Improved differential cryptanalysis of SPEEDY

Tim Beyne[1] and Addie Neyt[1,2]

[1] KU Leuven, Leuven, Belgium, name.lastname@esat.kuleuven.be
[2] RMA, Brussels, Belgium, addie.neyt@mil.be

Abstract. SPEEDY is a family of lightweight block ciphers designed by Leander et al. Several differential attacks have been reported on the SPEEDY variants. However, nearly all of these attacks are based on differential characteristics with probabilities that differ from their reported values. These discrepancies arise from incorrect calculations of the (key-averaged) probability, particularly in consecutive steps within one round without intermediate key addition. In this paper, we revisit all reported differential characteristics and accurately calculate their key-averaged probabilities using quasidifferential trails. We extend this to also estimate the fixed-key probability. Our analysis reveals several characteristics with zero or significantly altered probability, invalidating several proposed attacks. We further implement a search algorithm and find a 5.5-round differential distinguisher that can be used to mount a full-round key-recovery attack with a data complexity of 2^{183} and a time complexity of 2^{185}. The memory complexity varies: in the chosen-plaintext setting, it is 2^{156}, whereas in the chosen-ciphertext setting, it is 2^{36}.

Keywords: Differential cryptanalysis · SPEEDY · Quasidifferential trails · Key recovery

1 Introduction

Low-latency lightweight cryptography aims to balance security and minimal encryption delay. This requirement has led researchers to propose new cipher designs, such as SPEEDY, which is the focus of this paper.

The SPEEDY family of block ciphers, introduced by Leander, Moos, Moradi, and Rasoolzadeh at CHES 2021 [LMMR21], was designed to enable low-latency encryption on hardware platforms. The default instance, SPEEDY-r-192, has a 192-bit block size and 192-bit key with $r \in \{5, 6, 7\}$ denoting the number of rounds. The design of SPEEDY has been the subject of several cryptanalytic studies [BDBN23, YJZZ22, Zha24, WNL+23] across its different variants. Some of these works have proposed full-round key-recovery attacks, questioning the ciphers' security claims. These attacks are based on differential cryptanalysis.

Differential cryptanalysis was first published by Biham and Shamir [BS91] in 1990, and has since become a cornerstone in the security evaluation of block ciphers. The core principle involves studying how differences in plaintext pairs propagate through the cipher. By identifying and exploiting high-probability sequences of intermediate differences (differential characteristics), an attacker can recover the secret key.

The key-averaged probability of a characteristic is calculated under the assumption that the round keys that are xored with the state after each round are uniformly random and independent. This assumption has often been used together with the hypothesis of stochastic equivalence [LMM91]. It states that the probability of a differential for a specific key can be estimated by its average probability over all possible keys. However, in practice the probability can substantially vary among different keys [DR07]. To overcome these

Table 1: Overview of differential attacks on the three SPEEDY variants, with their reported and recalculated complexities (data, time, memory). CP and CC indicate data collection in the chosen-plaintext and chosen-ciphertext settings, respectively. r is the number of rounds attacked. Time complexities are measured in the equivalent of encryptions, and memory costs are expressed in block size units.

r	Claimed			Revisited			Ref.
	Data	Time	Memory	Data	Time	Memory	
SPEEDY-7-192 with security claim $2^{192}, 2^{192}$							
5	2^{109}CC	2^{109}	2^{109}	2^{107}	2^{107}	2^{109}	[YJZZ22]
6	2^{158}CC	2^{158}	2^{158}	2^{154}CC	2^{154}	2^{158}	[YJZZ22]
7	2^{187}CC	2^{188}	2^{42}		invalid		[BDBN23]
7	2^{187}CC	2^{187}	2^{36}	2^{190}CC	2^{191}	2^{36}	[WNL+23]
7	2^{187}CP	2^{187}	2^{156}	2^{190}CP	2^{191}	2^{156}	[WNL+23]
7	2^{183}CP	2^{185}	2^{156}	2^{183}CP	2^{185}	2^{156}	Section 5.2
7	2^{183}CC	2^{185}	2^{36}	2^{183}CC	2^{185}	2^{36}	Section 5.2
SPEEDY-6-192 with security claim $2^{128}, 2^{128}$							
5.5	2^{122}CC	2^{128}	2^{42}	2^{127}	2^{133}	2^{42}	[BDBN23]
6	2^{122}CC	2^{152}	2^{42}	2^{127}	2^{157}	2^{42}	[BDBN23]
SPEEDY-5-192 with security claim $2^{64}, 2^{128}$							
4	2^{61}CC	2^{120}	2^{83}	2^{59}CC	2^{118}	2^{83}	[WNL+23]
5	2^{102}CC	2^{108}	2^{42}	2^{107}	2^{113}	2^{42}	[BDBN23]

assumptions, Beyne and Rijmen [BR22] introduced quasidifferential trails, which allow for calculating fixed-key differential probabilities. This approach accounts for key-dependent behavior, leading to more accurate probability estimates.

Contributions. In this paper, we conduct a thorough analysis of the literature on the differential cryptanalysis of the SPEEDY block cipher family. In the interest of scientific reproducibility, and because cryptanalytic results can directly impact confidence in a cipher's security, it is important that published attacks are correct. Several differential attacks have been proposed against different round-reduced variants of SPEEDY [YJZZ22, BDBN23, WNL+23, Zha24]. However, upon closer inspection, we find that nearly all of them rely on incorrect assumptions resulting in flawed probability estimates. In some cases, the attack becomes invalid due to the use of a differential characteristic with a probability of zero. In others, the attack still works, but for reasons that differ from the original explanation. We also identify some instances where the attack is actually stronger than reported, due to underestimating the characteristics' probability.

To explain and resolve these flawed probability estimates, we begin by focusing on the underlying one-round differential characteristics[1] used in the reported multi-round characteristics. Several of these one-round characteristics have probability zero — even in the key-averaged setting. This renders the multi-round characteristics used in the key-recovery attacks invalid. We provide a detailed analysis of this issue, which arises from the incorrect application of assumptions used to estimate key-averaged probabilities in steps without key addition. Furthermore, we propose a method to accurately and efficiently calculate one-round probabilities using quasidifferential trails. Surprisingly, this is the first application of quasidifferential trails to compute differential probabilities in block ciphers

[1]There are *one-round characteristics*, because every SPEEDY round contains two S-box layers.

that compose nonlinear layers without intermediate key addition.

Using our corrected calculation of one-round differential characteristic probabilities, we develop a tool based on mixed-integer linear programming to search for optimal one-round characteristics. Using these optimal one-round characteristics we search for multi-round characteristics and find improved 5-round differential characteristics. We also identify some alternatives for the one-round characteristics with probability zero that were used in the multi-round characteristics of [YJZZ22, BDBN23, WNL+23, Zha24].

Building on this tool, we reevaluate all previously published key-recovery attacks on SPEEDY, carefully analyzing the corresponding characteristics and reevaluating the security margin. Table 6 presents an overview of the original and updated attack complexities, including data, time, and memory. For instance, we show that the 7-round attack by [BDBN23], originally based on a differential with 409 characteristics, is invalid due to the majority of those characteristics having a probability of zero. Our reevaluation reduces the actual differential probability by a factor of more than 2^{10}, rendering the attack invalid. Conversely, the 5-round attack of [YJZZ22] turns out to be more effective than reported, with the corrected characteristic offering a probability three times higher than initially claimed. Finally, we propose a new full-round chosen-plaintext key-recovery attack on SPEEDY-7-192. We adopt the key-recovery strategy from [WNL+23] and use a new 5.5-round truncated differential, extended with 1.5 key-recovery rounds. This results in a key-recovery attack with a data complexity of 2^{183}, a time complexity of 2^{185}, and a memory complexity of 2^{156}. Using the same truncated differential we can also mount a chosen-chipertext attack with same data and time complexity and a memory complexity of 2^{36}.

Earlier and parallel work. This paper is an extended version of our note [BN24] from February 2024, where we showed that the characteristic used in [BDBN23] has probability zero. While this paper was being finalized, a parallel and independent report was uploaded to ePrint by Boura et al. [BDG+25]. They also identify a valid differential characteristic for SPEEDY-7-192 and mount a key-recovery attack. Although we have not performed a detailed comparison, our characteristic has higher probability and leads to a slightly more efficient attack.

2 Preliminaries

2.1 SPEEDY

SPEEDY [LMMR21] is a family of low-latency block ciphers designed by Leander et al. SPEEDY-r-$6l$ denotes one instance of the family with block and key size $6l$ and r rounds. The designers suggest $l = 32$ as default with the number of rounds $r \in \{5, 6, 7\}$. Since all reported differential characteristics are on these instances, we denote them by SPEEDY-r-192. The internal state x is represented as a 32×6 array of 192 bits. We denote the bit at row i ($0 \leq i < 32$) and column j ($0 \leq j < 6$) by $x_{i,j}$, following the notation of the designers [LMMR21]. The zero bit is the most significant bit. The designers claim 128-bit security for SPEEDY-6-192 and 192-bit security for SPEEDY-7-192.

Round Function. The s^{th} round function $R_s : \mathbb{F}_2^{32\times 6} \to \mathbb{F}_2^{32\times 6}$ consists of five different operations in the order shown in Figure 1. We denote the input (resp. output) to each of the described operations as a vector x (resp. y). In the last round some linear operations are omitted, and there is an additional round key addition.

- AddRoundKey (ARK). The 192-bit round key k_s is XORed to the state. The details of the key schedule to derive these round keys from the 192-bit master key K are given below.

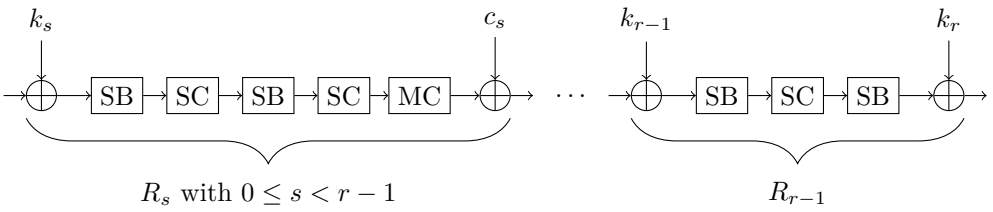

Figure 1: The round function of SPEEDY for the first $r-1$ rounds and the last round.

- SubBox (SB). Applies a 6-bit b-Box S to each row of the state. The S-box is given in Table 2.
- ShiftColumns (SC). The j-th column of the state is rotated upwards by j positions: $y_{i,j} = x_{i+j,j}$.
- MixColumns (MC). A circulant binary matrix is multiplied with each column of the state: $y_{i,j} = x_{i,j} \oplus x_{i+\alpha_1,j} \oplus x_{i+\alpha_2,j} \oplus x_{i+\alpha_3,j} \oplus x_{i+\alpha_4,j} \oplus x_{i+\alpha_5,j} \oplus x_{i+\alpha_6,j}$ with $\alpha = (1, 5, 9, 15, 21, 26)$.
- AddRoundConstant (ARC). A 192-bit constant c_s is XORed to the whole state. We refer to [LMMR21] for more details about c_s.

Table 2: The 6-bit S-box used in SPEEDY.

x_0x_1	$x_2x_3x_4x_5$															
	.0	.1	.2	.3	.4	.5	.6	.7	.8	.9	.a	.b	.c	.d	.e	.f
0.	08	00	09	03	38	10	29	13	0c	0d	04	07	30	01	20	23
1.	1a	12	18	32	3e	16	2c	36	1c	1d	14	37	34	05	24	27
2.	02	06	0b	0f	33	17	21	15	0a	1b	0e	1f	31	11	25	35
3.	22	26	2a	2e	3a	1e	28	3c	2b	3b	2f	3f	39	19	2d	3d

Key Schedule. The key schedule takes a 192-bit master key, used as the zero-th round key k_0. The following round keys are derived in a linear way from this master key using a bit permutation. We refer to [LMMR21] for the details of this bit permutation.

Round Structure. We emphasize an unconventional aspect of the SPEEDY round function: two nonlinear SubBox layers are separated solely by a linear ShiftColumns operation, without any key addition in between. This design choice is the reason why even key-averaged differential probability calculations fail, and it will be the focus of our detailed analysis.

2.2 Differential Cryptanalysis

The main principle of differential cryptanalysis is to investigate the propagation of a given plaintext difference through a cipher. A differential (a, b) consist of an input difference a and an output difference b. To set up a differential attack, one first finds good differentials, i.e. with probability much higher than for a random function. These differentials can then be used as a distinguisher and for key-recovery.

2.2.1 Differential characteristics

A differential characteristic $(a_1, a_2, \ldots, a_{r+1})$ with $a_1 = a$ and $a_{r+1} = b$ specifies, besides the input and output difference, every intermediate difference for a composition of r functions. The probability of a differential (a, b) is calculated by summing the probability of all possible characteristics with the same input and output difference.

The key-averaged probability of a characteristic is determined by multiplying the probabilities over each round, assuming that the round keys are uniformly random and independent [LMM91]. However, this assumption is not valid due to the use of a key schedule. More importantly, the key-average probability does not determine the data-complexity of a differential attack. For this reason, the hypothesis of stochastic equivalence [LMM91] is often used. This hypothesis states that the probability of a differential for a specific key can be estimated through the average probability over all possible keys. However, in practice the probability can vary substantially across different keys [DR07, BR22].

2.2.2 Quasidifferential trails

To overcome the hypothesis of stochastic equivalence and being able to calculate the probability of a differential characteristic in a fixed-key model, quasidifferential trails were introduced by Beyne and Rijmen in 2022 [BR22]. For a function $F\colon \mathbb{F}_2^n \to \mathbb{F}_2^n$ the quasidifferential transition matrix is defined by

$$D^F_{(v,b),(u,a)} = \big(2 \Pr_x[v^T F(x) = u^T x \mid F(x+a) = F(x) + b] - 1\big) \Pr_x[F(x+a) = F(x) + b].$$

A quasidifferential trail is a sequence of mask-difference pairs $((u_1, a_1), (u_2, a_2), \ldots, (u_{r+1}, a_{r+1}))$ with $(u_1, a_1) = (u, a)$ and $(u_{r+1}, a_{r+1}) = (v, b)$. The correlation of a quasidifferential trail is defined as $\prod_{i=1}^r D^{F_i}_{(u_{i+1},a_{i+1}),(u_i,a_i)}$. By summing over all possible intermediate masks we can find the probability of a characteristic:

$$\Pr_{x_0}[\textstyle\bigwedge_{i=1}^r F_i(x_i + a_i) = F_i(x_i) + a_{i+1}] = \displaystyle\sum_{u_2,\ldots,u_r} \prod_{i=1}^r D^{F_i}_{(u_{i+1},a_{i+1}),(u_i,a_i)},$$

with $x_i = F_i(x_{i-1})$ for all $i \geq 1$. Lastly, the probability of the differential is found by summing over all possible quasidifferential trails:

$$D^F_{(0,b),(0,a)} = \sum_{\substack{u_2,\ldots,u_r \\ a_2,\ldots,a_r}} \prod_{i=1}^r D^{F_i}_{(u_{i+1},a_{i+1}),(u_i,a_i)}.$$

In the remainder of this paper we will often make use of [BR22, Theorem 4.2 (2)] which states that if we have a quasidifferential trail with differences $a_1, \ldots, a_{r+1}$ and maximum correlation p, that if we find a quasidifferential trail with the same differences but with correlation $-p$, then the probability of the characteristic is zero.

Theorem 1. *[BR22, Theorem 4.2] For a function $F = F_r \circ \cdots \circ F_1$ and a sequence of differences $a_1, \ldots, a_{r+1}$ with correlation p (as quasidifferential trail), it holds that:*

(1) If $(u_1, a_1), \ldots, (u_{r+1}, a_{r+1})$ is a quasidifferential trail with correlation $(-1)^b p$ where $b \in \{0, 1\}$, then for any quasidifferential trail $((v_1, a_1), \ldots, (v_{r+1}, a_{r+1}))$ with correlation c, the correlation of the quasidifferential trail $((u_1 + v_1, a_1), \ldots, (u_{r+1} + v_{r+1}, a_{r+1}))$ is $(-1)^b c$.

(2) If the correlations of any number of quasidifferential trails with differences $a_1, \ldots, a_{r+1}$ and correlation $\pm p$ sum to zero, then the probability of the characteristic $(a_1, \ldots, a_{r+1})$ is zero.

We end this section with a general property of quasidifferential transition matrices that will be used in Section 3. This is a new result.

Lemma 1. *Let $((u,a),(v,b))$ be a quasidifferential for a function $F\colon \mathbb{F}_2^n \to \mathbb{F}_2^n$ which consists of differences (a,b) and masks (u,v). If $u^T a \neq v^T b$, then*

$$D^F_{(v,b),(u,a)} = 0$$

Proof. Recall the definition of the quasidifferential transition matrix:

$$D^F_{(v,b),(u,a)} = \big(2 \Pr_x[u^T x = v^T F(x) \mid F(x+a) = F(x) + b] - 1\big) \Pr_x[F(x+a) = F(x) + b].$$

Rewriting the first factor yields

$$\begin{aligned}
&2 \Pr_x\big[u^T x = v^T F(x) \mid F(x+a) = F(x) + b\big] - 1 \\
&= \Pr_x\big[u^T x + v^T F(x) = 0 \mid \cdots\big] - \Pr_x\big[u^T x + v^T F(x) = 1 \mid \cdots\big] \\
&= \Pr_x\big[u^T (x+a) + v^T F(x+a) = 0 \mid \cdots\big] - \Pr_x\big[u^T x + v^T F(x) = 1 \mid \cdots\big] \qquad (1) \\
&= \Pr_x\big[u^T x + v^T F(x) = u^T a + v^T b \mid \cdots\big] - \Pr_x\big[u^T x + v^T F(x) = 1 \mid \cdots\big]
\end{aligned}$$

In step (1), we leverage the symmetry of the relation $F(x+a) = F(x) + b$, which ensures that if x satisfies the condition, so does $x+a$. If $u^T a + v^T b = 1$, then the result is zero. □

3 Revisiting reported differential characteristics

The literature on the differential cryptanalysis of SPEEDY contains several differential characteristics used in various key-recovery attacks, see Table 6. Section 5 provides a more detailed discussion of the key-recovery attacks. In this section, we re-examine the reported probability of the characteristics in Table 6 by analyzing the one-round characteristics they comprise.

We start this section by illustrating, for a one-round characteristic, the discrepancy between the actual probability and the probability reported in the literature — even in the key-averaged case. As discussed in the introduction, this discrepancy arises from an incorrect application of assumptions used to estimate the key-averaged probability in steps without a key addition. Specifically, the assumption that the probability over two consecutive nonlinear SubBox applications can be computed as the product of their individual probabilities is flawed. We show how quasidifferential trails can be used to calculate the probability of such one-round characteristics.

To demonstrate the scope of this issue, Table 3 lists all one-round characteristics extracted from multi-round differential characteristics in the literature where such discrepancies occur between the actual and reported probabilities. We begin by examining the characteristics with probability zero, and then proceed to those with a nonzero but different probability. Finally, for the characteristics with probability zero, we propose alternative characteristics for the same differential when possible.

3.1 Differential properties of one round

The core challenge in analyzing SPEEDY's differential characteristics arises from the operation $SB \circ SC \circ SB$, where two nonlinear SubBox applications are separated only by the ShiftColumns operation and no key addition. Previous works [BDBN23, WNL+23, Zha24, YJZZ22] calculated the probability by multiplying the probabilities over the two SubBox applications. However, we will show that this approach is not valid when differences occur

Table 3: Overview of discussed one-round characteristics which are part of multi-round characteristics in the literature. We give for each characteristic the reported probability, which corresponds to the correlation of the trail with zero masks, the probability calculated using quasidifferential trails, the number of rounds of the multi-round characteristic it is part of and the reference.

Name	Correlation with zero masks	Probability	# rounds	Reference
A_0	$2^{-70.42}$	0	4.5	[YJZZ22, Figure 1]
A_1	2^{-42}	0	4 and 5	[Zha24, Table 6, 7]
			4.5 and 5.5	[BDBN23, Figure 8, 4]
A_2	$2^{-44.2}$	0	6	[Zha24, Figure 4]
A_3	$2^{-64.83}$	0	5.5	[WNL+23, Figure 3]
A_4	$2^{-29.41}$	0	3.5	[BDBN23, Figure 7]
A_5	$2^{-42.23}$	0	5.5	[BDBN23, Figure 4]
A_6	$2^{-44.87}$	0	5.5	[BDBN23, Figure 4]
B_0	$2^{-45.08}$	$2^{-44.97}$	4.5	[YJZZ22, Figure 1]
			5.5	[WNL+23, Figure 3]
B_1	2^{-42}	$2^{-40.42}$	3.5	[YJZZ22, Figure 4]
			3.5	[WNL+23, Figure 7]
B_2	$2^{-32.05}$	$2^{-32.66}$	5.5	[BDBN23, Figure 4]
B_3	$2^{-44.87}$	$2^{-44.49}$	3.5	[BDBN23, Figure 7]
S_0	$2^{-67.8}$	$2^{-66.42}$	alternative for A_0 and A_3, see §3.4	
S_1	2^{-49}	2^{-49}	alternative for A_1, see §3.4	
S_2	$2^{-47.8}$	$2^{-47.83}$	alternative for A_2, see §3.4	

within six rows of each other. The first SubBox application constrains possible outputs, which are only linearly shifted by the ShiftColumns, affecting the input distribution of the second SubBox application and potentially affecting the probability.

Quasidifferential trails allow us to accurately calculate the probability by accounting for this relation between the output after the first SubBox application and the input before the second SubBox application. We note that the theory of plateau characteristics [DR07] cannot fully describe this issue: the S-box has differentials with probability $6/2^6$, for which the input and outputs satisfying the differential do not form an affine space (which is necessary for plateau characteristics).

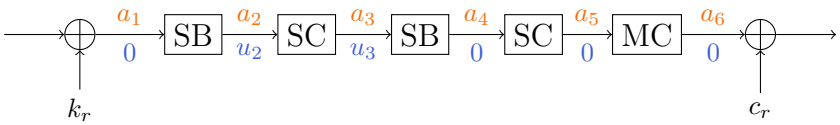

Figure 2: Quasidifferential trail over one round of SPEEDY.

Given a one-round differential characteristic $(a_1, \ldots, a_6)$ as shown in Figure 2, we determine its probability by identifying all quasidifferential trails and summing their correlations. We begin by eliminating mask values that lead to correlation zero without considering the specifics of the S-box. The input and output masks are set to zero. Since ShiftColumns and MixColumns are invertible linear operations, the intermediate masks before the MixColumns and before the ShiftColumns operations must also be zero. To

determine possible values for the masks u_2 and u_3, we apply the general principle that for inactive S-boxes, the input mask (respectively output mask) must be zero when the output mask (respectively input mask) is zero if the correlation is nonzero. Hence, the nonzero bits in u_2 should appear in nonzero rows of a_2 and the nonzero bits of u_3 have to align with the nonzero rows of a_3. However, we also have $u_3 = SC(u_2)$, which further restricts the possible values of u_2 and u_3. The preceding argument does not restrict the mask bits corresponding to nonzero difference bits. Lemma 1 from Section 2.2.2 provides some conditions on the masks that hold for all S-boxes. Taking into account all of these constraints, it is feasible to enumerate all quasidifferential trails with nonzero correlation for sparse differential characteristics. Note that the resulting probability is independent of the key because the input and output masks for the round function are set to zero.

Example 1. As a first example, we analyze the differences in the first and last row of characteristic A_1 (see Figure 3), in which the same differential appears twice. For the first SubBox application, we have the differential $(0x01, 0x10)$ with probability 2^{-3}, and after ShiftColumns, for the second SubBox we have the differential $(0x10, 0x04)$ with the same probability. The product of these probabilities corresponds to the correlation of the quasidifferential trail with all-zero masks, i.e. 2^{-6} (only taking these two active rows into account). However, looking at quasidifferential trails with nonzero masks, we find quasidifferentials $(0x0, 0x01) \to (0x10, 0x10)$ with correlation 2^{-3} and $(0x10, 0x10) \to (0x0, 0x04)$ with correlation -2^{-3}. The corresponding quasidifferential trail has correlation -2^{-6}, not taking into account other rows. Hence, by Theorem 1, the probability is zero. This corresponds to the fact that all valid differential pairs in the first SubBox produce outputs with the highest bit set to one (preserved by ShiftColumns), whereas the second SubBox transition requires the highest input bit to be zero, leading to a contradiction. This contradiction is what the nonzero mask captures, leading to a probability of zero. ▷

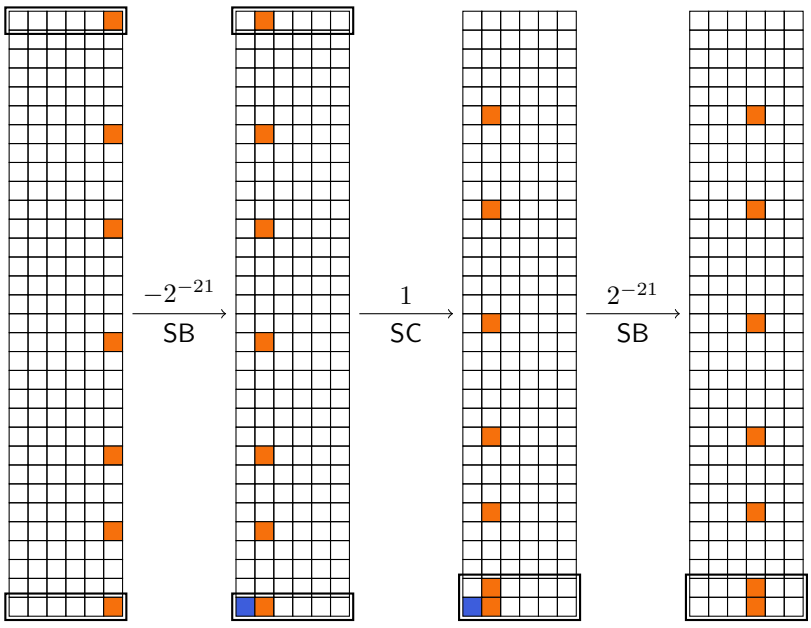

Figure 3: Quasidifferential trail for characteristic A_1 with an opposite correlation of the quasidifferential trail with all zero masks. The correlations are computed using (indices in hexadecimal): $D^S_{(20,10),(00,01)} = -2^{-3}$ and $D^S_{(00,04),(20,10)} = 2^{-3}$.

The phenomenon discussed in Example 1 is more general. Figure 4 illustrates the propagation of differences and masks through the composition $SB \circ SC \circ SB$. Nonzero

difference bits are represented by orange cells, and potentially nonzero mask bits are represented by the light blue cells. The potentially nonzero mask bits are spread across six different rows by the ShiftColumns operation.

For trails with nonzero correlation and nonzero masks, the mask after the second S-box layer can only be zero if at least two of the six S-boxes corresponding to the gray-colored rows in the rightmost state are active. If only one S-box is active, then the correlation is zero (for non-zero masks) due to Lemma 1. To identify the input differences for which quasidifferential trails with nonzero masks can have nonzero correlation, we trace back the gray cells in the input of the second SubBox (SB) application through the inverse ShiftColumns (SC) and the inverse SubBox (SB). This shows that the second nonzero row of the input difference should be within a range of six rows from the first active row.

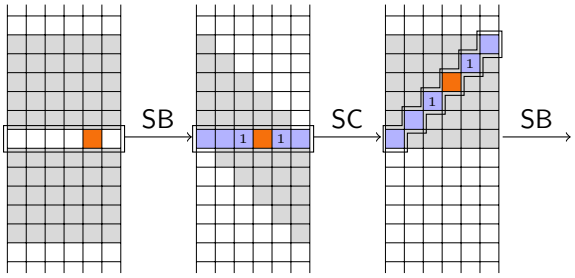

Figure 4: Propagation of a difference through $SB \circ SC \circ SB$.

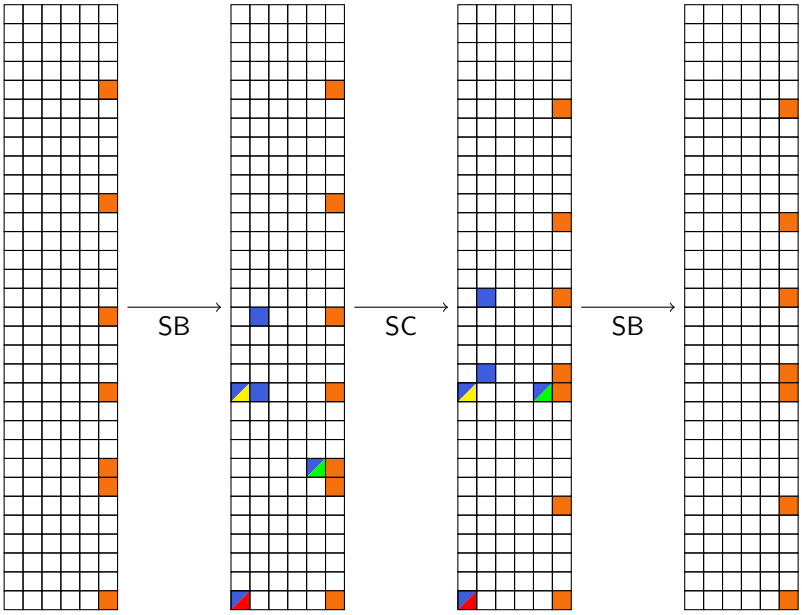

Figure 5: A one-round characteristic with probability 2^{-53}, based on eight quasidifferential trails with correlation 2^{-56}. These trails have nonzero masks after the first S-box layer, with $(u_{20,0}, u_{24,4}, u_{31,0})$ arbitrary in $\mathbb{F}_2^3$ and all other mask bits fixed to zero.

Example 2. The phenomenon described above can also lead to characteristics with higher probability than what might be expected based on multiplying one-round probabilities. Figure 5 depicts such a characteristic. Orange cells indicate a difference bit equal to one. The correlation of the quasidifferential trail with zero masks is 2^{-56}. Using the process

described above, we identify five potentially nonzero bits in the mask after the first SubBox application, indicated by blue cells in Figure 5. By examining all 32 possible values for this intermediate mask, we find eight quasidifferential trails with a nonzero correlation of 2^{-56}. The masks of these eight trails are illustrated in Figure 5 using different colors. Specifically, the eight masks are obtained by setting three mask bits to all possible values $(u_{20,0}, u_{24,4}, u_{31,0})$ in $\mathbb{F}_2^3$, and all other bits to zero. The probability is obtained by summing all correlations, i.e. 2^{-53} for this characteristic. ▷

3.2 Characteristics with probability zero

The characteristics $A_0, A_1, \ldots, A_6$ in Table 3 all have probability zero. Characteristic A_1 was already analyzed in Example 1. Characteristic A_0 is shown in Figure 6. It was estimated to have a probability of $3 \cdot 2^{-72}$ in [YJZZ22], corresponding to the correlation of the quasidifferential trail with all masks equal to zero. Similar to the analysis of characteristic A_1 in Example 1, we identify one more trail with a correlation of $-3 \cdot 2^{-72}$. The nonzero mask bits in this trail are represented by blue cells in Figure 6. Therefore, by Theorem 1, the probability of the entire characteristic is zero.

For the other characteristics A_2, A_3, A_4, A_5 and A_6, we can also identify a quasidifferential trail with correlation the opposite of to the correlation of the trail with all-zero masks. Hence, the probability of all these characteristics is zero. For details on these characteristics, see Figures 3, 12, 13, 14, 15, and 16 in the annex.

However, even if a specific characteristic has probability zero, this does not necessarily imply that the corresponding differential has probability zero. In Section 3.4, we explore alternative characteristics corresponding to the same differential.

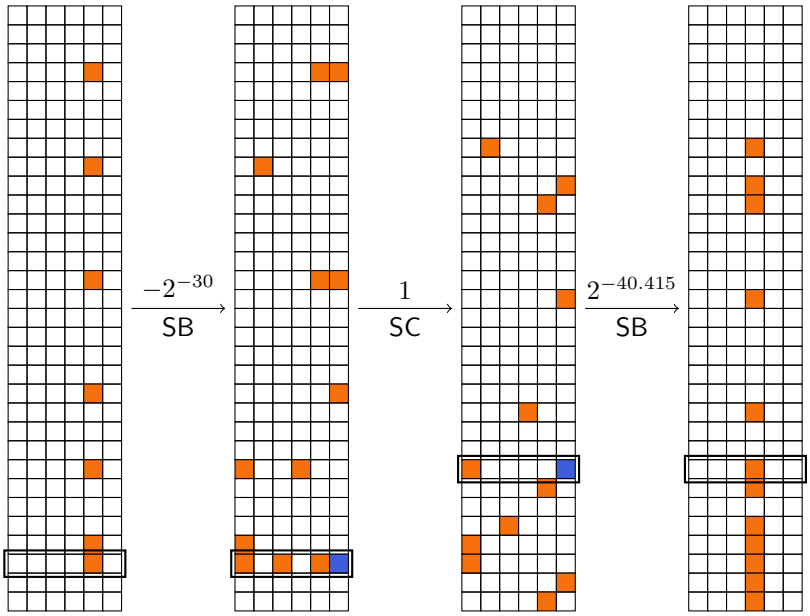

Figure 6: Quasidifferential trail for characteristic A_0 with the correlations computed using (indices in hexadecimal) $D^S_{(01,2a),(00,02)} = -2^{-5}$ and $D^S_{(00,04),(01,20)} = 2^{-4}$.

3.3 Characteristics with a different but nonzero probability

For the characteristics B_0, B_1, B_2 and B_3 we find multiple quasidifferential trails, resulting in a nonzero probability different from the reported probability. For characteristic B_0 we

find four quasidifferential trails, see Figure 7. The sum of the correlations of these found trails is $2^{-44.91} + 2^{-46.49} + 2^{-46.49} + 2^{-48.08} = 2^{-45.08}$, the probability of the characteristic. For the details of characteristics B_1, B_2 and B_3, we refer to Figures 17, 18 and 19 in the appendix.

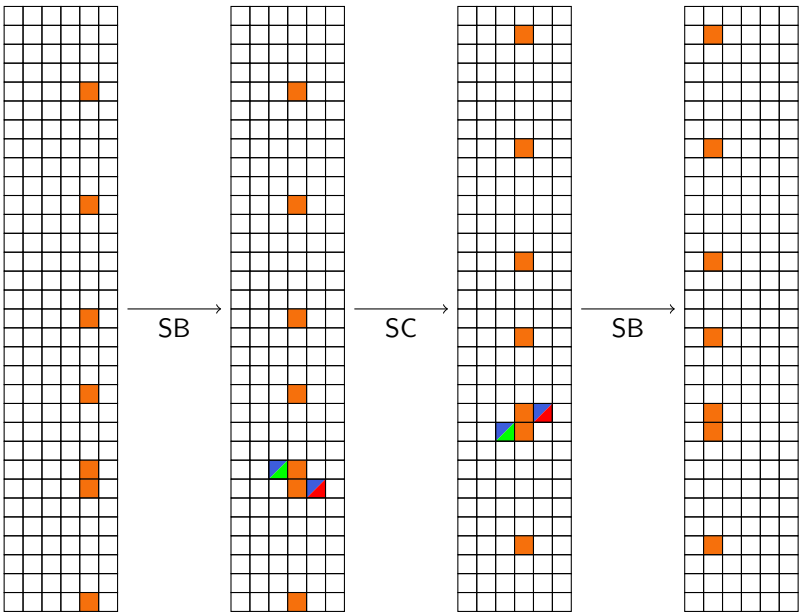

Figure 7: Characteristic B_0 with probability $p = 2^{-45.08}$ with 4 quasidifferential trails with correlations $2^{-44.91}, 2^{-46.49}, -2^{-46.49}, -2^{-48.08}$ for nonzero bits $(u_{24,2}, u_{25,4}) \in \mathbb{F}_2^2$ in the intermediate mask after the first SubBox application.

3.4 Alternatives for characteristics with probability zero

In Section 4, we search for multi-round differential characteristics by first constructing numerous one-round differential characteristics. Based on these results, we identify alternative characteristics with nonzero probability for the same differentials as the probability-zero characteristics from Table 3. Specifically, we found alternatives for characteristics A_0, A_1, A_2 and A_3. These are illustrated in Figure 8 below and Figures 22 and 23 in the appendix.

The impact of these results on key-recovery attacks is discussed in Section 5, in particular for attacks that rely on multi-round characteristics containing at least one one-round characteristic with probability zero. Since a single zero-probability one-round characteristic renders the entire multi-round characteristic invalid, this directly impacts the effectiveness of such attacks. Where possible, we substitute the probability-zero characteristics with the alternatives we found and re-evaluate the key-recovery attacks.

4 Searching for improved differential characteristics

In this section, we develop methods to find improved multi-round characteristics for `SPEEDY`. Our approach is based on finding a shortest path in a graph with edges corresponding to one-round characteristics.

To construct one-round characteristics, we first analyze the differential behavior of the MixColumns operation and demonstrate that the necessary condition for the existence of nontrivial quasidifferential trails discussed in Section 3.1, is always satisfied for two

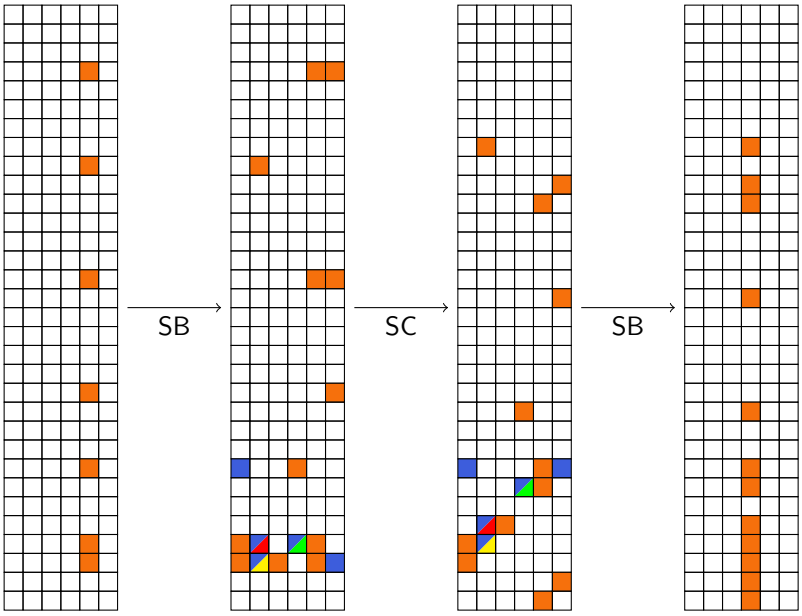

Figure 8: Characteristic S_0 with probability $2^{-66.42}$ based on eight quasidifferential trails with correlations $2^{-67.83}$ (twice), $2^{-69.42}$ (four times), $-2^{-69.42}$ (twice) with nonzero mask bits $(u_{28,1}, u_{28,3}, u_{29,1})$ in $\mathbb{F}_2^3$ in the mask after the first S-box layer.

rounds of SPEEDY. Furthermore, we show that in most cases, this condition leads to an actual discrepancy in the probability. We then construct an MILP model that searches for optimal one-round characteristics using quasidifferential trails to estimate probabilities.

For the multi-round differential characteristics we find, we extend our analysis to estimate the probability in the fixed-key model. We identify a 5-round characteristic with a correlation of $2^{-160.64}$ for the quasidifferential trail with all-zero masks. The actual differential probability of this characteristic is between $2^{-157.542}$ and $2^{-157.518}$, depending on the key.

4.1 Differential properties of the MixColumns operation

In Section 3.1, we examined differential characteristics over a single round and explained the necessity of using quasidifferential trails to accurately compute probabilities. Although we identified when quasidifferential trails with nonzero masks exist, we did not analyze how often this would happen in the context of a multi-round characteristic. This section provides a more detailed discussion of that aspect. To do so, we focus on the MixColumns operation. According to [LMMR21], the MixColumns operation has a branch number of eight, which is realized for inputs of Hamming weight one. In the following lemma, we take a closer look at MixColumns and the Hamming weights of its inputs and outputs.

Lemma 2 (MixColumns)**.** *For all x_1 and x_2 such that $x_2 = MC(x_1)$, at least one of $HW(x_1)$ or $HW(x_2)$ is greater than or equal to six, where $HW(x)$ denotes the Hamming weight of the state x.*

Proof. Using an MILP model, we can calculate for all possible x_1 and x_2 with $HW(x_1) \leq 6$ or $HW(x_2) \leq 6$, what the corresponding lowest Hamming weight of x_2, respectively x_1 is. The results are given in Table 4. □

Lemma 2 demonstrates that within two rounds of SPEEDY, either the first or the

Table 4: Lowest Hamming weight of the input (or output) of the MixColumns operation when the output (or input) Hamming weight is less than or equal to six.

HW input	1	2	3	4	5	6	7	8	9	12	19
HW output	7	8	7	8	7	6	5	4	3	2	1

second round inevitably activates at least six S-boxes. Given that there are 32 rows, this ensures that two nonzero rows will be positioned within six rows of each other, satisfying the necessary but not sufficient condition required for the existence of nontrivial quasidifferential trails as explained in Section 3.1.

To investigate this further, we examine how often this results in discrepancies of the probabilities. We limit ourselves to the scenario where the input and the output differences are one-bit differences located in the same column. This simplifies the analysis and matches the approach in the next subsection, where the multi-round characteristics are designed to start and end each round with only one active column. Examining all these relevant cases, we identify 26050 possible one-round differential characteristics. Among these, 27.88% have a probability of zero, 27.94% exhibit a higher probability, and 9.04% have a lower but nonzero probability compared to calculations based only on the quasidifferential trail with all-zero masks.

4.2 Finding multi-round differential characteristics

In this section, we search for differential characteristics by extending the methods from [BDBN23, WNL+23] and incorporating calculations based on quasidifferential trails.

Following [BDBN23], we focus on optimal one-round characteristics in which only a single column is active during the MixColumns operation. This restriction is motivated by the observation that multiple active columns, when propagated through the ShiftColumns operation, quickly result in numerous active rows, significantly reducing the overall probability of the characteristic. Our search starts from all possible pairs $(x, MC(x))$ satisfying $HW(x) + HW(MC(x)) \leq 12$. Since any of the 32 possible rotations within a column yields the same result, we select a representative from each equivalence class, ultimately identifying 64 distinct pairs. Note that, [BDBN23] constrained both input and output Hamming weights to at most seven, and [WNL+23] considered only cases with an input Hamming weight of at most eight.

For each identified pair, we generate six possible tuples $(x, MC(x), c)$, where $c \in \{0, \ldots, 5\}$ indicates the column in which the difference occurs. We then search for the optimal one-round differential characteristic for all pairs of such tuples defining the input and output differences. This is based on an MILP model to search for the quasidifferential trail with the highest correlation. Once a candidate is found, we compute its probability by calculating by enumerating all quasidifferential trails. If the resulting probability is zero, we continue with the next-best characteristic for the same differential and repeat the process. This continues until we find a characteristic with nonzero probability or exhaust all options. For every characteristic with nonzero probability, we also examine the next-best characteristics for the same differential. This analysis is limited to those characteristics for which the corresponding zero-mask quasidifferential trail has correlation at most eight times lower than the highest-correlation trail.

Using our MILP model, we identified 13799 differential characteristics, of which 6894 have probability zero. By analyzing the next-best characteristics, we found that 1569 of these zero-probability characteristics could be replaced with other characteristics having nonzero — but lower — probability.

All identified one-round differentials are used to construct a directed graph, where each node represents a difference. An edge from one node to another exists if there is a

one-round differential connecting the corresponding tuples, and the edge weight reflects the probability of that differential. We then use a branch-and-bound algorithm to search for a minimum-weight (i.e. maximum probability) path in this graph, corresponding to an optimal multi-round characteristic. In particular, our goal is to identify the best 4-round differential characteristic. With the key-recovery already in mind, we require that the characteristic starts with a difference that allows for a MixColumns transition from Hamming weight one to seven. This facilitates the addition of a high-probability round at the beginning of the characteristic. Additionally, we minimize the number of active rows after the final MixColumns operation, as the characteristic will also be extended with half a round at the end.

The optimal 4-round characteristic we identified is shown in Figure 9. To be able to use the same key-recovery strategy as in [WNL+23] we prepend the same first-round characteristic (shown in Figure 10). In Section 5, we also extend this characteristic with half a round at the end to obtain a truncated differential that can be used in a full-round key-recovery attack.

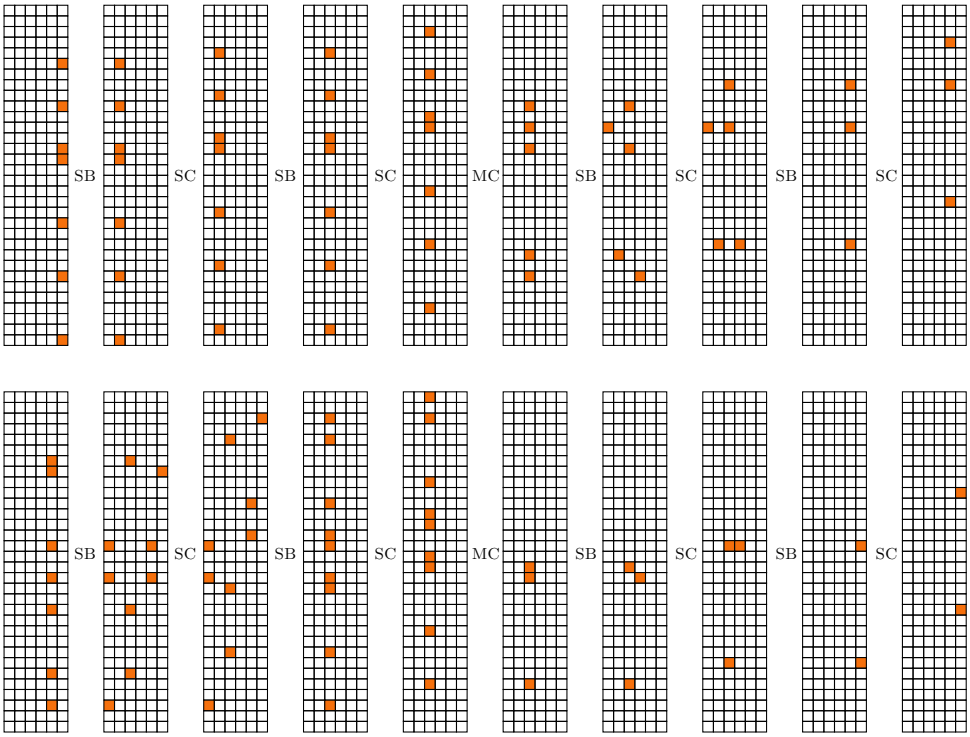

Figure 9: Four-round differential characteristic.

4.3 Probability calculation

To calculate the fixed-key probability of multi-round differential characteristics, we develop an MILP model designed to identify quasidifferential trails corresponding to a specific differential characteristic. To reduce the search space, we impose certain constraints and adopt an iterative approach to enumerate as many quasidifferential trails as feasible.

Quasidifferential trails with key-independent correlations. We first find quasidifferential trails with zero masks before and after the MixColumns operation, i.e. trails with key-independent correlations. These can also be found by using the earlier results for the

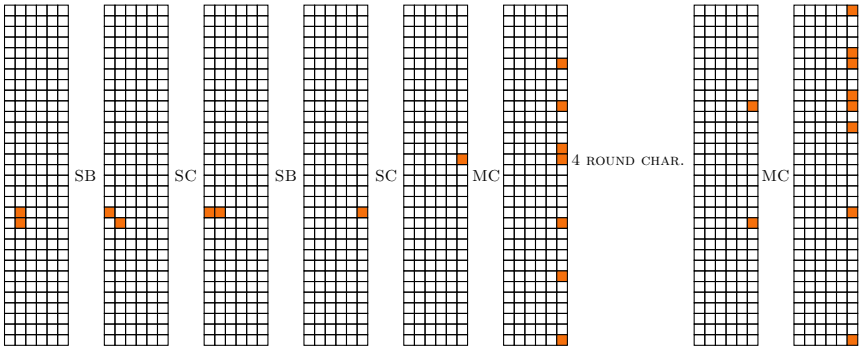

Figure 10: First-round extension for the four-round characteristic from Figure 9

one-round characteristics B_1, B_4 and B_5 (see Figures 17, 20 and 21 in the appendix) that are present in our 5-round characteristic.

For the 5-round characteristic in Figure 10, we identify 256 quasidifferential trails. The probability is obtained by summing the correlations of all identified trails (see Table 5), yielding a total value of $2^{-157.53}$. For comparison, the correlation of the quasidifferential trail with all-zero masks is $2^{-160.64}$.

Table 5: Number of found quasidifferential trails of each correlation.

Correlation	# Trails	Correlation	# Trails	Correlation	# Trails
$2^{-160.64}$	8	$2^{-161.64}$	8	$-2^{-162.23}$	8
$2^{-163.23}$	8	$-2^{-163.23}$	16	$2^{-163.81}$	8
$-2^{-164.23}$	8	$2^{-164.23}$	8	$2^{-164.81}$	16
$-2^{-164.81}$	8	$-2^{-165.40}$	8	$-2^{-165.81}$	16
$2^{-165.81}$	8	$2^{-166.40}$	8	$-2^{-166.40}$	16
$2^{-166.81}$	8	$2^{-167.40}$	16	$-2^{-167.40}$	8
$2^{-167.98}$	8	$-2^{-167.98}$	8	$2^{-168.40}$	8
$-2^{-168.98}$	16	$2^{-168.98}$	8	$-2^{-169.98}$	8
$2^{-170.57}$	8	$2^{-171.57}$	8		

Quasidifferential trails with key-dependent correlations. Once all key-independent quasidifferential trails have been determined, we relax the constraints on the masks. Rather than enforcing conditions on the masks, we now consider all quasidifferential trails with a correlation of at least 2^{-172}. This yields 104 key-dependent quasidifferential trails, which we use to identify potential dependencies on the subkeys.

Based on the masks of the additional trails, we refine our correlation calculation by enumerating all trails with the same masks at the corresponding SubKey additions. This refined approach allows us to account for a total of 11648 quasidifferential trails and compute a more accurate estimate of the probability of the characteristic. Our final estimate of the probability of the characteristic is

$$2^{-157.53} - (-1)^{K^1} 2^{-164.79} - (-1)^{K^2} 2^{-178.15} - (-1)^{K^3} 2^{-168.70},$$

where the linear combinations K^1, K^2 and K^3 of key bits are given by

$$K^1 = k_{4,5} + k_{5,5} + k_{10,5} + k_{11,5} + k_{15,5} + k_{19,5} + k_{31,5},$$
$$K^2 = k_{2,1} + k_{8,1} + k_{11,1} + k_{20,1} + k_{21,1} + k_{28,1} + k_{31,1},$$
$$K^3 = k_{8,2} + k_{13,2} + k_{14,2} + k_{19,2} + k_{20,2} + k_{24,2} + k_{28,2}.$$

The first term in this formula arises from the key-independent quasidifferential trails discussed in the previous paragraph. The subsequent terms originate from key-dependent trails. Due to the linear key schedule, we can determine the master key bits corresponding to linear combinations of subkey bits. Depending on the key, the probability varies slightly, between $2^{-157.542}$ and $2^{-157.518}$.

Multiple characteristics. Starting from our best five-round characteristic found using the method described in Section 4.2, we identify, for two of its five rounds, other one-round characteristics that correspond to the same differential over one round. By combining these, we construct three distinct 5-round characteristics for the same overall differential. Taking this into account refines the estimated probability to $2^{-157.27}$. We note that this refinement is based on a limited search and is not the result of exhaustive enumeration of all possible characteristics.

5 Key Recovery attack

In this section, we examine the impact of our analysis in the previous sections on key-recovery attacks. We begin by reevaluating attacks in the literature, and then present an improved full-round key-recovery attack on SPEEDY-7-192 using our 5-round differential characteristic from Section 4.

5.1 Revisiting reported key-recovery attacks

We revisit the key-recovery attacks summarized in Table 6. For each attack, we analyze the characteristic that was used and assess its probability. Where possible, we considered the alternative characteristics introduced in Section 3.4 to recalculate the cost as accurately as possible. We also examine the probability of the entire differential in other cases, since key-recovery attacks depend on differentials rather than on individual characteristics. The reported characteristics are typically dominant, but in some cases the effect of multiple characteristics is not negligible.

Attacks on SPEEDY-7-192 The 7-round key-recovery attack by Boura et al. [BDBN23] is depends on the one-round characteristics A_1, A_5, and A_6. Alternatives for these characteristics have significantly lower probability. Since the original paper took into account multiple differential characteristics, we also examined the possibility of other characteristics with a good probability. The authors of [BDBN23] identified 409 differential characteristics for the same differential. In our attempt to reproduce the results, using the same criteria as the paper, we identified 584 characteristics. Due to the absence of a detailed listing of the characteristics of the paper, a direct comparison was not possible.

Using the method from Section 3.1, we find that 545 of the identified characteristics have probability zero. Summing the nonzero probabilities shows that the overall differential probability of the 5-round characteristic is $2^{-180.60}$. This is $2^{10.65}$ times lower than the reported probability, rendering the attack invalid.

We did not identify an alternative for characteristic A_3, which is used in the key-recovery attacks from [WNL+23] on 7 rounds of SPEEDY-7-192. However, to construct an alternative characteristic replacing the full 5.5-round characteristic used in [WNL+23], we could

Table 6: Overview of differential key-recovery attacks on the three SPEEDY variants from the literature, with their reported costs (data, time, memory). CP and CC indicate data collection in the chosen-plaintext and chosen-ciphertext settings, respectively. r is the number of rounds attacked. Time complexities are measured in the number of encryptions, and memory costs are expressed in block size units.

r	Data	Time	Memory	Ref.	Contains
SPEEDY-7-192 with security claim $(2^{192}, 2^{192})$					
5	$2^{108.91}$CC	$2^{108.95}$	$2^{108.91}$	[YJZZ22]	B_1
6	$2^{158.04}$CC	$2^{158.06}$	$2^{158.04}$	[YJZZ22]	A_0, B_0
7	$2^{187.28}$CC	$2^{187.84}$	2^{42}	[BDBN23]	A_1, A_5, A_6, B_2
7	$2^{186.53}$CC	$2^{187.39}$	2^{36}	[WNL+23]	A_3, B_0
7	$2^{186.53}$CP	$2^{187.39}$	2^{156}	[WNL+23]	A_3, B_0
SPEEDY-6-192 with security claim $(2^{128}, 2^{128})$					
5.5	$2^{121.65}$CC	$2^{127.8}$	2^{42}	[BDBN23]	A_1, A_5, B_2
6	$2^{121.65}$CC	$2^{151.67}$	2^{42}	[BDBN23]	A_1, A_5, B_2
SPEEDY-5-192 with security claim $(2^{64}, 2^{128})$					
4	2^{61}CC	$2^{119.69}$	2^{83}	[WNL+23]	B_1
5	$2^{101.65}$CC	$2^{107.8}$	2^{42}	[BDBN23]	A_4, B_3

use S_0, with the additional requirement that the round preceding A_3 be modified using characteristic E_0 (see Figure 24 in the appendix). These adjustments reduce the probability of the characteristic by $2^{3.21}$. Although the attack remains valid, the complexities are worse than claimed. In this case no multiple characteristics where used, and we did not search for other 5.5-round characteristics that could improve the probability.

The first characteristic, S_0 (Figure 8), serves as a valid alternative to characteristic A_0. In fact, its probability is higher. Notably, when only considering the quasidifferential trail with zero masks for the characteristic S_0, we find a higher correlation than for the characteristic A_0. The attack on the 6-round SPEEDY-7-192 in [YJZZ22] stays valid and the complexities would improve slightly because the probability of the differential characteristic improves by a factor of $2^{4.026}$.

The attack on 5-round SPEEDY-7-192 presented in [YJZZ22] uses the characteristic B_1. Since the probability of the characteristics is three times higher than its reported value, the cost of the attack is also lower.

Attacks on SPEEDY-6-192 The 5.5-round and 6-round attacks on SPEEDY-6-192 presented in [BDBN23] use characteristic A_5, for which we were unable to find a good alternative. However, since the paper also utilizes multiple characteristics we again searched all characteristics and found 5 characteristics of which only one has a nonzero probability. The differential has a probability of approximately $2^{-88.12}$, which is $2^{5.10}$ times lower than the reported probability.

Attacks on SPEEDY-5-192 The 5-round attack on SPEEDY-5-192 presented in [BDBN23] employs the characteristic A_4, for which we were unable to find a good alternative. For the same differential, we found 54 candidate characteristics of which 14 have nonzero probability. The differential has a probability of $2^{-132.37}$, which is 2^7 times worse than the reported probability.

The 4-round attack on SPEEDY-5-192 presented in [WNL+23] relies on the characteristic B_1, which improves the attack because its probability is three times higher than reported

in [WNL+23] .

Table 7: Overview of the reported differential attacks on the three SPEEDY variants, with their claimed and recalculated complexities (data, time, memory). CP and CC indicate data collection in the chosen-plaintext and chosen-ciphertext settings, respectively. r is the number of rounds attacked. Time complexities are measured in the number of encryptions, and memory costs are expressed in block size units. See also Table 1.

r	Claimed			Revisited			Ref.
	Data	Time	Memory	Data	Time	Memory	
SPEEDY-7-192 with security claim $2^{192}, 2^{192}$							
5	$2^{108.91}$CC	$2^{108.95}$	$2^{108.91}$	$2^{107.33}$CC	$2^{107.37}$	$2^{108.91}$	[YJZZ22]
6	$2^{158.04}$CC	$2^{158.06}$	$2^{158.04}$	$2^{154.01}$CC	$2^{154.03}$	$2^{158.04}$	[YJZZ22]
7	$2^{187.28}$CC	$2^{187.84}$	2^{42}		invalid		[BDBN23]
7	$2^{186.53}$CC	$2^{187.39}$	2^{36}	$2^{189.74}$CC	$2^{190.60}$	2^{36}	[WNL+23]
7	$2^{186.53}$CP	$2^{187.39}$	2^{156}	$2^{189.74}$CP	$2^{190.60}$	2^{156}	[WNL+23]
SPEEDY-6-192 with security claim $2^{128}, 2^{128}$							
5.5	$2^{121.65}$CC	$2^{127.8}$	2^{42}	$2^{126.75}$CC	$2^{132.9}$	2^{42}	[BDBN23]
6	$2^{121.65}$CC	$2^{151.67}$	2^{42}	$2^{126.75}$CC	$2^{156.77}$	2^{42}	[BDBN23]
SPEEDY-5-192 with security claim $2^{64}, 2^{128}$							
4	2^{61}CC	$2^{119.69}$	2^{83}	$2^{59.42}$CC	$2^{118.11}$	2^{83}	[WNL+23]
5	$2^{101.65}$CC	$2^{107.8}$	2^{42}	$2^{106.75}$CC	$2^{112.90}$	2^{42}	[BDBN23]

5.2 New key-recovery attack

We adopt the key-recovery strategy from [WNL+23] to construct a chosen-plaintext attack on the full 7 rounds of SPEEDY-7-192. This attack leverages a 5.5-round truncated differential, preceded by one key-recovery round and extended with half a key-recovery round.

In Section 4, we identified the best 4-round characteristic, which we already extended forward by one round. Figure 11 illustrates how this is extended with half a round to get a full 5.5-round truncated differential (light blue cells represent undetermined difference bits, constrained so that not all bits in the same row are zero). The figure also shows the key-recovery rounds, with red cells indicating difference bits that are fixed to zero, which occurs with probability approximately $2^{-3.04}$. Taking all components into account, the overall probability of the truncated differential used in our attack is approximately $2^{-181.56}$.

In the first phase of the attack, we collect chosen plaintexts and use ciphertext differences to reduce the number of possible values of subkey k_7. The remaining key bits are recovered using the sieving functions from [WNL+23].

Data collection. Based on the input difference, we construct structures of 2^{156} plaintexts with rows 0,1,2,3,19,31 constant and all possible values for the remaining rows (indicated by dark red cells in Figure 11). These plaintexts are encrypted, and the ciphertexts are stored in a hash table indexed by the 156 zero bits of the ciphertext difference. Given that there are approximately $2^{2\times156-1}$ possible plaintext-ciphertext pairs, we obtain an average of 2^{155} collisions. These serve as the pairs in the sieving step.

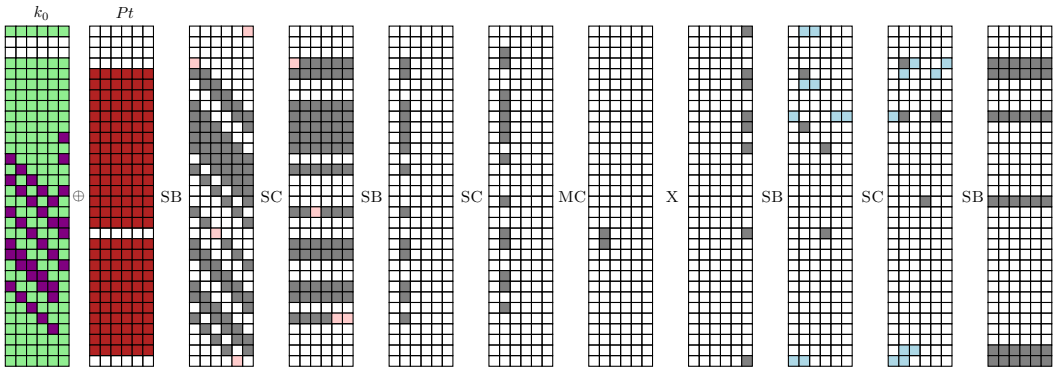

Figure 11: Key recovery attack on SPEEDY-7-192, with X the 5-round characteristic from Figure 10

Number of possible candidates for the targeted part of k_7. From the ciphertext difference of the collected pairs, we deduce the possible candidates for part of k_7. For a fixed input and output difference of the S-box, on average only one value is possible for the corresponding 6-bit part of the subkey k_7. In cases with undetermined difference bits, we analyze the second-to-last SubBox application. The following truncated differentials over the S-box have to be considered:

1. (on row 31) $\texttt{000001} \to **\texttt{0000}$ has two instantiations with nonzero probability: output difference 100000 or 010000.
2. (on row 0 and 5) $\texttt{000001} \to \texttt{0}**\texttt{000}$ has two instantiations with nonzero probability: output difference 010000 or 001000.
3. (on row 8) $\texttt{000001} \to *\texttt{000}**$ has four instantiations with nonzero probability: output difference 000001, 100000, 000011, or 100011.

Thus, on average we obtain a total of $2 \times 2 \times 2 \times 4 = 32$ different possible candidates for the 36-bit targeted part of k_7. Due to the linear key schedule, the guessed key bits of k_7 can be linearly mapped to the key bits of k_0, as indicated by the purple cells in Figure 11. For each candidate for the 36-bit part of k_7, we still need to guess 145 bits of k_0.

Sieving candidate keys. Using the FirstSboxSieve and SecondSBoxSieve functions from [WNL+23], we reduce the possible values of k_0. The number of remaining key bits of k_0 to be guessed are shown in green in Figure 11. First, sieving of pairs and keys is based on the output difference of the first S-box layer. A second step then considered the second S-box to further refine the remaining key candidates. After sieving, we are left with $2^{138.46}$ pairs and $2^{6.50}$ 180-bit partial keys per pair on average, yielding $2^{144.96}$ candidate keys for the 180-bit partial master key per structure. Finally we still have to exhaustively search for the remaining 12 bits, for each candidate key. Additional details are given in Appendix A.

Cost estimate. The time required for data processing is 2^{156} per structure. Accessing the hash table storing the 2^{155} pairs is counted as one encryption. The 7-round SPEEDY utilizes $7 \times 2 \times 32 = 448$ S-boxes. Hence, the time required per structure is $2^{163.81}$ S-box operations. The exhaustive search step has a time complexity of $2^{156.96}$. The total time complexity is estimated as follows:

$$2^{156} + 2^{155} + \frac{2^{163.81}}{448} + 2^{156.96} \approx 2^{157.965}.$$

Since our differential has probability $2^{-181.56}$, the number of required structures is $2^{26.56}$ (the exponent is $181.56 - 156 + 1$).

For the full attack, we obtain a time complexity of $2^{184.53}$, a data complexity of $2^{182.56}$ and a memory complexity of 2^{156}. To mount a chosen-ciphertext attack, we can reuse the same distinguisher and follow a similar strategy, achieving identical time and data complexities while reducing the memory complexity to 2^{36}.

Acknowledgements. Tim Beyne is supported by a junior postdoctoral fellowship from the Research Foundation – Flanders (FWO) with reference number 1274724N.

References

[BDBN23] Christina Boura, Nicolas David, Rachelle Heim Boissier, and María Naya-Plasencia. Better steady than speedy: Full break of SPEEDY-7-192. In Carmit Hazay and Martijn Stam, editors, Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part IV, volume 14007 of Lecture Notes in Computer Science, pages 36–66. Springer, 2023.

[BDG+25] Christina Boura, Patrick Derbez, Baptiste Germon, Rachelle Heim Boissier, and María Naya-Plasencia. SPEEDY: Caught at last. Cryptology ePrint Archive, Paper 2025/890, 2025.

[BN24] Tim Beyne and Addie Neyt. Note on the cryptanalysis of speedy. Cryptology ePrint Archive, Report 2024/262, 2024.

[BR22] Tim Beyne and Vincent Rijmen. Differential cryptanalysis in the fixed-key model. In Yevgeniy Dodis and Thomas Shrimpton, editors, Advances in Cryptology – CRYPTO 2022, Part III, volume 13509 of Lecture Notes in Computer Science, pages 687–716, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[BS91] Eli Biham and Adi Shamir. Differential cryptanalysis of DES-like cryptosystems. Journal of Cryptology, 4(1):3–72, January 1991.

[DR07] Joan Daemen and Vincent Rijmen. Plateau characteristics. IET Inf. Secur., 1(1):11–17, 2007.

[LMM91] Xuejia Lai, James L. Massey, and Sean Murphy. Markov ciphers and differential cryptanalysis. In Donald W. Davies, editor, Advances in Cryptology – EUROCRYPT'91, volume 547 of Lecture Notes in Computer Science, pages 17–38, Brighton, UK, April 8–11, 1991. Springer Berlin Heidelberg, Germany.

[LMMR21] Gregor Leander, Thorben Moos, Amir Moradi, and Shahram Rasoolzadeh. The SPEEDY family of block ciphers engineering an ultra low-latency cipher from gate level for secure processor architectures. IACR Transactions on Cryptographic Hardware and Embedded Systems, 2021(4):510–545, 2021.

[WNL+23] Jinliang Wang, Chao Niu, Qun Liu, Muzhou Li, Bart Preneel, and Meiqin Wang. Cryptanalysis of SPEEDY. In Leonie Simpson and Mir Ali Rezazadeh Baee, editors, ACISP 23: 28th Australasian Conference on Information Security and Privacy, volume 13915 of Lecture Notes in Computer Science, pages 124–156, Brisbane, QLD, Australia, July 5–7, 2023. Springer, Cham, Switzerland.

[YJZZ22] Qingyuan Yu, Keting Jia, Guangnan Zou, and Guoyan Zhang. Differential cryptanalysis of round-reduced SPEEDY family. In Yi Deng and Moti Yung, editors, Information Security and Cryptology - 18th International Conference, Inscrypt 2022, Beijing, China, December 11-13, 2022, Revised Selected Papers, volume 13837 of Lecture Notes in Computer Science, pages 272–291. Springer, 2022.

[Zha24] Lei Zhang. Observations on the branch number and differential analysis of SPEEDY. Des. Codes Cryptogr., 92(5):1175–1199, 2024.

A Details on the key-recovery attack from Section 5.2

Table 8 lists the steps of the sieving process. In this appendix, we explain in more detail what happens in every step.

The first phase calculates the set $\mathcal{K}^7$, containing on average 32 possible candidates for the 36-bit targeted part of k_7. To allow for meaningful comparison, the estimation of the time complexity follows the strategy of [WNL+23]. For each of the 32 cases, two S-box applications are counted; for the other active S-boxes with a fixed input difference, only one S-box application is counted. We find a time complexity of $32 \times 2 + 1 = 65$ S-box evaluations per pair.

In the next phase, we apply the `FirstSboxSieve` function to each relevant row i over the first S-box layer. For each row, this yields 6-bit key candidates along with corresponding S-box outputs. Since not all output bits are relevant for the next stage, we reduce the outputs to their difference pattern since these are used in the sieving for the second S-box layer. As an example, consider the application of `FirstSboxSieve` to rows 20 through 25. Each of these yields a set of key candidates and corresponding S-box outputs. After the ShiftColumns operation, the input to the second S-box on row 20, consists of the first bit from the possible outputs of row 20, the second bit from the possible outputs of row 21, and so on. This transformed input is then used in the `SecondSboxSieve` function, which performs a further filtering step based on the propagation through the second S-box layer.

In Table 8, the column I_k contains the number of newly guessed key bits of k_0 for the corresponding row (shown in green in Figure 11). The filtering probability is estimated based on the number of inactive bits in the S-box output. The columns labeled by N_p and N_k contain the remaining number of plaintext pairs and remaining number of key candidates, respectively.

The time complexity of each sieving step is determined by N_p and N_k. The total time complexity of the sieving phase is obtained by summing the complexities across all lines, yielding a total time complexity for the sieving of approximately $2^{163.81}$.

Table 8: Details of the sieving step from Section 5.2.

`FirstSboxSieve` on row	`SecondSboxSieve` on row	I_k	Time complexity	Filtering probability	N_p	N_k
Constructing $\mathcal{K}^7$			65×2^{155}	1	2^{155}	2^5
18		3	$2^{159.0}$	$2^{-4.0}$	$2^{155.0}$	$2^{4.0}$
20		4	$2^{160.0}$	$2^{-4.0}$	$2^{155.0}$	$2^{4.0}$
21		3	$2^{159.0}$	$2^{-3.0}$	$2^{155.0}$	$2^{4.0}$
22		3	$2^{159.0}$	$2^{-3.0}$	$2^{155.0}$	$2^{4.0}$
23		4	$2^{160.0}$	$2^{-4.0}$	$2^{155.0}$	$2^{4.0}$
24		3	$2^{159.0}$	$2^{-3.0}$	$2^{155.0}$	$2^{4.0}$
25		4	$2^{160.0}$	$2^{-2.0}$	$2^{155.0}$	$2^{6.0}$
	20	0	$2^{162.0}$	2^{-6}	$2^{155.0}$	1
26		5	$2^{161.0}$	$2^{-3.0}$	$2^{155.0}$	$2^{2.0}$
	21	0	$2^{158.0}$	2^{-6}	$2^{151.0}$	1
17		4	$2^{156.0}$	$2^{-4.0}$	$2^{151.0}$	1
19		4	$2^{156.0}$	1	$2^{151.0}$	$2^{4.0}$
	17	0	$2^{156.0}$	$2^{-5.54}$	$2^{149.46}$	1
27		5	$2^{155.46}$	$2^{-3.0}$	$2^{149.46}$	$2^{2.0}$

Continued on next page

Table 8 — continued from previous page

`FirstSboxSieve` on row	`SecondSboxSieve` on row	I_k	Time complexity	Filtering probability	N_p	N_k
28		5	$2^{155.46}$	$2^{-3.0}$	$2^{149.46}$	$2^{4.0}$
29		6	$2^{156.46}$	$2^{-3.0}$	$2^{149.46}$	$2^{7.0}$
	24	0	$2^{157.46}$	2^{-6}	$2^{149.46}$	$2^{1.0}$
30		6	$2^{156.46}$	$2^{-4.0}$	$2^{149.46}$	$2^{3.0}$
	25	0	$2^{153.46}$	2^{-6}	$2^{146.46}$	1
16		4	$2^{151.46}$	$2^{-4.0}$	$2^{146.46}$	1
15		3	$2^{150.46}$	$2^{-3.0}$	$2^{146.46}$	1
14		4	$2^{151.46}$	$2^{-2.0}$	$2^{146.46}$	$2^{2.0}$
13		5	$2^{152.46}$	$2^{-1.0}$	$2^{146.46}$	$2^{6.0}$
	13	0	$2^{153.46}$	2^{-6}	$2^{146.46}$	1
11		6	$2^{153.46}$	$2^{-1.0}$	$2^{146.46}$	$2^{5.0}$
12		4	$2^{151.46}$	$2^{-1.0}$	$2^{146.46}$	$2^{8.0}$
	11	0	$2^{155.46}$	2^{-6}	$2^{146.46}$	$2^{2.0}$
10		5	$2^{152.46}$	$2^{-2.0}$	$2^{146.46}$	$2^{5.0}$
	10	0	$2^{152.46}$	2^{-6}	$2^{145.46}$	1
9		6	$2^{152.46}$	$2^{-2.0}$	$2^{145.46}$	$2^{4.0}$
	9	0	$2^{150.46}$	2^{-6}	$2^{143.46}$	1
8		6	$2^{150.46}$	$2^{-2.0}$	$2^{143.46}$	$2^{4.0}$
	8	0	$2^{148.46}$	2^{-6}	$2^{141.46}$	1
7		6	$2^{148.46}$	$2^{-3.0}$	$2^{141.46}$	$2^{3.0}$
	7	0	$2^{145.46}$	2^{-6}	$2^{138.46}$	1
31		6	$2^{145.46}$	1	$2^{138.46}$	$2^{6.0}$
0		6	$2^{145.46}$	1	$2^{138.46}$	$2^{12.0}$
	27	0	$2^{151.46}$	$2^{-5.67}$	$2^{138.46}$	$2^{6.33}$
4		6	$2^{145.46}$	$2^{-4.0}$	$2^{138.46}$	$2^{8.33}$
5		6	$2^{145.46}$	$2^{-4.0}$	$2^{138.46}$	$2^{10.33}$
6		6	$2^{145.46}$	$2^{-4.0}$	$2^{138.46}$	$2^{12.33}$
	4	0	$2^{151.79}$	2^{-6}	$2^{138.46}$	$2^{6.33}$
3		6	$2^{145.46}$	1	$2^{138.46}$	$2^{12.33}$
	3	0	$2^{151.79}$	$2^{-5.83}$	$2^{138.46}$	$2^{6.5}$

B Characteristics with probability zero

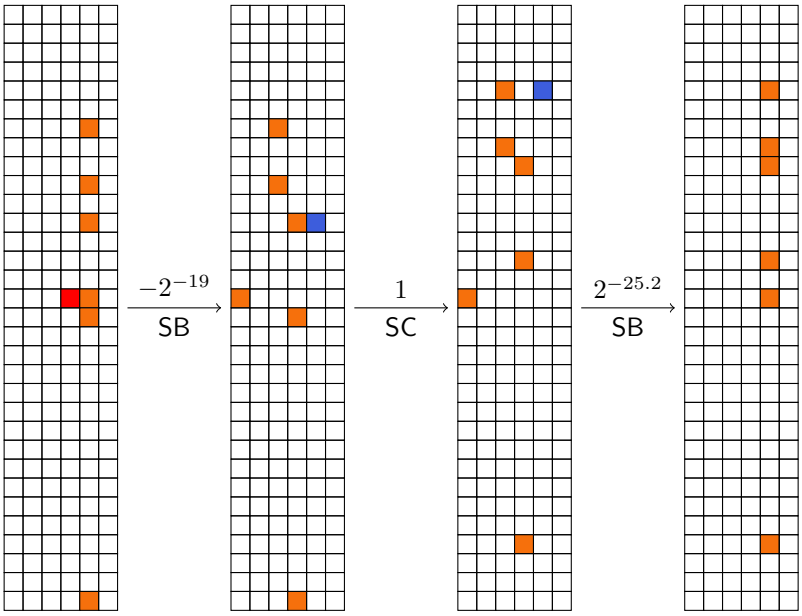

Figure 12: Quasidifferential trail for characteristic A_2 with an opposite correlation of the quasidifferential trail with all zero masks. The correlations are computed using (indices in hexadecimal) $D^S_{(02,04),(00,02)} = -2^{-3}$ and $D^S_{(00,02),(02,08)} = 2^{-5}$.

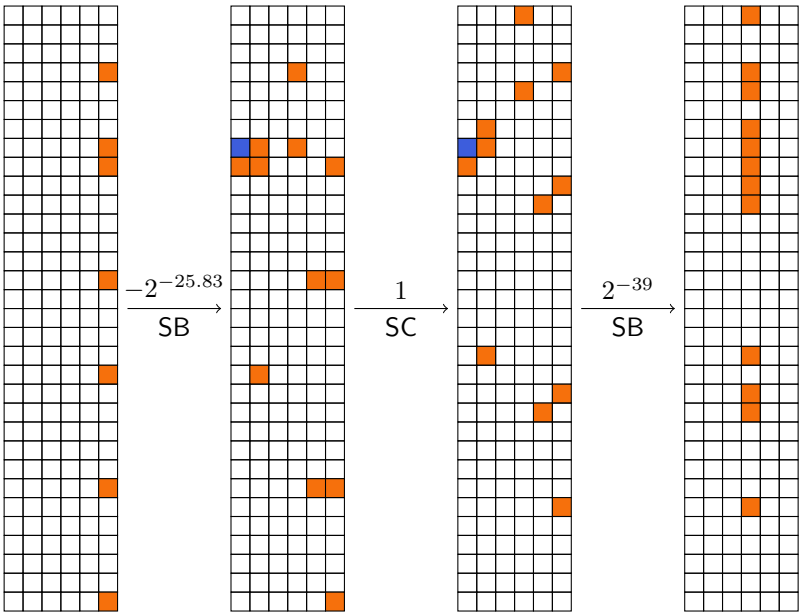

Figure 13: Quasidifferential trail for characteristic A_3 with an opposite correlation of the quasidifferential trail with all zero masks. The correlations are computed using (indices in hexadecimal) $D^S_{(20,14),(00,01)} = -2^{-5}$ and $D^S_{(00,04),(20,10)} = 2^{-3}$.

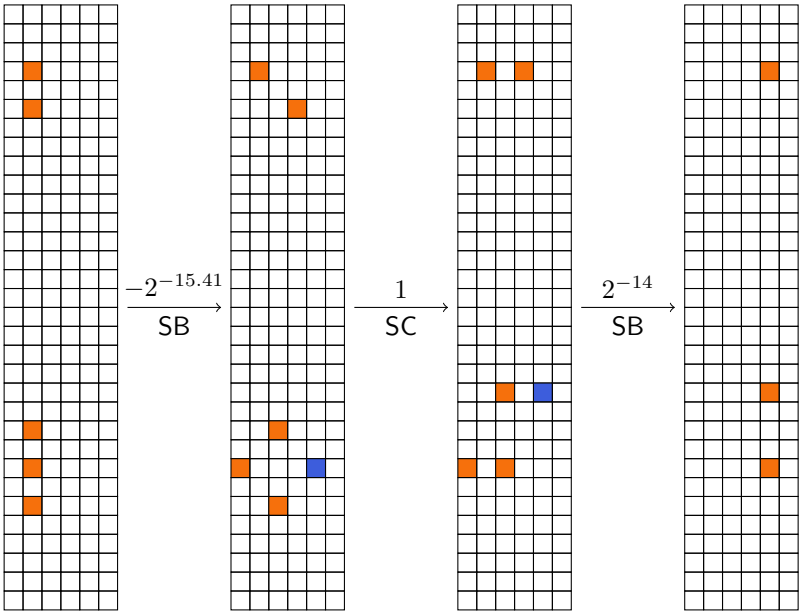

Figure 14: Quasidifferential trail for characteristic A_4 with an opposite correlation of the quasidifferential trail with all zero masks. The correlations are computed using (indices in hexadecimal) $D^S_{(02,20),(00,10)} = -2^{-3}$ and $D^S_{(00,02),(02,08)} = 2^{-5}$.

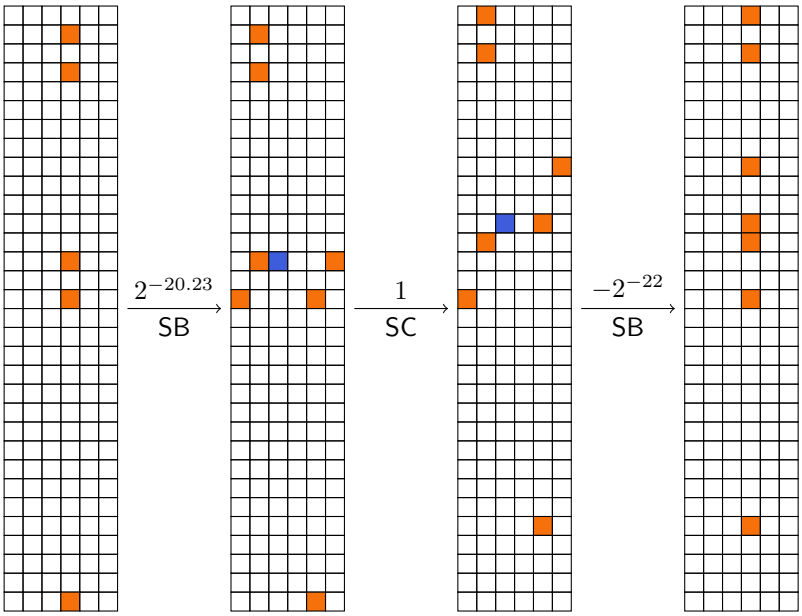

Figure 15: Quasidifferential trail for characteristic A_5 with an opposite correlation of the quasidifferential trail with all zero masks. The correlations are computed using (indices in hexadecimal) $D^S_{(08,11),(00,04)} = 2^{-5}$ and $D^S_{(00,04),(08,02)} = -2^{-3}$.

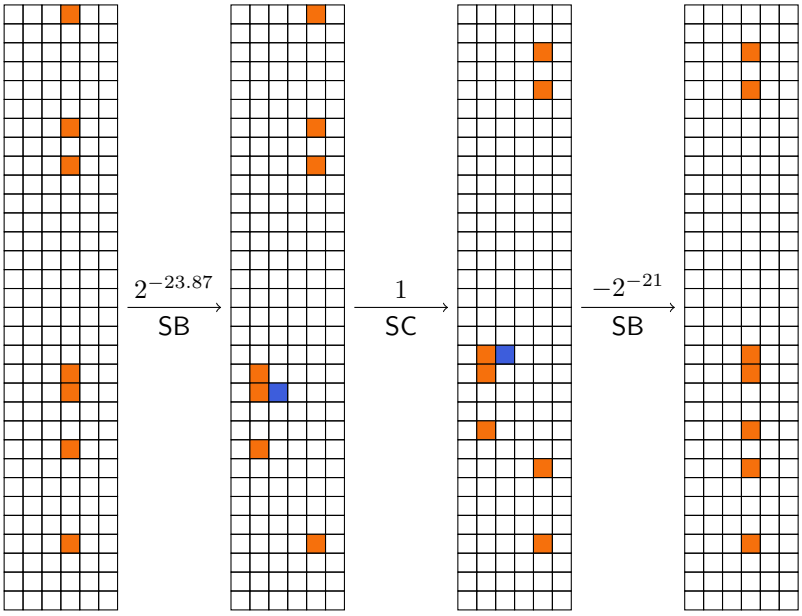

Figure 16: Quasidifferential trail for characteristic A_6 with an opposite correlation of the quasidifferential trail with all zero masks. The correlations are computed using (indices in hexadecimal) $D^S_{(08,10),(00,04)} = 2^{-3.415}$ and $D^S_{(00,04),(08,10)} = -2^{-3}$.

C Characteristics with a different but nonzero probability

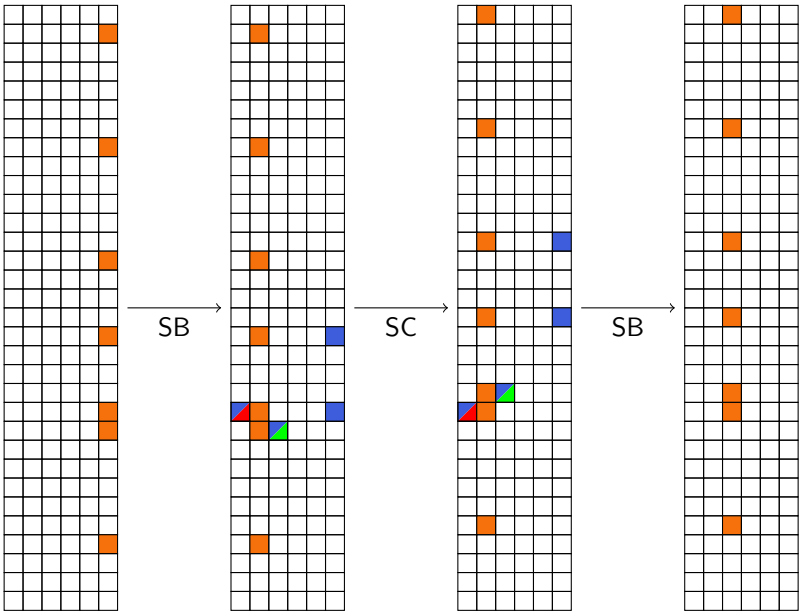

Figure 17: Characteristic B_1 with probability $2^{-40.42}$ based on four quasidifferential trails with correlations 2^{-42}, 2^{-43}, 2^{-42} and 2^{-43} for nonzero bits $(u_{21,0}, u_{22,2})$ in $\mathbb{F}_2^2$ of the intermediate mask after the first SubBox application.

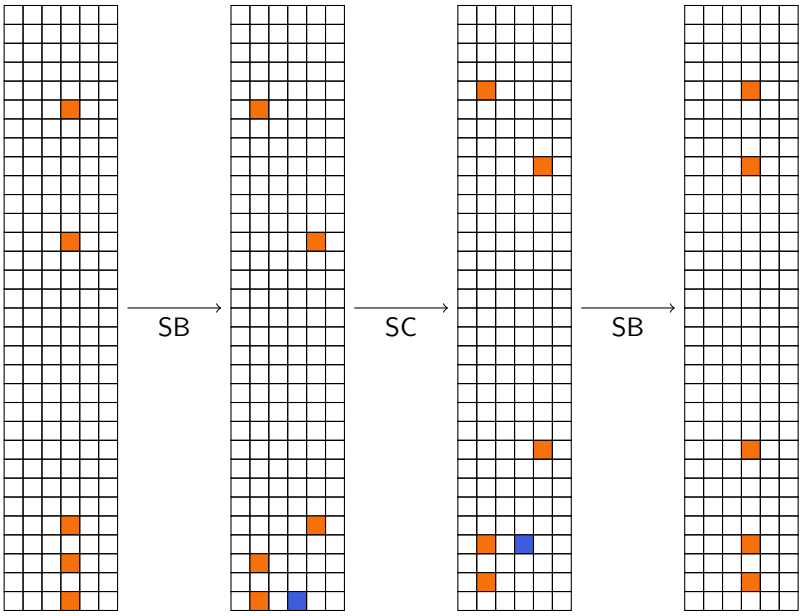

Figure 18: Characteristic B_2 with probability $2^{-32.66}$ based on two quasidifferential trails with correlations $2^{-32.08}$ and $2^{-33.66}$ for nonzero bit $u_{31,3}$ in $\mathbb{F}_2$ of the intermediate mask after the first SubBox application.

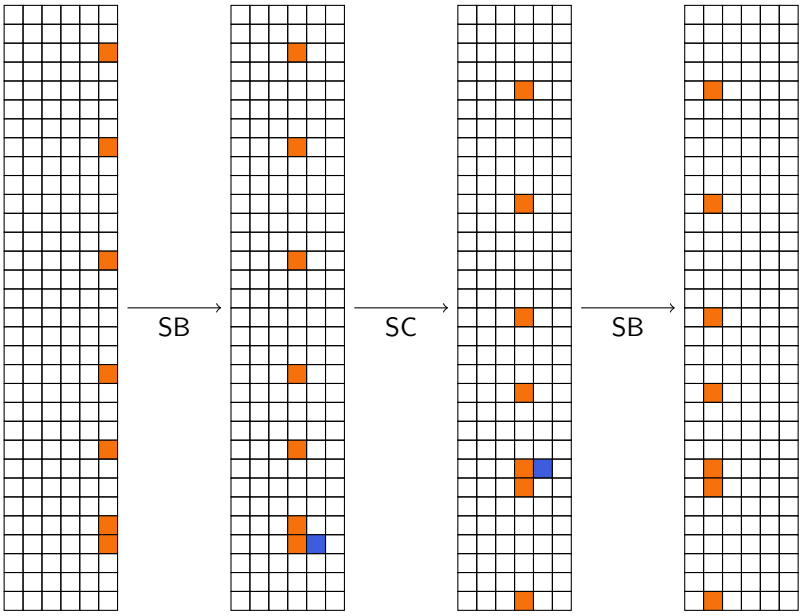

Figure 19: Characteristic B_3 with probability $2^{-44.49}$ based on two quasidifferential trails with correlations $2^{-44.91}$ and $2^{-46.49}$ for nonzero bit $u_{28,4}$ in $\mathbb{F}_2$ of the intermediate mask after the first SubBox application.

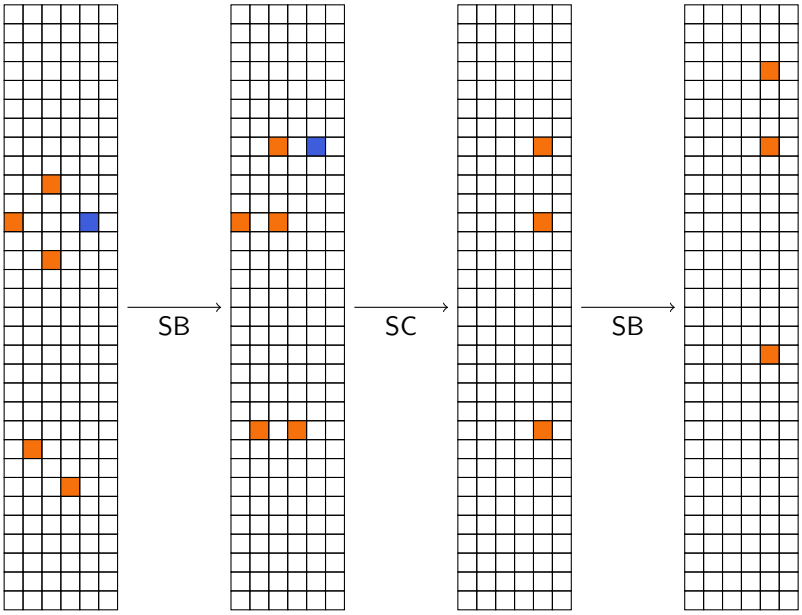

Figure 20: Characteristic B_4 with probability $2^{-33.25}$ based on two quasidifferential trails with correlations $2^{-34.25}$ and $2^{-34.25}$ for nonzero bit $u_{7,4}$ in $\mathbb{F}_2$ in the intermediate mask after the first SubBox application.

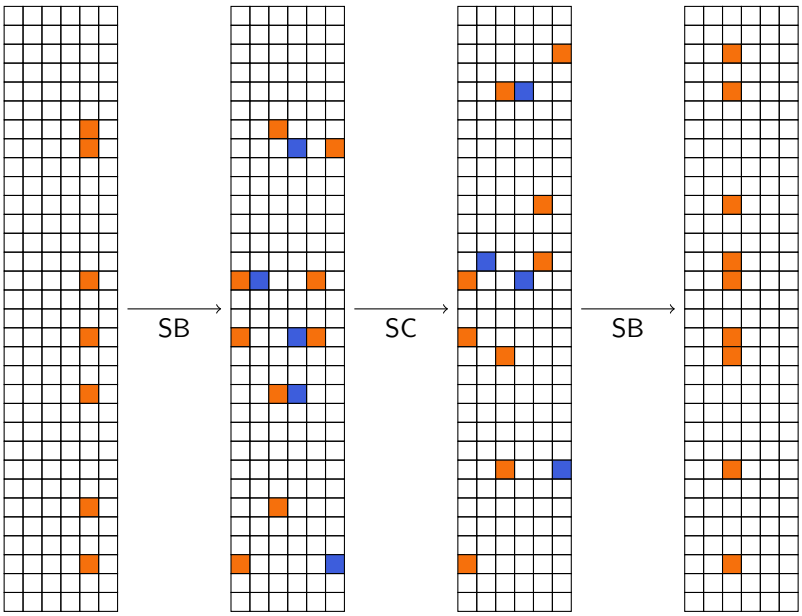

Figure 21: Characteristic B_5 with probability $2^{-54.79}$ based on 32 quasidifferential trails with correlations $2^{-55.3}$, $-2^{-56.9}$, $2^{-57.9}$, $-2^{-57.9}$, $2^{-58.5}$, $2^{-59.5}$, $-2^{-59.5}$, $-2^{-60.1}$, $-2^{-60.5}$, $2^{-61.1}$, $-2^{-61.1}$, $2^{-62.1}$, $2^{-62.7}$, $-2^{-62.7}$, $-2^{-63.7}$, $2^{-65.2}$ (each twice) for nonzero bit $(u_{7,3}, u_{14,1}, u_{17,3}, u_{20,3}, u_{29,5})$ in $\mathbb{F}_2^5$ of the intermediate mask after the first SubBox application.

D Alternatives for characteristics with probability zero

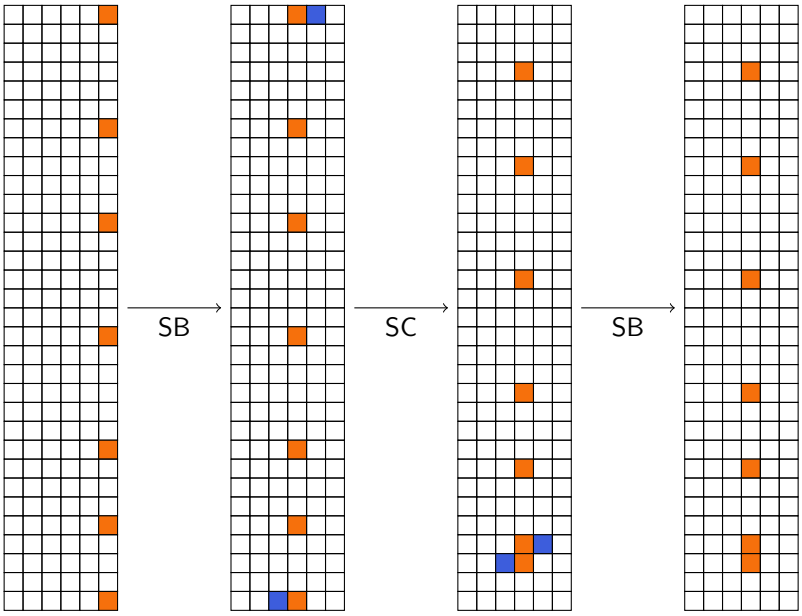

Figure 22: Characteristic S_1 with probability 2^{-49}, with only the quasidifferential trail with all-zero masks.

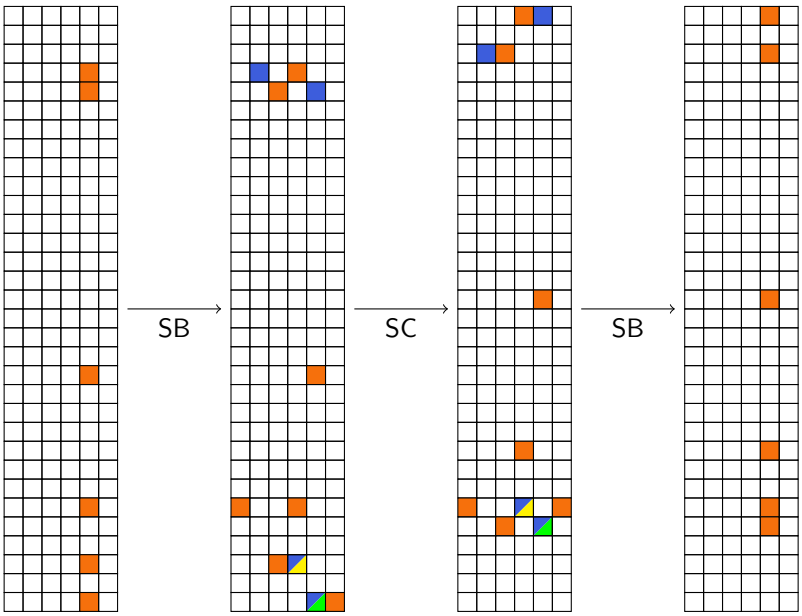

Figure 23: Characteristic S_2 with probability $2^{-47.83}$ based on four quasidifferential trails with correlations $2^{-47.83}$,$2^{-47.83}$,$2^{-48.83}$ and $-2^{-48.83}$ for nonzero bits $(u_{29,3}, u_{31,4})$ in $\mathbb{F}_2^2$ of the intermediate mask after the first SubBox application.

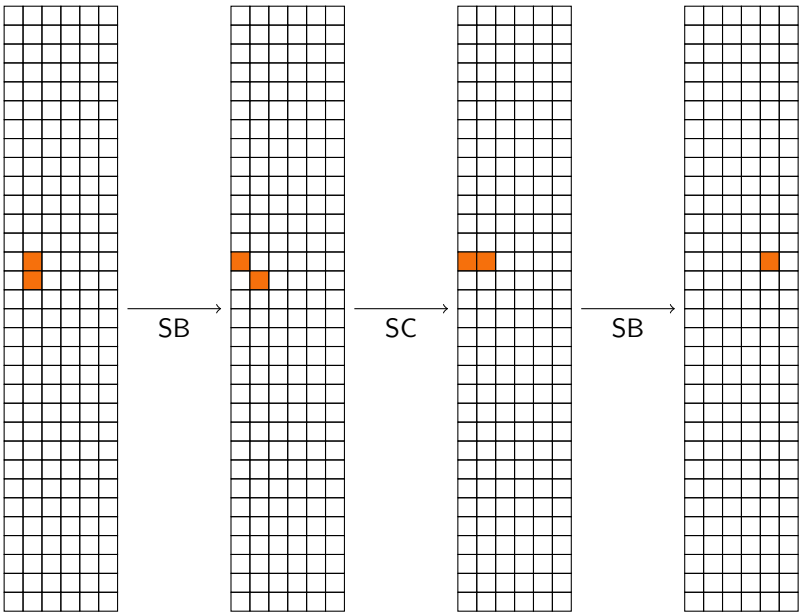

Figure 24: Characteristic E_0 with probability $2^{-11.415}$