on reduced-round variants of the `SKINNY` constructions. As further work, we plan to automatize the tool including more evolved structures, as the ones used in the MITM attacks from [2,3,10], and we expect we might be able to reach more rounds in both variants.

Fig. 9. 9-round truncated differential characteristic of probability 2^{52} for `SKINNY-64-192`.

7.1 Attack on 23-Round `SKINNY-64-192`

Since `SKINNY` key schedule is linear, it would be an efficient approach to guess some subkey bits and retrieve the whole key after guessing enough independent round key bits. Moreover, the key schedule makes the evaluation of the dimension of any set of round key nibbles easy since a round key nibble depends on exactly three master key nibbles, $\mathbf{TK1}[i], \mathbf{TK2}[i]$, and $\mathbf{TK3}[i]$, for a specific $i \in \{0, \ldots, 15\}$.

An Attack on `SKINNY-64-192` Without Parallel Partitioning. We first propose a 21-round truncated differential-MITM attack which is followed by two additional rounds using the parallel partitioning method, resulting in a 23-round attack. We conducted an automatic search for the optimum 21-round core attack on `SKINNY-64-192`, enhanced with the state-test key guessing technique, based on the MILP method proposed in Sect. 5. We set $r_m = 9$, $r_{in} = 6$, and $r_{out} = 6$. The core attack is represented in Fig. 14 in Appendix F of the longer version of this paper [1], with the following parameters:

Rounds	p	s_{in}	s_{out}	δ_{in}	δ_{out}	$|k_{in}|$	$|k_{out}|$
22	52	4	4	4	8	128	116

The clustering effect increases p to 51.78, which is not very significant. Although we have included the probabilistic key recovery method in our search, the optimum solution returned $p_{in} = p_{out} = 0$, meaning that a deterministic key guessing is more efficient for `SKINNY-64-192`.

Table 9 in Appendix F of [1] describes all the subkey nibbles of k_{in} and k_{out} needed in the 21-round attack of `SKINNY-64-192` which are also reflected in Fig. 14 from [1] (that also includes this information in the needed key nibbles for the 23-round full attack in Table 10). It also indicates which nibble of the master key each subkey nibble of k_{in} and k_{out} depends on, and presents the total number of linear relations in each nibble of the master key, given k_{in} and k_{out}. In this way, we can determine the number of common linear relations, i.e. the size of the intersection $|k_{in} \cap k_{out}|$ which is $15 \times 4 = 60$ bits.

State-Test Ttechnique. We could use the state-test technique to reduce $|k_{in}|$ and $|k_{out}|$ by testing the 3 and 2 respective nibbles of Table 5, instead of guessing the 5 and 4 respective subkeys nibbles described on the table. Thanks to this

technique, the number of bits in k_{in} could be reduced by 8 bits, and the number of bits in k_{out} could also be reduced by 8 bits. The optimal time complexity is nevertheless reached when we only consider one state-test technique for each part (the non-crossed ones in Table 5).

Table 5. Non-linear relations available in the state-test technique on the 21 and 23 rounds of `SKINNY-64-192`. The crossed cells will not be used in the 23-round attack.

| Wanted nibble	RoundKey nibbles involved	Nibbles needed from the precedent state
~~$X_4[2]$~~	$K_3[2]$, $K_2[7]$	$K_3[2] \oplus \text{SC}(X_3[2]) \oplus \text{SC}(X_3[15]) \oplus \text{SC}(K_2[7] \oplus \text{SC}(X_2[7]) \oplus \text{SC}(X_2[10]))$
~~$X_4[10]$~~	$K_3[5]$, $K_2[7]$	$K_3[5] \oplus \text{SC}(X_3[5]) \oplus \text{SC}(K_2[7] \oplus \text{SC}(X_2[7]) \oplus \text{SC}(X_2[10]))$
$X_5[9]$	$K_4[4]$, $K_3[6]$	$K_4[4] \oplus \text{SC}(X_4[4]) \oplus \text{SC}(K_3[6] \oplus \text{SC}(X_3[6]) \oplus \text{SC}(X_3[9]))$
~~$X_{17}[5]$~~	$K_{17}[5]$, $K_{18}[6]$	$K_{17}[5] \oplus \text{SC}^{-1}(X_{18}[10]) \oplus \text{SC}^{-1}(X_{18}[14]) \oplus \text{SC}^{-1}(K_{18}[6] \oplus \text{SC}^{-1}(X_{19}[7]) \oplus \text{SC}^{-1}(X_{19}[11]) \oplus \text{SC}^{-1}(X_{19}[15]))$
$X_{16}[5]$	$K_{16}[5]$, $K_{17}[6]$	$K_{16}[5] \oplus \text{SC}^{-1}(X_{17}[10]) \oplus \text{SC}^{-1}(X_{17}[14]) \oplus \text{SC}^{-1}(K_{17}[6] \oplus \text{SC}^{-1}(X_{18}[7]) \oplus \text{SC}^{-1}(X_{18}[11]) \oplus \text{SC}^{-1}(X_{18}[15]))$

Attack Steps. We describe now the core attack on 21 rounds of `SKINNY-64-192`. The guesses needed for this attack are summarized in Table 9 from [1] and the state test equations to guess are given in Table 5.

1. Ask for the encryption of the whole codebook (we will explain later how to apply the data reduction of Sect. 2.3 to this case).
2. Pick one plaintext/ciphertext pair (P, C).
3. First we guess the 44 subkey bit common relations shared between k_{in} and k_{out}.
4. Compute all possible tuples $(P, \tilde{P}, i)$ for each value i of the remaining 84 bits of k_{in} such that the difference after the 6th S-box layer is according to Δ_{in} of Fig. 9. At the end of this step, we have 2^{84+4} possible candidates. For all $\tilde{P}$, compute $E(\tilde{P}) = \hat{C}$ and store them in a hash table.
5. Similarly, for each value j of the remaining 72 bits of k_{out}, compute all possible tuples $(C, \tilde{C}, j)$ so that the difference on the state before the 15th S-box layer is according to Δ_{out} of Fig. 9. At the end of this step, we have 2^{72+8} possible candidates for the tuples $(C, \tilde{C}, j)$. And check for possible matches on the hash table. The match is performed on both the new ciphertext $\hat{C}$ and $\tilde{C}$ so that $(\tilde{P}, \tilde{C})$ is a valid plaintext/ciphertext pair.
6. Repeat from Step 1 until the right key is retrieved.

Adjustment for the Number of Candidate Triplets. If we consider what complexity such an attack would have, we obtain:

$$2^{52-4}2^{44}\left(2^{128-44+4} + 2^{116-44+8} + 2^{168-64}\right) = 2^{196},$$

that exceeds the exhaustive search.

Fig. 10. The two last rounds of the 23-round attack on `SKINNY-64-192`, using parallel partitioning with fixed values for $X[1] \oplus X[11]$, $X[3] \oplus X[9]$, $Y[2] \oplus Y[13]$ and $Y[7] \oplus Y[13]$.

This is a possible side effect of the state test technique, that allows less sieving regarding the keybits. To compensate this, we will consider less state-test equations: we will guess in addition $K_2[7]$, that will determine because of the two first state-test equations $K_3[2]$ and $K_3[5]$; and also $K_{18}[6]$ and $K_{17}[6]$ that will determine because of the fourth and fifth state test equations $K_{17}[5]$ and $K_{16}[5]$. These three words guesses provide an additional sieving of 6 words, so 2^{-24}. The complexity will then be below exhaustive search:

$$2^{52-4}2^{68}\left(2^{128+4-68+4} + 2^{116+8-68+8} + 2^{132-64}\right) = 2^{184}.$$

Attack on 23-Round of SKINNY-64-192 Now we will explain how to extend for two rounds this 21-round attack thanks to the improvement with the parallel partitioning method. In addition to the 21-round attack, we guess words $K_{22}[6]$, $K_{21}[0]$, $K_{21}[3]$ and $K_{21}[1]$, as they will be needed to sieve with respect to all the available potential during the parallel partitioning. In Fig. 10, we give the scheme of the two last rounds of the attack. We represent in red the internal state and the subkey words that we will use to build the parallel partitions and compute the upper part of the attack. Similarly, we represent in blue the internal state and the subkey words used to build the initial structure and compute the lower part of the attack. The remaining attack procedure of the 23-round attack is similar to the 21-round one:

1. We fix the values of the words F, $A \oplus B$ and $C \oplus D$ of W_{21} in Fig. 10, and of the words $a \oplus d$ and $b \oplus d$ of Y_{22} in Fig. 10.
2. We guess the 76 bits of common linear relations of $(k_{in} \cap k_{out})$, detailed in Table 10 in [1].
3. Then for each value i of the remaining 56+4=60 bits of $(k_{in} - k_{in} \cap k_{out})$, we can compute from the values of step 1, the fixed nibbles $C[5, 12, 14, 15]$,

$C[2] \oplus C[13]$ and $C[7] \oplus C[13]$ as shown in Fig. 10. Then we have 2^{40} possible states for C as we have imposed the values of 6 nibbles thus there are 10 nibbles that can take all possible values. For all the possible values for C, we compute the ciphertext and get the corresponding plaintext P.

4. Compute all possible tuples $(P, \tilde{P}, i)$ for each $2^{|k_{in} - k_{in} \cap k_{out}|} \times 2^{s_{in}} = 2^{56+4}$ values i of $k_{in} - |k_{in} \cap k_{out}|$ and the state test relation, and each P from the structure defined at the previous step, such that they generate Δ_{in} after six rounds. At the end of this step, we have $2^{40} \times 2^{60+4} = 2^{104}$ possible candidates, and store them in a hash table.

5. Similarly, for each of the $2^{|k_{out} - k_{in} \cap k_{out}|} \times 2^{s_{out}} = 2^{52+4}$ possible values j of $(k_{out} - k_{in} \cap k_{out})$ and state-test relations, we can compute from the values of step 1, the fixed values $X_{21}[0, 8, 10, 15]$, $X_{21}[1] \oplus X_{21}[11]$ and $X_{21}[3] \oplus X_{21}[9]$, in blue in Fig. 10. We then pick all the 2^{40} possible states of X_{21} and we compute the 2^{40} possible $W_{20} = \mathtt{MC}^{-1}(X_{21})$.

6. For each $2^{52+4} = 2^{56}$ values j of $(k_{out} - k_{in} \cap k_{out})$ and state-test relation and for each value of state W_{20}, compute all possible tuples $(W_{20}, \tilde{W}_{20}, j)$ so that there is a difference on the state before the 15th S-box layer on both nibbles. At the end of this step we have $2^{40} \times 2^{56+8} = 2^{104}$ possible candidates for the tuples $(W_{20}, \tilde{W}_{20}, j)$.

7. Check for possible matches on the hash table. The match is performed on two quantities:
   - The values of the nibbles $\tilde{X}_{21}[8]$ and $\tilde{X}_{21}[15]$ can be fully computed from $\tilde{C}[2] \oplus \tilde{C}[13], \tilde{C}[7] \oplus \tilde{C}[13]$ and the guessed k_{in} and similarly the values of the nibbles $\tilde{C}[5], \tilde{C}[12], \tilde{C}[14]$ and $\tilde{C}[15]$ can be fully computed from $\tilde{X}_{21}[0], \tilde{X}_{21}[10], \tilde{X}_{21}[1] \oplus \tilde{X}_{21}[11], \tilde{X}_{21}[3] \oplus \tilde{X}_{21}[9]$ and the guessed k_{out}: a $6 \times 4 = 24$ bit filter;
   - The linear relations between X_{21} and C, and $\tilde{X}_{21}$ and $\tilde{C}$: a 80-bit filter (10×2 equations on 4 bits each), summarized in Table 11 from [1].

8. Repeat from Step 1 with different values for the fixed nibbles until the right key is retrieved.

Linear Relations to Match with the Parallel Partitioning Improvement. At the end of the attack procedure we have to match X_{21} and C and their associated pair $\tilde{X}_{21}$ and $\tilde{C}$. After two rounds of `SKINNY-64-192`, if we know the 4 first subkey nibbles then the relations between the words of X_{21} and C and $\tilde{X}_{21}$ and $\tilde{C}$, of Fig. 10, are linear. In our application, we have guessed the subkey nibbles $K_{21}[0], K_{21}[1]$ and $K_{21}[3]$. We do not know the subkey nibble $K_{21}[2]$ but we have guessed the subkey nibbles $K_{22}[2]$ and $K_{22}[6]$ and thus we obtain the linear relations of Column 3. To match both sides of the equations when the subkey word is not completely determined, we add on each side the subkey information known by each side, respectively. Thus the relations are linear as we can compute each side of the equations independently. The linear relations (given for the pair (X_{21}, C), but that are also true for $(\tilde{X}_{21}, \tilde{C})$) to match are given in Table 11 of

Appendix F in [1]. The time complexity of this step will be:

$$\mathcal{T} = 2^{52-4-40}2^{76}\left(2^{56+4+4}2^{40} + 2^{52+4+8}2^{40} + 2^{104+104-24-40-40}\right) = 2^{188}.$$

To reduce the data complexity of our attack, we use the improvement of [11] to impose the value of $\frac{64-48}{2} = 8$ bits of the ciphertexts. Thus we fix the nibbles $C[6]$, $C[1] \oplus C[13]$, $\tilde{C}[6]$ and $\tilde{C}[1] \oplus \tilde{C}[13]$, which is a 8 bits condition. Moreover, we compute and build a stored table with the data needed to perform the attack to avoid doing the decryption during the upper part of the attack. Finally the data complexity of this attack is $\mathcal{D} = 2^{64-8} = 2^{56}$. The memory complexity is determined by Step 6 storing $\mathcal{M} = 2^{104}$ words of $64 + 64 = 128$.

Thanks to the truncated differential, we manage to extend the characteristic for more rounds than if it was a concrete differential characteristic since the subkeys around the characteristic are not needed in the attack. Moreover in the case of `SKINNY`, we can use, with little cost, the parallel partitioning improvement to reach two more rounds and thus reach the same best number of rounds in the single tweak setting as the best known attack.

7.2   Improved Attacks on 25-Round `SKINNY-128-384`

We consider the 24-round and the 25-round attack given in [11] in Sect. 3.4 and 3.5 respectively. Each uses a different differential. By considering the core attack on 23 rounds that is extended by one round in the paper, we will apply our improved structure with a two-round extension, reducing the data complexity of the best known attacks, that cover 25 rounds. By considering the 25-round attack also considered in the paper, and including the last key recovery round in the final structure, reaching now 2 rounds instead of 1, we are able to reduce the overall time complexity, as we have fewer key bits to guess, providing the 25-round `SKINNY-128-384` attacks with the lowest complexity. These results are shown in Table 1, compared to all the previous results.

Reducing the Data Needed. The application is quite straightforward given the previous attack, so we will omit the details and refer to the original paper [11]. We consider exactly the same core over 23 rounds, but we will add two rounds instead of one. For this, we will additionally guess nibbles 5, 2, and 4 from K_{24} in k_{in}, and nibbles 14 and 11 from K_{23} in K_{out}. All these guesses add equations in the common part of the key to guess. This will allow us to build structures of size 2^{64}, with 2 fixed words from the first and the last columns, 3 from the second, and one from the third – see Fig. 15 from [1]. The full filtering can be applied as we can rewrite all equations as linear ones in the unknown parts of the key. The time complexity becomes:

$$\mathcal{T} = 2^{105.9-64}2^{120+40}\left(2^{128-16}2^{64} + 2^{136-24}2^{64} + 2^{128-16+64+136-24+64-128-64}\right)$$

So we have $\mathcal{T} = 2^{201.9}\left(2^{176} + 2^{176} + 2^{160}\right) = 2^{378.9}$, and the memory complexity, with $x = (128 - 105.9)/2 = 11.05$ becomes $2^{176-11.05} = 2^{165}$, and data

$\mathcal{D} = 2^{128-11.05} = 2^{117}$, providing the best data complexity for a 25-round attack, as can be seen in Table 1.

Reducing the Time. We consider the attack from [11] on 25 rounds, which added 4 rounds at the top and 5 rounds at the bottom of a 15-round distinguisher, plus one added through a structure. As the bottleneck term is the five rounds at the lower part, where we guess 8 more keybits than the upper part, we propose our attack with a 2-round structure, yielding a configuration of $4+15+4+2$ rounds. As we are guessing 1 less word from k_{out} (all but the word 12 from K_{23}) than before, as we add 7 words needed in order to verify linearly the equations in the structure and to build small structures, all the complexities stay the same, but the time reduces of about a factor of 2^8 as shown in Table 1. We can build now 4 fixed relations between the input and the output of the structure, given by equations A_1+C_1, A_2+C_2, A_3+C_3, A_4+C_4, B_1+C_1, B_2+C_2, B_3+C_3 from Fig. 16 from [1], and the structures will have a size of 2^{72}. We do not guess the word 12 from round K_{23} anymore, which reduces by 1 word the lower guess and leaves the number of common keybits the same. The time complexity becomes:

$$\mathcal{T} = 2^{116.5-72}2^{120}\left(2^{128+72} + 2^{136-64+56+72} + 2^{128+72+136-64+56+72-72-128}\right),$$

giving $\mathcal{T} = 2^{164.5}\left(2^{200} + 2^{200} + 2^{200}\right) = 2^{366}$, while previous best time was $2^{372.5}$. Applying the same idea as in the previous attack for reducing the data and memory needed, we obtain, without modifying the time and memory complexities a data of also $2^{128-x} = 2^{122.25}$, as $x = \frac{128-116.5}{2} = 5.75$, with a memory complexity of $2^{194.25}$, that stays the same.

8   Conclusion

We have implemented a tool, based on MILP modeling, that finds the distinguishers that produce the best overall attack complexities when considering the key-guessing rounds and their two related improvements. The inclusion of the structures is left as an open problem for the future.

We have been able to apply the variety of results to `CRAFT`, `SKINNY-64-192` and `SKINNY-128-384`. For `CRAFT` we managed to improve by two rounds the previous attacks with a truncated version of differential MTIM. For `SKINNY-128-384` we managed to improve the complexities of the best known attacks in the single tweak setting, and for `SKINNY-64-192` we matched the same number of rounds while the time stays comparable and the memory and data are worse.

We have in particular shown that differential MITM attacks have a different nature than differential attacks that allow them to be combined with MITM-related techniques that can not be combined with differential attacks, like the parallel partitions that are closely related to initial structures and bicliques.

Actually, we leave as an open problem to produce a tool that will combine differential MITM attacks with parallel partitioning technique [4] and bicliques [19] over more rounds, as the one in [14].

In addition, we showed that differential MITM attacks can be easily combined with the state-test technique, thanks to the fact that each key is tested for a fixed data. In differentials attacks it would be much harder to apply.

References

1. Ahmadian, Z., Khalesi, A., M'foukh, D., Moghimi, H., Naya-Plasencia, M.: Improved differential meet-in-the-middle cryptanalysis. Cryptology ePrint Archive, Paper 2024/351 (2024). https://eprint.iacr.org/2024/351

2. Aoki, K., Sasaki, Y.: Preimage attacks on one-block MD4, 63-step MD5 and more. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) Selected Areas in Cryptography. LNCS, vol. 5381, pp. 103–119. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04159-4_7

3. Aoki, K., Sasaki, Y.: Meet-in-the-middle preimage attacks against reduced SHA-0 and SHA-1. In: Halevi, S. (ed.) Advances in Cryptology. CRYPTO 2009. LNCS, vol. 5677, pp. 70–89. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03356-8_5

4. Aoki, K., Sasaki, Y.: Meet-in-the-middle preimage attacks against reduced SHA-0 and SHA-1. In: Halevi, S. (ed.) Advances in Cryptology. CRYPTO 2009. LNCS, vol. 5677, pp. 70–89. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03356-8_5

5. Banik, S., Pandey, S.K., Peyrin, T., Sasaki, Y., Sim, S.M., Todo, Y.: GIFT: a small present: towards reaching the limit of lightweight encryption. In: Fischer, W., Homma, N. (eds.) Cryptographic Hardware and Embedded Systems. CHES 2017. LNCS, vol. 10529, pp. 321–345. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66787-4_16

6. Bao, Z., Guo, J., Shi, D., Tu, Y.: Superposition meet-in-the-middle attacks: updates on fundamental security of AES-like hashing. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology. CRYPTO 2022. LNCS, vol. 13507, pp. 64–93. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15802-5_3

7. Beierle, C., et al.: The SKINNY family of block ciphers and its low-latency variant MANTIS. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology. CRYPTO 2016. LNCS, vol. 9815, pp. 123–153. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_5

8. Beierle, C., Leander, G., Moradi, A., Rasoolzadeh, S.: CRAFT: lightweight tweakable block cipher with efficient protection against DFA attacks. IACR Trans. Symmet. Cryptol. **2019**(1), 5–45 (2019)

9. Biham, E., Shamir, A.: Differential cryptanalysis of des-like cryptosystems. In: Menezes, A., Vanstone, S.A. (eds.) Advances in Cryptology. CRYPTO 1990. LNCS, vol. 537, pp. 2–21. Springer, Heidelberg (1990). https://doi.org/10.1007/3-540-38424-3_1

10. Bogdanov, A., Khovratovich, D., Rechberger, C.: Biclique cryptanalysis of the full AES. In: Lee, D.H., Wang, X. (eds.) Advances in Cryptology. ASIACRYPT 2011. LNCS, vol. 7073, pp. 344–371. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_19

11. Boura, C., David, N., Derbez, P., Leander, G., Naya-Plasencia, M.: Differential meet-in-the-middle cryptanalysis. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology. CRYPTO 2023. LNCS, vol. 14083, pp. 240–272. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38548-3_9

12. Boura, C., Naya-Plasencia, M., Suder, V.: Scrutinizing and improving impossible differential attacks: applications to clefia, camellia, lblock and simon. In: Sarkar, P., Iwata, T. (eds.) Advances in Cryptology. ASIACRYPT 2014. LNCS, vol. 8873, pp. 179–199. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45611-8_10

13. Diffie, W., Hellman, M.E.: Special feature exhaustive cryptanalysis of the NBS data encryption standard. Computer **10**(6), 74–84 (1977)

14. Dong, X., Hua, J., Sun, S., Li, Z., Wang, X., Hu, L.: Meet-in-the-middle attacks revisited: key-recovery, collision, and preimage attacks. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology. CRYPTO 2021. LNCS, vol. 12827, pp. 278–308. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84252-9_10

15. Dunkelman, O., Sekar, G., Preneel, B.: Improved meet-in-the-middle attacks on reduced-round DES. In: Srinathan, K., Rangan, C.P., Yung, M. (eds.) Progress in Cryptology. INDOCRYPT 2007. LNCS, vol. 4859, pp. 86–100. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-77026-8_8

16. Eichlseder, M., Leander, G., Rasoolzadeh, S.: Computing expected differential probability of (truncated) differentials and expected linear potential of (multidimensional) linear hulls in SPN block ciphers. In: Progress in Cryptology. INDOCRYPT 2020, pp. 345–369. Springer (2020)

17. Hadipour, H., Bagheri, N., Song, L.: Improved rectangle attacks on SKINNY and CRAFT. IACR Trans. Symmetric Cryptol. **2021**(2), 140–198 (2021)

18. Hadipour, H., Sadeghi, S., Eichlseder, M.: Finding the impossible: automated search for full impossible-differential, zero-correlation, and integral attacks. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology. EUROCRYPT 2023. LNCS, vol. 14007, pp. 128–157. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30634-1_5

19. Khovratovich, D., Rechberger, C., Savelieva, A.: Bicliques for preimages: attacks on skein-512 and the SHA-2 family. In: Canteaut, A. (ed.) Fast Software Encryption. FSE 2012. LNCS, vol. 7549, pp. 244–263. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34047-5_15

20. Knudsen, L.R.: Truncated and higher order differentials. In: Preneel, B. (ed.) Fast Software Encryption: Second International Workshop. LNCS, vol. 1008, pp. 196–211. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-60590-8_16

21. Knudsen, L.R., Wagner, D.A.: Integral cryptanalysis. In: Daemen, J., Rijmen, V. (eds.) Fast Software Encryption. FSE 2002. LNCS, vol. 2365, pp. 112–127. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45661-9_9

22. Lallemand, V., Naya-Plasencia, M.: Cryptanalysis of KLEIN. In: Cid, C., Rechberger, C. (eds.) Fast Software Encryption. FSE 2014. LNCS, vol. 8540, pp. 451–470. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-46706-0_23

23. Leander, G., Abdelraheem, M.A., AlKhzaimi, H., Zenner, E.: A cryptanalysis of printcipher: the invariant subspace attack. In: Rogaway, P. (ed.) Advances in Cryptology. CRYPTO 2011. LNCS, vol. 6841, pp. 206–221. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_12

24. Ma, Z., Li, M., Chen, S.: Meet-in-the-middle attacks on round-reduced CRAFT based on automatic search. IET Inf. Secur. **17**(3), 534–543 (2023)
25. Matsui, M.: Linear cryptanalysis method for DES cipher. In: Helleseth, T. (ed.) Advances in Cryptology. EUROCRYPT 1993. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-48285-7_33
26. Moghaddam, A.E., Ahmadian, Z.: New automatic search method for truncated-differential characteristics application to midori. SKINNY and CRAFT. Comput. J. **63**(12), 1813–1825 (2020)
27. Rasoolzadeh, S., Ahmadian, Z., Salmasizadeh, M., Aref, M.R.: An improved truncated differential cryptanalysis of KLEIN. Tatra Mount. Math. Publ. **67**(1), 135–147 (2016)
28. Shi, D., Sun, S., Derbez, P., Todo, Y., Sun, B., Hu, L.: Programming the demirci-selçuk meet-in-the-middle attack with constraints. In: Peyrin, T., Galbraith, S.D. (eds.) Advances in Cryptology. ASIACRYPT 2018. LNCS, vol. 11273, pp. 3–34. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03329-3_1
29. Shi, D., Sun, S., Song, L., Hu, L., Yang, Q.: Exploiting non-full key additions: full-fledged automatic demirci-selçuk meet-in-the-middle cryptanalysis of SKINNY. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology. EUROCRYPT 2023. LNCS, vol. 14007, pp. 67–97. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30634-1_3
30. Song, L., et al.: Optimizing rectangle attacks: a unified and generic framework for key recovery. In: Agrawal, S., Lin, D. (eds.) Advances in Cryptology. ASIACRYPT 2022. LNCS, vol. 13791, pp. 410–440. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22963-3_14
31. Tolba, M., Abdelkhalek, A., Youssef, A.M.: Impossible differential cryptanalysis of reduced-round SKINNY. In: Joye, M., Nitaj, A. (eds.) Progress in Cryptology. AFRICACRYPT 2017. LNCS, vol. 10239, pp. 117–134. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-57339-7_7