

Improved Differential and Linear Trail Bounds for ASCON

Solane El Hirsch, Silvia Mella, Alireza Mehrdad and Joan Daemen

Radboud University, Nijmegen, The Netherlands

solane.elhirsch@ru.nl, silvia.mella@ru.nl, alireza.mehrdad@ru.nl, joan@cs.ru.nl

Abstract. ASCON is a family of cryptographic primitives for authenticated encryption and hashing introduced in 2015. It is selected as one of the ten finalists in the NIST Lightweight Cryptography competition. Since its introduction, ASCON has been extensively cryptanalyzed, and the results of these analyses can indicate the good resistance of this family of cryptographic primitives against known attacks, like differential and linear cryptanalysis.

Proving upper bounds for the differential probability of differential trails and for the squared correlation of linear trails is a standard requirement to evaluate the security of cryptographic primitives. It can be done analytically for some primitives like AES. For other primitives, computer assistance is required to prove strong upper bounds for differential and linear trails. Computer-aided tools can be classified into two categories: tools based on general-purpose solvers and dedicated tools. General-purpose solvers such as SAT and MILP are widely used to prove these bounds, however they seem to have lower capabilities and thus yield less powerful bounds compared to dedicated tools.

In this work, we present a dedicated tool for trail search in ASCON. We arrange 2-round trails in a tree and traverse this tree in an efficient way using a number of new techniques we introduce. Then we extend these trails to more rounds, where we also use the tree traversal technique to do it efficiently. This allows us to scan much larger spaces of trails faster than the previous methods using general-purpose solvers. As a result, we prove tight bounds for 3-rounds linear trails, and for both differential and linear trails, we improve the existing upper bounds for other number of rounds. In particular, for the first time, we prove bounds beyond 2^{-128} for 6 rounds and beyond 2^{-256} for 12 rounds of both differential and linear trails.

Keywords: Differential Trail Search · Linear Trail Search · Trail Weight Bounds · ASCON

1 Introduction

ASCON is a family of cryptographic algorithms for authenticated encryption (AE) and hashing [DEMS21a]. It is currently one of the ten finalists in the NIST lightweight cryptography (LWC) competition for lightweight AE [TMC⁺21] and was selected in the final portfolio of the CAESAR competition [com14] as primary choice for lightweight AE [DEMS16]. The AE schemes are based on the duplex construction [BDPV11a], while the hashing functions are based on the sponge construction [BDPV07, BDPV08]. All family members are based on the ASCON permutation, which is also used in ISAP [DEM⁺20], another finalist in the NIST LWC competition.

The ASCON permutation has been extensively cryptanalyzed since its introduction, giving confidence on the security of the schemes based on it. However a thorough effort to prove bounds on the differential probability (DP) and squared correlation (C^2) of its

44 trails was conducted only recently [GPT21, EME22, MR22]. Before that, only bounds
 45 for 3-round trails were proved in [DEMS15] and for more rounds, the authors performed
 46 heuristic searches showing small DP and small C^2 .

47 Proving bounds for trails is an important task in the evaluation of the security of a
 48 permutation. The cost of a differential attack based on a given trail is inversely proportional
 49 to its DP. Similarly, the cost of a linear attack is inversely proportional to the C^2 . Therefore,
 50 the smaller the DP or C^2 is, the higher the cost of the attack is. Bounds on the DP or
 51 C^2 of trails are usually proven by bounding the number of active S-boxes of the trails
 52 or its weight. Roughly speaking, the weight w of a differential trail relates to its DP as
 53 $DP \approx 2^{-w}$. Similarly, the weight of a linear trail relates to its C^2 as $C^2 \approx 2^{-w}$. Therefore,
 54 the higher the number of active S-boxes or the higher the weight is, the more costly the
 55 attack is.

56 For some primitives, bounds can be proved analytically. An example is the AES with
 57 its simple proof that a 4-round differential trail has weight at least 150 [DR20]. For other
 58 primitives they are obtained by computer-aided proofs. In this case, a program scans the
 59 space of all r -round trails satisfying a given requirement. The requirement is usually that
 60 the number of active S-boxes in the trail is below a given threshold, or that the weight of
 61 the trail is below a given threshold. Large state size and weak alignment contribute in
 62 making the search space very large and thus the cost of scanning it very costly. It follows
 63 that the bounds that one can prove are limited by the capability of the tool for scanning
 64 such spaces.

65 Automated tools that are often used to prove bounds on the number of active S-boxes
 66 are based on general-purpose solvers like Boolean satisfiability (SAT) [MP13, EME22],
 67 (mixed) integer linear programming ((M)ILP) [SHW⁺14, BPP⁺17, BJK⁺16, WH19] or
 68 Satisfiability Modulo Theories (SMT) [DEMS15]. Dedicated tools were used to prove
 69 lower bounds on the weight of trails in NOEKEON [DPAR00], KECCAK- p [DV12, MDV17],
 70 XOODOO [DHVV18b], and SUBTERRANEAN [MMGD22]. Such dedicated programs allow
 71 to better exploit the structural properties of the primitive and usually allow to scan
 72 larger spaces, leading to better results than tools based on general-purpose solvers. Before
 73 2022, the best result obtained with tools based on general-purpose solvers that we are
 74 aware of is the work of Mouha and Preneel, who used a SAT-based method to scan
 75 the space of all 3-round characteristics up to weight 26 in the ARX primitive Salsa20,
 76 which implies a weight per round below 9 [MP13]. The dedicated search for Noekeon
 77 in [DPAR00] and KECCAK- f [1600] in [DV12] both reached a weight per round of 12, while
 78 the improvements of [MDV17] allowed to reach a weight per round of 15. The dedicated
 79 search on SUBTERRANEAN reached a weight per round beyond 14 [MMGD22]. In the last
 80 months, better results have been achieved with both solvers-based tools and dedicated
 81 tools. In [EME22] Erlacher et al. reached a weight per round of 17 with their SAT-based
 82 method to scan the space of trails in ASCON. While the most recent improvements to the
 83 dedicated tool for XOODOO allowed to reach a weight of 21 per round [DMA22].

84 Inspired by the previous works on dedicated tools and their results compared to
 85 automated tools based on general-purpose solvers, in this work we introduce a dedicated
 86 tool for ASCON. We present a number of techniques that deeply exploit the properties of
 87 the linear and non-linear layer of ASCON to generate trails very efficiently. Such techniques
 88 allow us to scan larger spaces of trails at a smaller computational cost compared to previous
 89 work, that results in improved bounds. In particular, we reach a weight per round of 21.

90 **Related work.** Exact values for the DP and C^2 of trails over 1 and 2 rounds of ASCON
 91 can be derived by the fact that the S-box has maximum DP of 2^{-2} and maximum C^2 of
 92 2^{-2} , and that the linear layer has branch number $\mathcal{B} = 4$. For more rounds, lower bounds
 93 were proven in [DEMS15] and [EME22]. Both works are based on SAT solvers and prove
 94 bounds on the number of active S-boxes. Directly bounding the probability would require

95 a more expensive model for the SAT solver compared to bounding the number of active
96 S-boxes, which already requires a major computational effort.

97 In [DEMS15], Dobraunig et al. presented an SMT model and used it to prove that a
98 3-round differential trail has a minimum of 15 active S-boxes and a 3-round linear trail
99 has a minimum of 13 active S-boxes. These bounds automatically give bounds of 2^{-30}
100 and 2^{-26} for the DP and the C^2 of 3-round trails, respectively. Bounds for more rounds
101 were proven later in [EME22], where Erlacher et al. presented a SAT model and used it to
102 prove bounds on the number of active S-boxes for 4 and 6 rounds, from which they derived
103 bounds for 8 and 12 rounds. In addition, by using these results and the bound on 1 round,
104 we can derive bounds for 5, 7, 9, 10, and 11 rounds. We summarize such bounds in the
105 second column of Table 1.

106 To overcome the computational limitation of SAT solvers, the authors of [EME22] aim
107 at reducing the search space as much as possible and split it in sub-spaces that can be
108 scanned in parallel. To this end, they introduced a number of techniques similar to those
109 usually used in dedicated tools, like starting from shorter trails with minimum number of
110 active S-boxes, building long trails from short trails in an incremental way, and taking
111 advantage of the translation symmetry of the primitive [DV12].

112 A significant effort has been also performed to find trails with the highest DP or C^2 .
113 Such searches are based on heuristic tools and provide upper bounds. In [DEMS15], the
114 authors used a dedicated guess-and-determine tool (`nldtool`) to find differential trails up
115 to 5 rounds, while a heuristic tool (`lineartrails`) to find linear trails for 4 and 5 rounds
116 was introduced in [DEM15]. In [GPT21], the authors used constrained programming (CP)
117 to find best differential trails for 5 and 6 rounds. The authors in [MR22] presented an
118 MILP-based approach that allowed them to find a new 5-round linear trail with best
119 known C^2 and proved tight bound for differential trails over 3 rounds. We report the best
120 known trails found by these tools in the first column of Table 1.

121 In dedicated tools, bounds on the weight of trails are derived, instead of evaluating
122 the number of active S-boxes. The first dedicated tool for trail search was introduced
123 as early as 2000 for NOEKEON [DPAR00]. It was later improved and refined in [DV12]
124 and [MDV17] for KECCAK- p and then adapted to XOODOO in [DHVV18b] and SUB-
125 TERRANEAN in [MMGD22]. In each of these works, the authors presented a number of
126 techniques specific for the permutation under analysis that deeply exploit the structure
127 of its linear and non-linear layers. However, the approach underlying these works is the
128 same and is generic, so it can in principle be applied to other ciphers. In a few words,
129 the goal of such approach is to reduce as much as possible the search space and define
130 methods to scan it efficiently. To this end, trails are split into classes where the weight
131 of trails in the same class can be easily bounded by generating only one representative
132 trail per class, called *trail core*. By exploiting the symmetry properties of the permutation,
133 trail cores can be further split into classes where each trail core in a class is the translated
134 version of another trail in the class and trail cores in the same class have the same weight.
135 Therefore, only one representative is generated, that is called *canonical* (or *necklace* to use
136 the terminology of [EME22]). Trail cores over multiple rounds are built by first generating
137 the shortest possible trail cores, that are those over 2 rounds, and by extending them one
138 round at the time each time checking if the weight is below the expected limit. In [MDV17]
139 a generic method is introduced to generate such 2-round trail cores efficiently as a tree
140 search.

141 **Our contribution.** In this work we present a dedicated tool for trail search in ASCON,
142 based on the tree-based approach introduced in [MDV17]. To obtain an efficient instantia-
143 tion of the tree-based approach, we introduce a number of techniques that deeply exploit
144 the structure of the linear and non-linear layers in ASCON. We also introduce methods
145 to efficiently extend trails over multiple rounds. We implemented such techniques in a

Table 1: Previous and new bounds for the differential probability (DP) of differential trails and squared correlation (C^2) of linear trails in ASCON. R denotes the number of rounds; min #S denotes the minimum number of active S-boxes.

(a) Differential trails

R	best known probability			previous lower bound			new bound DP
	DP	method	reference	DP	method	reference	
1	2^{-2}	DDT		2^{-2}	DDT		
2	2^{-8}	DDT+ \mathcal{B}		2^{-8}	DDT+ \mathcal{B}		
3	2^{-40}	nldtool	[DEMS15]	2^{-40}	MILP	[MR22]	
4	2^{-107}	nldtool	[DEMS15]	$\leq 2^{-72}$	SAT+min #S	[EME22]	$\leq 2^{-86}$
5	2^{-190}	CP	[DEMS15, GPT21]	$\leq 2^{-74}$	combine 1R+4R		$\leq 2^{-100}$
6	2^{-305}	CP	[GPT21]	$\leq 2^{-108}$	SAT+min #S	[EME22]	$\leq 2^{-129}$
7				$\leq 2^{-110}$	combine 1R+6R		$\leq 2^{-131}$
8				$\leq 2^{-144}$	SAT+min #S	[EME22]	$\leq 2^{-172}$
9				$\leq 2^{-146}$	combine 1R+8R		$\leq 2^{-186}$
10				$\leq 2^{-180}$	combine 4R+6R		$\leq 2^{-215}$
11				$\leq 2^{-182}$	combine 1R+10R		$\leq 2^{-229}$
12				$\leq 2^{-216}$	SAT+min #S	[EME22]	$\leq 2^{-258}$

(b) Linear trails

R	best known squared correlation			previous lower bound			new bound C^2
	C^2	method	reference	C^2	method	reference	
1	2^{-2}	LAT		2^{-2}	DDT		
2	2^{-8}	LAT+ \mathcal{B}		2^{-8}	DDT+ \mathcal{B}		
3	2^{-28}	lineartrails	[DEM15]	$\leq 2^{-26}$	SMT+min #S	[DEMS15]	2^{-28}
4	2^{-98}	lineartrails	[DEM15]	$\leq 2^{-72}$	SAT+min #S	[EME22]	$\leq 2^{-88}$
5	2^{-184}	MILP	[MR22]	$\leq 2^{-74}$	combine 1R+4R		$\leq 2^{-96}$
6				$\leq 2^{-108}$	SAT+min #S	[EME22]	$\leq 2^{-132}$
7				$\leq 2^{-110}$	combine 1R+6R		$\leq 2^{-134}$
8				$\leq 2^{-144}$	SAT+min #S	[EME22]	$\leq 2^{-176}$
9				$\leq 2^{-146}$	combine 1R+8R		$\leq 2^{-184}$
10				$\leq 2^{-180}$	combine 4R+6R		$\leq 2^{-220}$
11				$\leq 2^{-182}$	combine 1R+10R		$\leq 2^{-228}$
12				$\leq 2^{-216}$	SAT+min #S	[EME22]	$\leq 2^{-264}$

146 dedicated tool, called `AsconTrailTool`, that we used to prove bounds for differential and
 147 linear trails for different number of rounds. Though a comparison of the computational
 148 costs of our method and the method of [EME22] is not straightforward, due to the different
 149 machines employed in the two works, our techniques allowed us to scan a larger space at a
 150 lower cost. The most direct consequence is that we can improve over known bounds. We
 151 report our improved bounds in the third column of Table 1. Notably, for linear trails, we
 152 prove tight bound for 3 rounds, closing the gap between the lower bound and the best
 153 known trail. For 4 rounds, we can prove the bound of 2^{-86} for differential trails in 13
 154 CPU days, and of 2^{-88} for linear trails in 110 CPU days. Our method is more efficient in
 155 comparison to the previous methods where the cost estimation for proving the bound of
 156 2^{-80} is 6688 CPU days in [EME22] and 3898 CPU days in [MR22].

157 Given the aforementioned 4-round trails, proving bounds for 6 rounds required us 6
 158 additional CPU days to prove the bound of 2^{-129} for differential trails and 21 additional
 159 CPU days to prove the bound of 2^{-132} for linear trails. Our method performs better than
 160 the one in [EME22] where the authors indicated that it required 2 additional CPU months
 161 to prove the bound of 2^{-108} . For 12 rounds, we can prove for the first time bounds beyond
 162 2^{-256} . We also prove better bounds for other numbers of rounds, which can be useful
 163 information for designers when they have to choose the number of rounds to use in the
 164 different phases of a given construction.

165 **Organization of the paper.** In Section 2 we first recall some concepts about trails and
 166 trail cores, then we recall the strategy used in previous dedicated tools to prove trail
 167 bounds and the generic tree-based method. Then, in Section 3 we present the specification
 168 of ASCON round function and propagation properties through it. In Section 4, we introduce
 169 the tree-based method applied to ASCON to generate 2-round trail cores and provide new
 170 techniques to traverse the tree in a more efficient way. After that, we explain how we
 171 efficiently perform trail core extension using the techniques introduced in Section 5. Finally,
 172 we present our practical results and improved bounds in Section 6 and in Section 7 we
 173 provide some final remarks.

174 2 Trails and trail search strategy

175 In this section we first recall some concepts related to differential and linear cryptanalysis.
 176 Then we explain the general strategy for performing trail search using the tree-based
 177 approach.

178 2.1 Trails and trail cores

179 We start by defining differential trails and trail cores over iterative cryptographic primitives.
 180 Then, we do the same for linear trails and we introduce a unified notation for both cases.

181 2.1.1 Differentials and differential trails

182 Let x_1 and x_2 be two inputs to a transformation α over \mathbb{F}_2^n , and $y_1 = \alpha(x_1)$ and $y_2 = \alpha(x_2)$
 183 be their corresponding outputs. We say $b = x_1 \oplus x_2$ is an input difference of α and
 184 $a = y_1 \oplus y_2$ is an output difference and we call the pair (b, a) a *differential* over α . The
 185 difference probability (DP) of a differential (b, a) is defined as

$$186 \quad \text{DP}(b, a) = \frac{|\{x \in \mathbb{F}_2^n \mid \alpha(x - b) - \alpha(x) = a\}|}{2^n}.$$

187 When $\text{DP}(b, a) > 0$, we say that a is *compatible* with b through α . The *restriction weight*
 188 of a differential, denoted by w_r , is defined as

$$189 \quad w_r(b, a) = -\log_2 \text{DP}(b, a).$$

190 Let α be an iterative mapping, that consists of the repetition of a number of rounds p_i :
 191 $\alpha = p_r \circ \dots \circ p_2 \circ p_1$. A differential over p_i is called a *round differential*. An *r-round*
 192 *differential trail* over α is a sequence of r round differentials.

193 Let the round function be defined as the composition of a linear layer p_L and a non-
 194 linear layer p_S . We use a redundant representation of trails where we specify the difference
 195 after each layer:

$$196 \quad Q = a^0 \xrightarrow{p_L} b^0 \xrightarrow{p_S} a^1 \xrightarrow{p_L} b^1 \xrightarrow{p_S} a^2 \xrightarrow{p_L} \dots \xrightarrow{p_S} a^r.$$

197 The restriction weight of a trail is the sum of the weight of its round differentials: $w_r(Q) =$
 198 $\sum_{i=1}^r w_r(a^{i-1} \xrightarrow{p_i} a^i)$. Since p_L is linear, the weight of a trail only depends on the weight
 199 over the non-linear layers: $w_r(Q) = \sum_{i=1}^r w_r(b^{i-1} \xrightarrow{p_S} a^i)$. If the non-linear layer p_S has
 200 algebraic degree 2 (as in ASCON), the weight of a differential over p_S only depends on its
 201 input difference b [Dae95]. Hence, the weight of the trail is given by $w_r(Q) = \sum_{i=1}^r w_r(b^{i-1})$.

202 Since the weight of an r -round trail Q is independent of the first and last differences
 203 of the trail, the sequence of differences $(b^0, a^1, \dots, a^{r-1}, b^{r-1})$ – which is Q with the first
 204 and last differences removed – defines a set of r -round trails with the same weight $w_r(Q)$.
 205 On the other hand, for a given a^1 there exist several differences b^0 that are compatible
 206 with a^1 through p_S^{-1} . The minimum weight over all these compatible states b^0 is called
 207 the *minimum reverse weight* of a^1 and it is denoted by $w_{\text{rev}}(a^1)$ [DV12]. It follows that
 208 the sequence $\tilde{Q} = (a^1, \dots, a^{r-1}, b^{r-1})$ defines a set of r -round trails with weight at least
 209 $w_{\text{rev}}(a^1) + \sum_{i=2}^r w_r(b^{i-1})$. \tilde{Q} is called *r -round differential trail core* [DV12].

210 2.1.2 Correlation and linear trails

211 Let α be a transformation over \mathbb{F}_2^n . A linear approximation over α consists of a pair (a, b) of
 212 selection vectors over \mathbb{F}_2^n , called *input mask* and *output mask*, respectively. The *correlation*
 213 C of a linear approximation (a, b) is the correlation between the Boolean functions $a^T \cdot x$
 214 and $b^T \cdot \alpha(x)$:

$$215 C(a, b) = \frac{|\{x \in \mathbb{F}_2^n \mid a^T x + b^T \alpha(x) = 0\}|}{2^{n-1}} - 1.$$

216 The *correlation weight* is denoted by $w_c(a, b)$ and is defined as

$$217 w_c(a, b) = -\log_2 C^2(a, b).$$

218 Similar to a differential trail, an *r -round linear trail* is defined as a sequence of linear
 219 masks. As in [BDPV11b, DHVV18b] we study linear propagation from the output to the
 220 input. To this end, we rephrase the round function so that the trail first encounters p_L and
 221 then p_S of each round (as in the differential case). Notice that such rephrasing does not
 222 affect the trail analysis.

223 A linear trail is represented as

$$224 Q = a^0 \xrightarrow{p_L^\top} b^0 \xrightarrow{p_S^{-1}} a^1 \xrightarrow{p_L^\top} b^1 \xrightarrow{p_S^{-1}} a^2 \xrightarrow{p_L^\top} \dots \xrightarrow{p_S^{-1}} a^r.$$

225 where a_0 is the output mask (after the last round) and a_r is the input mask (before the
 226 first round). A mask a_i at the output of p_L maps to a mask $b_i = p_L^\top(a_i)$ before p_L . If
 227 the linear mapping p_L is seen as the multiplication by a matrix M , then p_L^\top denotes the
 228 linear mapping obtained by the multiplication by M^\top . To denote the propagation from
 229 the output of p_S to its input, we use p_S^{-1} .

230 The correlation weight of a linear trail is the sum of the correlation weights of the
 231 round linear approximations composing the trail. Given that p_L^\top is linear and that, when
 232 p_S has algebraic degree 2, the correlation weight depends only on the value of the output
 233 mask [Dae95], the weight of a linear trail is given by $w_c(Q) = \sum_{i=1}^r w_c(b^{i-1})$

234 Similar to the differential case, an *r -round linear trail core* [DHVV18b] is a sequence
 235 $\tilde{Q} = (a^1, \dots, a^{r-1}, b^{r-1})$ that defines a set of r -round linear trails with weight at least
 236 $w_{\text{rev}}(a^1) + \sum_{i=2}^r w_r(b^{i-1})$.

237 2.1.3 Unified representation of trail cores

238 As done in [BDPV11b] with KECCAK- p and in [DHVV18b] with XOODOO, we use a unified
 239 representation of trails and trail cores. In fact, also in the case of ASCON, there are strong

240 similarities in the study of propagation of differential and linear trails. For differential
 241 trails we consider the propagation of differences from input to output and for linear trails
 242 we consider the propagation of masks from output to input. A trail core is specified by:

$$243 \quad \tilde{Q} = a^1 \xrightarrow{p_L^*} b^1 \xrightarrow{p_S^*} a^2 \xrightarrow{p_L^*} b^2 \xrightarrow{p_S^*} a^3 \xrightarrow{p_L^*} \dots \xrightarrow{p_S^*} b^{r-1}.$$

244 where

- 245 • $p_L^* = p_L$, and $p_S^* = p_S$ for differential trails, and
- 246 • $p_L^* = p_L^\top$, and $p_S^* = p_S^{-1}$ for linear trails.

247 We refer to differences and masks as *state patterns*, or only *states* or *patterns*, when we
 248 generally talk about trails. A pattern a_i represents a difference at the output of p_S in a
 249 differential trail and a mask at the input of p_S in a linear trail. A pattern b_i represents
 250 a difference at the input of p_S in a differential trail and a mask at the output of p_S in a
 251 linear trail. We use the term weight, denoted by w , when we generically refer to w_r and
 252 w_c .

253 2.2 Strategy of the trail search

254 In our trail search, we aim to scan the space of all r -round trails with weight below a certain
 255 threshold T_r , where r is usually a small number like 3,4, or 6. A naive way to generate them
 256 would be to generate all 1-round trails (i.e. round differentials and linear approximations)
 257 with weight below $\lfloor T_r/r \rfloor$ and then extend them to r rounds. The value of T_r that can be
 258 achieved is limited by the quantity of such 1-round trails, which grows exponentially with
 259 the weight, and the cost of extending them. The number of 1-round trails can be reduced
 260 when symmetry properties are taken into account. For instance, in XOODOO it can be
 261 reduced roughly by a factor 128 thanks to the fact that both the linear and non-linear
 262 layers are invariant with respect to translations parallel to the planes [DHVV18b]. While
 263 in KECCAK- f [1600] it can be reduced by a factor 64 thanks to the translation invariance
 264 along the lanes [MDV17]. Even with such reductions, it is shown that this number still
 265 grows exponentially with the weight [MDV17, DHVV18b].

266 However, as demonstrated in [MDV17, DHVV18b], the number of trails with a given
 267 weight per round decreases with the number of rounds. That is, the number of 2-round
 268 trails with weight below $\lfloor 2T_r/r \rfloor$ is smaller than the number of 1-round trails with weight
 269 below $\lfloor T_r/r \rfloor$. Therefore, a more convenient approach for KECCAK- p like primitives consists
 270 in starting from 2-round trails and extend them. This allows to achieve much higher values
 271 of T_r for the same number of rounds r .

272 Actually, to prove bounds, it is not necessary to generate all r -round trails. We can
 273 limit ourselves to r -round trail cores, since the weight of a trail core lower bounds the
 274 weight of all trails in it. Therefore, we can start from 2-round trail cores and extend them.

275 This strategy was used for KECCAK [DV12, MDV17], XOODOO [DHVV18b], and
 276 SUBTERRANEAN [MMGD22] and we will use it also in this work. In fact, also in the case
 277 of ASCON, starting from 2-round trail cores instead of 1-round trails significantly reduces
 278 the number of patterns to extend. The symmetry properties of ASCON allows us to reduce
 279 the number of 1-round trails and the number of 2-round trail cores with weight per round
 280 ($w/\#r$). In particular, the linear and non-linear layers of ASCON are invariant with respect
 281 to translation along the horizontal axis and it allows to reduce them by a factor 64. In
 282 Fig. 1, we depict these reduced numbers with weight per round ($w/\#r$).

283 2.2.1 Generating 2-round trail cores as a tree search

284 A method to generate all 2-round trail cores with weight below a given threshold T_2
 285 was introduced in [DPAR00], applied to KECCAK- p in [DV12], and improved and refined

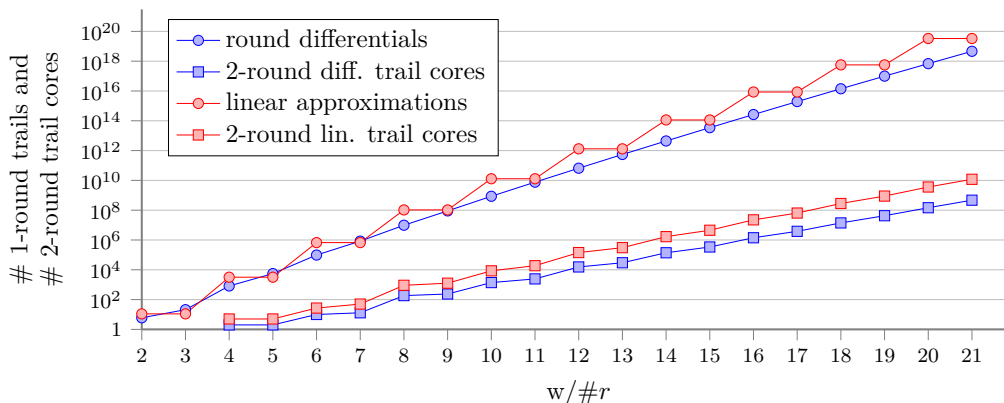


Figure 1: Number of 1-round trails and 2-round trail cores with weight per round ($w/\#r$), divided by 64.

286 in [MDV17]. Later, similar method was applied to XOODOO in [DHVV18b] and also
 287 SUBTERRANEAN in [MMGD22].

288 We now recall the main idea at the basis of the refined method of [MDV17], which
 289 consists in seeing all 2-round trail cores as nodes of a tree that is properly traversed to get
 290 only those nodes with weight below T_2 . In Section 4, we will explain how to instantiate it
 291 for ASCON to perform an efficient search.

292 A 2-round trail core is a pair (a, b) with weight $w_{\text{rev}}(a) + w(b)$. To build them we
 293 have two choices: either we build a and then compute $b = p_L^*(a)$ or we build b and we
 294 compute $a = p_L^{*-1}$. Each node of the tree is encoded as an ordered list of *units*, called
 295 *unit-list*. A unit is a set of *active* bits at a (if we are building a) or at b (if we are building
 296 b), where a bit is called active if it equals one, otherwise it is called *passive*. For instance,
 297 in KECCAK- p and XOODOO a type of unit is the *orbital*, which is a pair of active bits in
 298 the same column at a [MDV17, DHVV18b], while in SUBTERRANEAN a unit is a single
 299 active bit at a [MMGD22].

300 The choice of building first a or b , the definition of units and their order relation
 301 influence the efficiency of the 2-round trail core generation. Therefore, it requires a
 302 good understanding of the linear and non-linear layers of the round function and their
 303 propagation properties.

304 **Traversing the tree.** The tree traversal is performed in a depth-first fashion, where a
 305 program iteratively calls the function `next()` (Algorithm 1) to generate the next valid
 306 node, as in [MMGD22]. The traversal starts by calling `next()` on an empty unit-list, and
 307 ends when it results again in the empty unit list.

308 The function `next()` traverses the tree with three possible moves: `toFirstChild()`,
 309 `toSibling()` and `toParent()`. If the node is an empty unit-list, then it adds the smallest
 310 possible unit. The function `toFirstChild()` returns false if adding a new unit is not
 311 possible. Otherwise it returns true. Then *additional conditions* are checked to see if we can
 312 prune the tree. If the `toFirstChild()` function returns false or the additional conditions
 313 are not satisfied, the routine will look for the next valid node in the tree by generating a
 314 sibling for the current node using the function `toSibling()`. The function `toSibling()`
 315 iterates the value of the last unit of the unit-list. If a sibling is found then the additional
 316 conditions are checked. If there are no valid siblings, the algorithm calls the function
 317 `toParent()` to remove the last unit from the unit-list and look for a valid sibling of the
 318 parent node in a recursive way.

Algorithm 1 `next()` function [MMGD22]

```

if (toFirstChild() == true) then
  if (additional conditions are satisfied) then
    return true;
do
  while (toSibling() == true) do
    if (additional conditions are satisfied) then
      return true;
while (toParent() == true)
return false;

```

319 **Pruning the tree.** To efficiently traverse the tree, at each move we check whether the
320 node satisfies some additional conditions or not. To this end, we make use of two tools:
321 *canonicity* and *score*, whose definition fully depends on the specification of the linear and
322 non-linear layers.

- 323 • **Canonicity:** Without considering round constant and key addition, the round
324 function of many cryptographic primitives exhibits translation symmetry. This
325 symmetry allows to divide the state space into equivalence classes where all patterns
326 in a class have the same properties and weight. Therefore, we aim to generate only
327 one pattern per equivalence class, that is called *canonical*.
- 328 • **Score:** The score of a node is defined as a lower bound on the weight of a node and
329 all its descendants. This tool allows us to prune entire sub-trees as soon as we reach
330 a node whose score is higher than T_2 . It should be tight enough to allow efficient
331 pruning, but also efficiently computable.

332 2.2.2 Trail core extension

333 After generating all 2-round trail cores with weight below T_2 , we need to extend them to
334 generate trail cores over more rounds. Extension is done incrementally one round at the
335 time. Namely, we first extend the 2-round trail cores by one round to generate 3-round
336 trail cores with weight below a given T_3 . Then we extend the obtained 3-round trail cores
337 by one round to generate 4-round trail cores with weight below a given T_4 and so on.

338 Given an r -round trail core $\tilde{Q} = (a^1, b^1, \dots, b^{r-1})$, one can extend it to $(r + 1)$ rounds
339 in both forward and backward directions. In the term forward extension, forward means
340 through p_S^* , so through p_S for differential trails and through p_S^{-1} for linear trails. Backward
341 means through p_S^{*-1} , so through p_S^{-1} for differential trails and through p_S for linear trails.

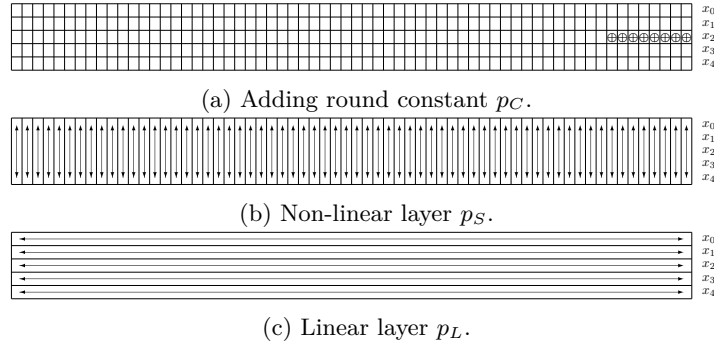
342 In forward extension, we generate all patterns a^r that are compatible with b^{r-1} through
343 p_S^* , compute $b^r = p_L^*(a^r)$ and finally append (a^r, b^r) to the end of \tilde{Q} . The weight of the
344 obtained cores is $w(\tilde{Q}) + w(b^r)$.

345 In backward extension, we generate all patterns b^0 compatible with a^1 through p_S^{*-1} ,
346 then compute the corresponding $a^0 = p_L^{*-1}(b^0)$, and prepend them to \tilde{Q} . The weight of these
347 3-round trail cores is obtained by subtracting $w_{\text{rev}}(a^1)$ and then adding $w_{\text{rev}}(a^0) + w(b^0)$.

348 By repeating the aforementioned process, one can extend a trail core over multiple
349 rounds in any direction.

350 3 The Ascon permutation

351 ASCON family includes the authenticated encryption schemes ASCON-128 and ASCON-
352 128A [DEMS21b], the hash functions ASCON-HASH and ASCON-HASHA and the extendable

Figure 2: ASCON's round function p .

353 output functions (XOF) ASCON-XOF and ASCON-XOFA. The AE schemes are based on
 354 the duplex construction [BDPV11a], while the hashing and XOF functions are based on
 355 the sponge construction [BDPV07, BDPV08]. All family members are based on the ASCON
 356 permutation, which is also used in ISAP [DEM⁺20], another finalist of the NIST LWC
 357 competition.

358 3.1 Ascon round specification

359 The ASCON permutation operates on a state of 320 bits arranged in five 64-bit rows
 360 x_0, \dots, x_4 . The number of rounds is a tunable parameter. It is 12 in the initialization and
 361 finalization phase of all ASCON schemes, while it changes for the data processing phase.
 362 It is 6 for ASCON-128, 8 for ASCON-128A, ASCON-HASHA, and ASCON-XOFA, and 12 for
 363 ASCON-HASH and ASCON-XOF.

364 The round function of ASCON is denoted by p and consists of three steps: $p = p_L \circ p_S \circ p_C$.
 365 The function p_C , that can be seen in Fig. 2a, adds a round constant to row x_2 of the state.
 366 The non-linear layer p_S applies 64 parallel 5-bit S-boxes, denoted \mathcal{S} , to the columns of the
 367 state, as in Fig. 2b. The non-linear part of the S-box \mathcal{S} is based on the χ shift-invariant
 368 mapping [Dae95]. We denote χ applied to an n -bit circle of bits as χ_n , so the S-box in
 369 KECCAK- p is χ_5 [BDPV11b]. We hence can describe \mathcal{S} as χ_5 preceded and followed by
 370 two linear mappings, each consisting of 3 bitwise additions. We depict it in Fig. 3.


371 Finally, p_L applies a linear function to each row independently as in Fig. 2c and is
 372 defined as follows:

$$\begin{aligned}
 x_0 &\leftarrow x_0 \oplus (x_0 \ggg 19) \oplus (x_0 \ggg 28) \\
 x_1 &\leftarrow x_1 \oplus (x_1 \ggg 61) \oplus (x_1 \ggg 39) \\
 x_2 &\leftarrow x_2 \oplus (x_2 \ggg 1) \oplus (x_2 \ggg 6) \\
 x_3 &\leftarrow x_3 \oplus (x_3 \ggg 10) \oplus (x_3 \ggg 17) \\
 x_4 &\leftarrow x_4 \oplus (x_4 \ggg 7) \oplus (x_4 \ggg 41)
 \end{aligned} \tag{1}$$

374 Clearly, in p_L there is no inter-row mixing and this is compensated by the linear
 375 mappings in p_S .

376 3.2 Propagation properties through the round

377 Since the S-box \mathcal{S} is based on the χ_5 mapping also used in KECCAK- p , it inherits some
 378 interesting properties from it that were discussed in [Dae95] and that we summarize here.

Figure 3: ASCON's S-box \mathcal{S} .

379 **Difference propagation properties.** Since p_S has algebraic degree 2, given a difference
 380 b at the input of p_S , the space of compatible differences a at the output of p_S form a
 381 linear affine space $\mathcal{A}(b)$ with $2^{w_r(b)}$ elements [Dae95]. We can compute offset and basis
 382 for such space starting from offset and basis over χ_5 , that are reported in [BDPV11b]. In
 383 particular, for a given difference b at the input of \mathcal{S} , we map it at the input of χ_5 through
 384 the first linear layer of bitwise additions, we take the offset and basis that determine the
 385 affine space at the output of χ_5 , and finally we map them through the second linear layer.
 386 We provide offset and basis vectors for all possible 31 non-zero differences at the input of
 387 \mathcal{S} in Table 7. Among the 31 non-zero differences, 5 have weight 2, 15 have weight 3, and
 388 11 have weight 4. Therefore, the weight of b is at least twice the number of active columns
 389 in b .

390 Difference propagating through the inverse of p_S is different. For a given difference a
 391 at the output of p_S , the set of compatible differences b at the input of p_S is not an affine
 392 space, but we can exhaustively list them. The list of the differences b compatible with a
 393 is needed to compute $w_{\text{rev}}(a)$ which is required for our trail search. Among the 31 non-zero
 394 differences, 10 have 9 compatible differences, 10 have 10 compatible differences, 6 have 11
 395 compatible differences, and 5 have 12 compatible differences. Moreover, 20 have minimum
 396 reverse weight 2, and 11 have minimum reverse weight 3.

397 **Mask propagation properties.** For a given output mask b , the space of input mask a
 398 with a non-zero correlation with b is a linear affine space with $2^{w_c(b)}$ elements [Dae95].
 399 Again, to build a representation of such space, we rely on the specification of offset and
 400 basis over χ_5 [BDPV11b]. We provide offset and basis vectors for all possible 31 non-zero
 401 masks at the output of \mathcal{S} in Table 8. Among the 31 non-zero masks, 10 have weight 2,
 402 and 21 have weight 4.


403 Given a mask a at the input of p_S , we can list the compatible masks b at the output
 404 of p_S , which do not form an affine space. Among the 31 non-zero masks, 10 have 10
 405 compatible masks, and 20 have 13 compatible masks, and 1 has 16 compatible masks.
 406 Moreover, 30 have minimum reverse weight 2, and 1 has minimum reverse weight 4.

407 Notice that a linear trail has always even weight.

408 As explained in Section 2.1.2, the propagation of masks through the linear layer p_L
 409 is deterministic: an output mask b fully determines the corresponding input mask a by
 410 $b = p_L^T(a)$. The transpose p_L^T has the same shape as p_L itself, the only difference is that
 411 the right shifts become left shifts.

412 4 Generating 2-round trail cores in Ascon as tree-search

413 In this section, we explain how we generate all 2-round trail cores in ASCON, with weight
 414 below a given target T_2 , using the tree-based approach of Section 2.2.1. To this end, we
 415 first define units and their order relation. Then we give a description of the techniques

Figure 4: Propagation through p_L

416 used to traverse the tree and, to do it in an efficient way, we define the score function and
 417 discuss canonicity. After identifying the techniques used in the tree-search in Section 4.1,
 418 we give a more detailed description on the two-level tree search in Section 4.2, and in
 419 Section 4.3 we give a description of an alternative representation of p_L^* .

4.1 Concepts and techniques

420 **Active bits as units.** For the tree-based approach we have to define units and their
 421 ordering and the most important criteria for this choice are the ability to define an efficient
 422 score function and deal with canonicity efficiently. The linear mapping p_L^* does not have
 423 a particular structure like the column parity mixers in XOODOO or KECCAK- p , and the
 424 obvious choice for units would be (coordinates of) active bits. We can choose to have the
 425 units be active bits in a or in b . In other words, we either build the state at a and we
 426 compute $b = p_L^*(a)$, or we build the state at b and we compute $a = p_L^{*-1}(b)$.
 427

428 **Active bits in a as units.** If the units are defined as active bits in a , adding a unit affects
 429 3 bits in b . If some of these bits are active in the parent, this addition cancels them. We
 430 call the effect of active bits in a parent that are not present in the child *cancellation*. The
 431 inverse of the row mapping p_L^* is *dense*: it maps a row with a single active bit to a row
 432 with many active bits. If the units are defined as active bits in b , adding a unit affects
 433 many active bits in a , risking the cancellation of many more active bits. We illustrate this
 434 asymmetry for the mapping p_L on row 0 in Fig. 4. It works similarly for p_L^\top . So with
 435 units defined at a an efficient score is more likely to be easy as there is less opportunity
 436 for cancellation. So we define our units as active bits in a . Note that cancellation only
 437 takes place in b and an active bit in a will be present in all its children.

438 **Score function based on number of active columns.** The non-linear layer p_S operates
 439 in parallel on 5-bit columns. This is similar to XOODOO where the non-linear layer is the
 440 parallel operation of χ_3 on 3-bit columns and KECCAK- p , where it is the parallel operation
 441 of χ_5 on 5-bit rows. χ_3 and χ_5 are instantiations of χ that has the property that adding
 442 an active bit to an input difference does not decrease the weight, and that adding an active
 443 bit at the output does not decrease the minimum reverse weight. This also holds for linear
 444 masks. In p_S this is not the case due to the presence of additional linear mappings in
 445 the S-box. So, adding an active bit to a column in a may decrease its minimum reverse
 446 weight and adding an active bit to a column in b may decrease its weight. Still, each active
 447 column in a contributes at least 2 to its minimum reverse weight and each active column
 448 in b contributes at least 2 to its weight. Moreover, adding active bits to a column in a
 449 or b cannot make it passive. So we can base the score function on the number of active
 450 columns.

451 **Row-index-first lexicographic ordering.** In a all the active columns can be accounted
 452 for in the score, in b only those that cannot become passive due to cancellation when
 453 adding units. This is where the ordering comes in. Units are defined by their coordinates

454 (i, j) and there are two natural orderings, both lexicographic: i -first or j -first. In i -first
 455 the active bits in row $i = 0$ come before those in row 1 etc., in j -first those in column
 456 $j = 0$ come before those in column 1, etc. The i -first ordering works well with p_L^* . This is
 457 because this mapping is the parallel application of 5 linear mappings that operate on the
 458 rows separately. In the i -first ordering units are added row by row, where units are always
 459 added in the row of the last unit or after it. Let us call the i -coordinate of the last unit i' .
 460 Then rows in a with $i < i'$ will be the same for all descendants of a state. As p_L^* operates
 461 on rows separately, this will also be the case for the rows in b with $i < i'$. That means
 462 that we can take as score function two times the number of active columns in a plus the
 463 number of columns in b that are active in the rows with $i < i'$.

464 **Two-level tree: active rows and active bits.** When we look at the children of a node
 465 we see two kinds. Children where a unit is added to a row that already contains active
 466 bits on the one hand and children where a unit is added to a row that does not on the
 467 other. In the former case the last active row of b cannot be taken into account for the
 468 score and in the latter case it can. We address this distinction by defining the units in a
 469 two-level structure. At the top level the units are active rows, where an active row groups
 470 all active bits in the same row. We will call the top level the *row tree* and its unit-lists
 471 *row-lists*. This means that the children of a node in the row tree have the first active rows
 472 in common with their parent, but have one more active row. The consequence is that when
 473 navigating in the row tree, for the score function we can count all active rows at a and at
 474 b . We call this score function the `Score-state()` function. An active row is a unit list
 475 too, where the units are active bits (within a specific row) listed in a so called *bit-list*. The
 476 consequence is that the children of a node in the row tree are also arranged according to a
 477 tree, that we will call a *bit tree*. More exactly, the children of a node in the row tree with
 478 last active row at i' are $4 - i'$ bit trees. For example if $i' = 2$, the children are grouped in
 479 two bit trees: one that groups the states with last row at row index $i = 3$ and one that
 480 groups the ones with last row index $i = 4$. The two-level tree search is detailed more in
 481 Section 4.2.

482 **Score in the bit tree: the case of index 2.** Each bit tree contains $2^{64} - 1$ nodes so it
 483 would be good to also prune these trees using a score function. Clearly, all active columns
 484 of a and the active columns at b due to all active rows but the last can be counted in
 485 this score. However, this does not help in states with a single active row and also not
 486 when these rows have sparse bit-lists. We will now explain that we can also include active
 487 bits from the last active row in b . Let us take a look at row $i = 2$. Adding a unit at
 488 position j affects three bits in b , in positions j , $j - 1$ and $j - 6$, so it affects bits in b
 489 only in the interval $[j - 6 \bmod 64, j]$. Here we adopt the following convention for intervals
 490 where we take into account the circular structure of the rows of the state: $[x, y]$ with
 491 $y \geq x$ is the set of indexes $\{x, x + 1, x + 2, \dots, y\}$ and $[x, y]$ with $y < x$ is the set of indexes
 492 $\{x, x + 1, \dots, 63, 0, 1 \dots, y\}$. Assume we have an active row (bit-list) where the j -coordinate
 493 of the last active bit is j' . The range of j for the last active bit in its children is $[j' + 1, 63]$,
 494 so if $j' > 5$ the range of corresponding affected bits in b is $[j' - 5 \bmod 64, 63]$. In other
 495 words, any bit in b in the interval $[0, j' - 5]$ will be there for all children in the bit tree and
 496 therefore the corresponding active columns can be counted in the score. This becomes
 497 interesting as soon as $j' > 5$.

498 **Score in the bit tree: general case.** The efficiency of this technique depends on the
 499 (circular) distance between the affected bits in b : the smaller the better. In $j = 2$ this
 500 distance is only 6 but for the other rows, these distances are much larger. For example
 501 for $j = 0$, the bit positions are 0, 19, 28 and the shortest interval that encloses all three
 502 is $[0, 28]$. We will call the length of this interval the *span*. For $j = 1$, the bit positions




Figure 5: The score of a column difference with the first two stable bits set to $(1, 1)$ is 3.

503 0, 39, 61 can be enclosed in an interval of length 25: $[39, 0]$. We can address this problem by
 504 adopting an alternative representation of the row that is used to compute the score in the
 505 so called `Score-row()` function. A more detailed explanation on the new representation is
 506 given in Section 4.3.

507 **Refining the score of b .** Computing the score based on twice the number of active
 508 columns in b is sub-optimal. In fact, while we are working on row i , all active bits at rows
 509 i' with $i' < i$ are stable and thus we can consider their contribution to the weight. In
 510 particular, for a given active column, only bits in rows i' with $i' \geq i$ can be added and this
 511 may potentially decrease the weight (though not below 2), but it may not. We define a
 512 lower bound on the weight of each active column, that we call *score* of the column, as the
 513 minimum among the weight of the column and the weight of all possible columns that can
 514 be obtained by adding bits in $i' \geq i$. Then, the score of a state is the sum over the score
 515 of all columns.

516 We illustrate an example in Fig. 5, with column differences and restriction weight. Let
 517 the first two bits of column c in Fig. 5 be set to $(1, 1)$. These bits are stable and we denote
 518 them in black, while we denote in red the three bits that can become 1 later in the search.
 519 On the right of c we list the six possible column values that we can obtain by adding bits
 520 to c in row 2, 3 or 4. The restriction weight of each column is reported below the column
 521 and we can see that the minimum weight among them is 3. So, we can define the score of
 522 c to be 3. If there are several active columns whose score is higher than 2, then the score
 523 of b will grow more quickly and pruning comes earlier.

524 **Pruning the tree using canonicity.** Clearly, both p_L and p_S are shift-invariant with
 525 respect to horizontal shifts (along the j -axis). A state that is the *smallest* in its class of
 526 states that are equivalent modulo horizontal shift is called *canonical*. The natural order
 527 to determine which state is smallest is a lexicographical ordering on the row-list: state
 528 X is smaller than state Y if the first row in its row-list is smaller than the first row in
 529 the row-list of Y . If they have equal first rows, we compare the 2nd row and so on. The
 530 order of rows is similarly defined using lexicographic ordering of their bit-lists, where we
 531 compare j -coordinates of active bits starting from the first one.

532 It was proven in [MDV17] that with such an order relation, the children of a non-
 533 canonical node are not canonical. This implies that whenever a non-canonical node is
 534 encountered, the full subtree can be pruned. For an active row we can define its *period*:




Figure 6: The row x_0 in the left states is canonical with different period. In the top figure, since the period of x_0 equals the row length, translation results in a non-canonical state. In the bottom figure since the period of x_0 is smaller than the row length, translation can result in a canonical state.

535 it is the smallest offset such that a shift of the row over that offset leaves it invariant.
 536 The period must be a divisor of 64 (the row length) and the vast majority of row values
 537 has period 64. If the first active row of a canonical state has period 64, all its children
 538 are canonical. This means that in that case we do not have to check for canonicity in
 539 subsequent active rows. Otherwise, we have to check canonicity by shifting any newly
 540 added active row over all multiples of the period and comparing. Examples are given in
 541 Fig. 6.

542 In general, only if the partial state consisting of the stable rows has period smaller
 543 than the row length, these checks must be done when adding an active bit.

544 4.2 Two-level tree

545 We represent a 2-round trail core (a, b) by the positions of its active bits in a . An active
 546 bit is determined by its coordinates (i, j) in the state with i the row coordinate and j the
 547 column coordinate and $0 \leq i < 5$ and $0 \leq j < 64$.

548 The bit-list of an active row is of the following form

$$549 \quad a_i = [(i, j_1), (i, j_2), \dots, (i, j_\ell)], \quad (2)$$

550 with $j_k < j_{k+1} \forall k \in \{1, \dots, \ell - 1\}$. We have that $a_{i, j_k} = 1$ if and only if $k \in \{1, \dots, \ell\}$.

551 At state-level, the row-list of a state a is a list of the form

$$552 \quad a = [a_{i_1}, a_{i_2}, \dots, a_{i_{r-1}}, a_{i_r}] \quad (3)$$

553 with $i_s < i_{s+1} \forall s \in \{1, \dots, r - 1\}$. We have that $a_{i, j} = 1$ if and only if $(i, j) \in \bigcup_s a_{i_s}$. The
 554 *smallest* value that an active row a_i can assume is $[(i, 0)]$.

555 We use two sets of functions to walk through the tree. One is the set of functions that
 556 operate on the bit-list of the last active row. The other is the set of functions that operate
 557 on the row-list.

558 We start by describing the former, where we assume the bit-list of the last active row
 559 is as in Eq. (2).

560 **toFirstChildRow()** If $1 + j_\ell < 64$, it adds $(i, 1 + j_\ell)$ to the bit-list and returns true. It
 561 returns false otherwise.

562 **toSiblingRow()** If $1 + j_\ell < 64$, it iterates the last bit in the list, i.e. (i, j_ℓ) becomes
 563 $(i, 1 + j_\ell)$ and returns true. It returns false otherwise.

564 **toParentRow()** It removes the last bit of the list, resulting in $a_i = [(i, j_1), (i, j_2), \dots, (i, j_{\ell-1})]$.
 565 If it leaves the bit-list empty, it returns false and true otherwise.

566 The following functions operate on the row-list, where the row-list of the current node
 567 is as in Eq. (3).

568 **toFirstChildState()** If $1 + i_r < 5$, it adds $a_{1+i_r} = [(1 + i_r, 0)]$ to the row-list and returns
 569 true. It returns false otherwise.

570 **ToSiblingState()** It calls **nextRow()** on the last active row and if that returns true, it
 571 returns true. Otherwise, it checks whether the last active row is the bottom row,
 572 i.e., $i_r = 4$. If so, it returns false. If not, it moves the last active row one row index
 573 down, i.e. $i_r = 1 + i_r$, and there takes the smallest active row value $a_{i_r} = [(i_r, 0)]$
 574 and returns true.

575 **toParentState()** It first removes the last active row from the list, resulting in $a =$
 576 $[a_{i_1}, a_{i_2}, \dots, a_{i_{r-1}}]$. If this leaves the row-list empty it returns false and the search is
 577 over. Otherwise, it returns true.

578 The complete search works as follows. The tree traversal starts by calling **nextState()**
 579 on a state with a single active row set with a single active bit at position 0 and ends when
 580 **nextState()** returns false, that is when the row-list is empty. Its behavior is similar to
 581 that of the function **next()**. To prune the row tree the procedure calls **Score-state()** on
 582 the current canonical state.

583 The function **nextRow()** in Algorithm 3 is called by **ToSiblingState()** to iterate
 584 the last active row through a bit tree. It starts by checking **Score-row()** and if it is
 585 below the budget then it calls **toFirstChildRow()**. Here, a canonicity check is done on
 586 the whole state to only return canonical states. If there is no valid child either because
 587 **Score-row()** is above the budget or a canonical child has not been found, the procedure
 588 will look for a sibling by calling the function **toSiblingRow()**. Here again, a canonicity
 589 check is performed and if a canonical sibling has been found then the procedure returns
 590 true, otherwise the function **toParentRow()** is called.

591 4.3 The alternative row representation

592 The active bits in a row are indexed by j , and we index them by an alternative coordinate
 593 k that has a relation with j as $k = j \times q \bmod 64$, with q odd. Then, the row component
 594 function of p_L can be reformulated in terms of the new representation and this gives a
 595 mapping that only differs in the shift offsets. For a good choice of q we obtain a mapping
 596 with minimum span that we call *alternative representation*. Minimizing the span requires
 597 a specific factor q per row so, we have alternative representation for each row of ASCON.
 598 For $a_j = a'_{jq}$ and $b_j = b'_{jq}$, the alternative representation is defined as follows:

$$\begin{aligned} p_L : b_j &\leftarrow a_j \oplus a_{j+s} \oplus a_{j+t} \\ b'_{jq} &\leftarrow a'_{jq} \oplus a'_{(j+s)q} \oplus a'_{(j+t)q} \\ b'_{jq} &\leftarrow a'_{jq} \oplus a'_{jq+sq} \oplus a'_{jq+ tq} \\ p'_L : b'_k &\leftarrow a'_k \oplus a'_{k+sq} \oplus a'_{k+tq} \end{aligned}$$

599 Since the alternative representation has the minimum span, more active bits in b are
 600 guaranteed to stay active after adding a unit. The active bits in b that remain active after
 601 adding new units to a are called *stable* bits. In the alternative representation, the bits in b
 602 become stable sooner than in the original representation and more active columns can be
 603 accounted in **Score-row()**. For instance, p_L acts on the first row as $b_j \leftarrow a_j \oplus a_{j+19} \oplus a_{j+28}$.
 604 After multiplying the shift offsets by all odd numbers, we found that $q = 7$ results in the
 605 minimum span. So, the alternative representation of the linear diffusion layer for the first
 606 row is defined as $b'_k \leftarrow a'_k \oplus (a'_k \ggg 5) \oplus (a'_k \ggg 4)$. Fig. 7 provides a comparison between
 607 the original and alternative representation of p_L over row 0 where the number of stable
 608 bits, that are depicted by blue cells, is higher in the case of alternative representation.

Algorithm 2 Functions to navigate through a row tree

```

function NEXTSTATE()
  if (toFirstChildState() == true) then
    if (Score-state() <  $T_2$ ) then
      return true;
  do
    while (ToSiblingState() == true) do
      if (Score-state() <  $T_2$ ) then
        return true;
    while (toParentState() == true)
      return false;
end function

function TOFIRSTCHILDSTATE()
  if ( $i_l = 4$ ) then                                ▷ Last active row index has reached the bottom row
    return false;
   $a \leftarrow a \cup [(1 + i_l, 0)];$                 ▷ Set the last active row to the smallest active row value
  return true
end function

function TOSIBLINGSTATE()
  if (nextRow() == true) then
    return true;
  if ( $i_l = 4$ ) then
    return false;
   $i_l \leftarrow 1 + i_l;$                             ▷ The last active row is moved one row index down
   $a_{i_l} = [(i_l, 0)]$                                ▷ Set the last active row to the smallest active row value
  return true
end function

```

Algorithm 3 Function to navigate through a bit tree

```

function NEXTROW()
  if (Score-row() <  $T_2$ ) then
    if ((toFirstChildRow()) && (is canonical) ) then
      return true;
  do
    while ((toSiblingRow()) && (is canonical) ) do
      return true;
    while (toParentRow() == true)
      return false;
end function

```

609 We denote by p'_L the alternative linear mapping of p_L for each row such that

$$610 \quad p_L = \pi_{q-1} \circ p'_L \circ \pi_q$$

611 where $\pi_q(j) = q \times j \bmod 64$. Fig. 8 illustrates how p'_L in the new representation works for
612 row 0. The list of parameters for the different rows of the alternative representation of
613 p_L with the minimum span are listed in Table 2. The alternative representation of p_L^\top
614 corresponds to the mapping obtained with $-q$.




Figure 7: The grey cells at b and b' represent the span in the original and alternative representation of row x_0 , respectively. The original representation (left figure) results in a lower number of stable bits at b (blue cells) compared to its alternative representation on the right.

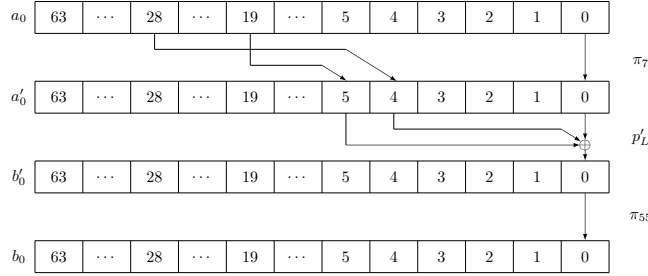


Figure 8: The linear mapping p_L for row 0 can be seen as its alternative representation p'_L surrounded by two multiplication layers, illustrated for bit b_0 .

5 Extension in Ascon

615

616 In this section, we explain how we perform trail core extension in ASCON. We partially
 617 rely on previous works on KECCAK- p [DV12,MDV17] and XOODOO [DHSV18b,DHP⁺20],
 618 given that extension deals with the non-linear layer of ASCON p_S which is based on χ_5 .

619

620 Given a trail core $\tilde{Q} = (a^1, \dots, b^{r-1})$, we recall that forward extension by one round
 621 consists in building all patterns a^r that are compatible with b^{r-1} over p_S^* and compute
 622 $b^r = p_L^*(a^r)$. While backward extension consists in building all patterns b^0 that are
 623 compatible with a^1 over p_S^* and compute $a^0 = p_L^{*-1}(b^0)$.

623

624 The non-linear layer of ASCON can be seen as the parallel application of 64 5-bit S-boxes,
 625 acting on each column independently. Therefore, we can treat extension at column level.
 626 If b^0 and a^1 are compatible over p_S^* , then the j -th column of b^0 is compatible with the
 627 j -th column of a^1 over \mathcal{S} , for any column index $0 \leq j < 64$. To build all states b^0 that

Table 2: List of parameters for the original and alternative representation of the linear diffusion layer of ASCON.

row	original representation			alternative representation			
	offset ₁	offset ₂	span	q	offset ₁	offset ₂	span
0	19	28	28	7	4	5	5
1	61	39	25	41	5	63	6
2	1	6	6	1	1	6	6
3	10	17	17	19	3	62	5
4	7	41	30	47	7	9	9

627 are compatible with a^1 , we first need to identify the active columns in a^1 , namely, the
 628 non-zero columns. Then, for each active column, we build all compatible column values at
 629 b^0 through \mathcal{S} . By combining them, we can finally build all compatible states b^0 .

630 Similarly, we can build all compatible state patterns a^r given b^{r-1} .

631 5.1 Extension as a tree search

632 Extension can be performed as a tree search [MDV17, DHVV18b], where we incrementally
 633 build b^0 or a^r . To this end we need to define units, their order relation, and a score function.
 634 In this case we don't have to deal with canonicity since canonical 2-round trail cores yields
 635 canonical r -round trail cores. A trail core is a sequence of state patterns. Translating each
 636 pattern of the sequence by a fixed offset results in an equivalent trail core with the same
 637 weight. We can define a canonical trail core as the smallest among its translated versions.
 638 We can say that a core $(a^1, b^1, \dots, b^{r-1})$ is smaller than a core $(\bar{a}^1, \bar{b}^1, \dots, \bar{b}^{r-1})$ if a^1 is
 639 smaller than \bar{a}^1 , or if $a^1 = \bar{a}^1$ and b^1 is smaller than \bar{b}^1 , etc. However, we can choose any
 640 intermediate pattern in the sequence instead of a^1 to start the comparison. It is then
 641 natural to start from the $(r-1)$ -round trail core from which the r -round core is generated.
 642 We say that an r -round trail core is canonical if the $(r-1)$ -round trail core from which it
 643 is generated is canonical. It follows that the generation of only canonical 2-round trail
 644 cores, yields to canonical r -round trail cores naturally.

645 Differently from the tree search for the generation of 2-round trail cores where a unit
 646 was an active bit, here units are determined by the compatible column values. At each
 647 move in the tree, we fix the value of an active column of the state. To efficiently traverse
 648 the tree we need a score function that lower bounds the weight of the $(r+1)$ -round trail
 649 cores obtained.

650 In forward extension this translates into lower bounding $w(b^r)$ while we are building
 651 a^r . The addition of a unit at a^r can cancel some bits at b^r because of the action of p_L^* . To
 652 define a good score function, we consider the stable bits at b^r , that are active bits that
 653 cannot be cancelled with the addition of any new unit. We represent stable bits by a
 654 stability mask \mathcal{M} , that is a state where a bit is 1 to indicate that the bit in that position
 655 is stable and 0 otherwise. Then $b^r \wedge \mathcal{M}$ gives the stable bits of b^r , and also the column of
 656 b^r that will be active in all its descendants. We can define the score as twice the number
 657 of active columns in $b^r \wedge \mathcal{M}$.

658 In backward extension we have to lower bound $w_{\text{rev}}(a^0) + w(b^0)$ while we are building
 659 b^0 . While the addition of a unit at b^0 cannot turn active bits into passive, adding a
 660 unit at b^0 can potentially cancel many bits at a^0 , since the inverse of p_L^* is dense. In
 661 KECCAK- p [DV12, MDV17], this problem was overcome by not considering the contribution
 662 of a^0 and by bounding $w_{\text{rev}}(a^0) + w(b^0)$ with a bound on $w(b^0)$ only. However, this is
 663 sub-optimal. In this work, we use stability masks to determine the stable bits of a^0 and
 664 thus consider also its contribution.

665 In general, the goal is to make the number of stable bits in the stability masks grow
 666 as quickly as possible while traversing the tree, so that more columns are counted in the
 667 score and pruning happens as early as possible. To this end, the order relation among the
 668 units must be carefully defined.

669 5.2 Forward Extension

670 For forward extension, we follow the approach used in [DHVV18b] for Xoodoo, that is
 671 the following. All patterns a^r that are compatible with b^{r-1} over p_S^* form an affine space
 672 $\mathcal{A}(b^{r-1})$ with $2^{w(b^{r-1})}$ elements. We represent such space through an offset and a basis.
 673 Each column at b^{r-1} defines an offset and basis for the space of compatible columns over
 674 \mathcal{S} , according to Table 7 and Table 8. The state offset, that we denote by \mathfrak{o} , is built by
 675 gathering together all the column offsets. It will be zero in all column positions that are

676 passive in b^{r-1} . For each column vector \mathbf{u} specified by each active column j , we build a
 677 state vector \mathbf{v} that is all zero except column j that has value \mathbf{u} . The basis has $w = w(b^{r-1})$
 678 elements that we denote by $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_w\}$. Therefore, $\mathcal{A}(b^{r-1}) = \mathbf{o} + \langle \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_w \rangle$.

679 Of course, brute-force scanning the whole affine space becomes unaffordable when
 680 $w(b^{r-1})$ is large. However, we only need to construct those states a^r such that the weight
 681 of b^r is below a given threshold. For this reason, it is practical to directly consider the
 682 affine space mapped through p_L^* , namely before the next p_S^* . We denote such space by
 683 $\mathcal{B}(b^{r-1}) = p_L^*(\mathcal{A}(b^{r-1})) = \mathbf{o}^* + \langle \mathbf{v}_1^*, \mathbf{v}_2^*, \dots, \mathbf{v}_w^* \rangle$, with $\mathbf{o}^* = p_L^*(\mathbf{o})$ and $\mathbf{v}_k^* = p_L^*(\mathbf{v}_k)$.

684 We scan the space $\mathcal{B}(b^{r-1})$ through a tree-based search as follows. The root of the
 685 tree is the offset \mathbf{o}^* . The units are the indexes of the basis vectors, ordered by the natural
 686 number ordering. A unit-list $\mathcal{K} = \{k_1, \dots, k_m\}$ encodes the element of the affine space
 687 given by $\mathbf{o} + \mathbf{v}_{k_1}^* + \dots + \mathbf{v}_{k_m}^*$. The children of \mathcal{K} are all nodes of the form $\mathcal{K} \cup k_{m+1}$ with
 688 $k_{m+1} \in \{k_m + 1, \dots, k_w\}$.

689 We need to define stability masks so that the number of stable bits increases quickly
 690 with k . A technique to do it consists in triangularizing the basis $\mathcal{V}^* = \{\mathbf{v}_1^*, \mathbf{v}_2^*, \dots, \mathbf{v}_w^*\}$. We
 691 perform triangularization in ASCON as follows. We start with an empty basis \mathcal{T} . We loop
 692 on all possible bit positions considering the lexicographic order relation on coordinates
 693 (i, j) . If a basis vector is found with an active bit in position (i, j) , then such basis vector is
 694 added to \mathcal{T} and removed from \mathcal{V}^* . The same vector is also added to all remaining vectors
 695 in \mathcal{V}^* that have bit (i, j) active, to make it passive. After triangularization, we obtain
 696 a new representation of $\mathcal{B}(b^{r-1})$ as $\mathbf{o}^* + \langle t_1, t_2, \dots, t_w \rangle$. If the first active bit in t_k is in
 697 position (i_k, j_k) , then, by construction, all bits in position $(i, j) \leq (i_k, j_k)$ are passive in
 698 all vectors t_{k+1}, \dots, t_w . We call (i_k, j_k) the pivot position of vector t_k . For each k , we
 699 define the stability mask \mathcal{M}_k as a state that is 1 in the pivot position and in all positions
 700 smaller than the pivot (i.e. in all $(i, j) \leq (i_k, j_k)$) and 0 otherwise. In addition we consider
 701 the position of the stable bits in the offset as $\mathcal{O} = \bigwedge_{i=1}^w \bar{t}_i$. We add them to each stability
 702 mask: $\mathcal{M}_k = \mathcal{M}_k \vee \mathcal{O}$.

703 If the last unit in the list of a node b^r is k , then all bits in $b^r \wedge \mathcal{M}_k$ will be active
 704 in all descendants of b^r . Therefore, all active columns of $b^r \wedge \mathcal{M}_k$ will be active in all
 705 descendants of b^r and each will contribute at least 2 to the weight. We define the score as
 706 twice the number of active columns of $b^r \wedge \mathcal{M}_k$.

707 5.3 Backward Extension

708 Given a^1 , the patterns b^0 that are compatible with a^1 over p_S^* do not form an affine space,
 709 so we shall use a different approach than the one for forward extension.

710 We present two methods to perform backward extension. In the first one, presented
 711 in Section 5.3.1, we follow the method used in [DV12] for KECCAK- p , that builds on the
 712 compatible column values, and we introduce some optimizations. Notice that [MDV17]
 713 presents some optimizations for backward extension in KECCAK- p , that exploit the structure
 714 of the linear step θ , which is a column parity mixer. Such techniques do not apply to
 715 ASCON since its linear layer has a different structure. In the second method, presented
 716 in Section 5.3.2, we build an envelope space that contains the set of compatible patterns,
 717 with the aim of growing the number of active columns in a^0 more quickly. The former
 718 method is more effective when the number of active columns in a^1 is small enough, say
 719 less than 12. The second method is more effective when there are many active columns in
 720 a^1 . In our code we use both of them, considering the number of active columns at hand.

721 5.3.1 Extension using compatible patterns

722 For each active column position j in a^1 , let $\mathcal{B}_j = \{\mathbf{v}_{j,1}, \dots, \mathbf{v}_{j,n(j)}\}$ denote the set of
 723 compatible column patterns at the input of p_S^* . The number of compatible patterns b^0
 724 is given by $\prod_j |\mathcal{B}_j|$. Since $n(j)$ ranges between 9 and 12 for compatible differences and

is 10, 13 or 16 for compatible masks, the number of patterns b^0 grows very quickly with the number of active columns in a^1 and it can be unaffordable to generate all of them. However, we need to generate only those such that $w_{\text{rev}}(a^0) + w(b^0)$ is smaller than a given threshold T . We can do it using a tree-based approach where the nodes of the tree are the patterns b^0 and units and score function are defined as follows.

The root of the tree is the fully passive state. The units are the indexes of the elements of the sets \mathcal{B}_j ordered by the lexicographic order over (j, k) . A unit-list can contain at most one element per set of column patterns for a given index j . At height h in the tree, all the first h active columns are set. Only the leaves of the tree give compatible patterns.

The score function shall bound the quantity $w_{\text{rev}}(a^0) + w(b^0)$ for a node and all its descendants. It is defined as $\text{score}_a + \text{score}_b$ with score_a that bounds $w_{\text{rev}}(a^0)$ and score_b that bounds $w(b^0)$.

We start with the explanation of score_b that we compute as in [DV12,MDV17]. We order the elements of each \mathcal{B}_j by increasing weight so that $w(\mathbf{v}_{j,k}) \leq w(\mathbf{v}_{j,k+1})$ for all k . We denote by w_j the minimum of such weights, that is $w_j = w(\mathbf{v}_{j,1})$. For a node at height h , the first h active columns are set and their value cannot change by the addition of a new unit. Each of the remaining active column will contribute to the weight by at least w_j . Therefore, for a node b^0 we define $\text{score}_b(b^0) = w(b^0) + \sum_{h < j} w_j$.

For KECCAK- p [DV12,MDV17], $\text{score}_a = 2$ since a non-passive state has weight at least 2. This is sub-optimal because it does not take into account the contribution of the active bits at a^0 . In this work, we define score_a based on the stable bits of a^0 in the following way. We map each set \mathcal{B}_j before p_L^* obtaining $\mathcal{A}_j = \{\mathbf{v}_{j,1}^*, \dots, \mathbf{v}_{j,n(j)}^*\}$, where $\mathbf{v}_{j,k}^* = p_L^{*-1}(\mathbf{v}_{j,k}^*)$. At height h , one element of each \mathcal{A}_j with $j \leq h$ has been added to a^0 and any element of \mathcal{A}_j can potentially be added for all $j > h$. The OR of the elements that can still be added gives the set of bits that can be potentially cancelled at a^0 . Its negation gives the stable bits. Therefore, for each h , we define the stability mask

$$\mathcal{M}_h = \overline{\bigvee_{h < j} \left(\bigvee_k \mathbf{v}_{j,k}^* \right)} = \bigwedge_{h < j} \left(\bigwedge_k \overline{\mathbf{v}_{j,k}^*} \right).$$

For a node a^0 at height h , all bits of $a^0 \wedge \mathcal{M}_h$ will be active in all descendants of a^0 . Therefore, all active columns of $a^0 \wedge \mathcal{M}_h$ will be active in all descendants of a^0 and each will contribute at least 2 to the weight. We define score_a as twice the number of active columns of the state $a^0 \wedge \mathcal{M}_h$.

The ordering of the elements in each \mathcal{B}_j by increasing weight implies that the right-siblings of a node have weight (resp. score) greater than or equal to the weight (resp. score) of that node. It follows that when a node is encountered whose score is greater than the given threshold all its descendants and also all its siblings can be pruned.

As an additional optimization, we observe that during the backward extension of a trail core $\tilde{Q}_r = (a^1, \dots, b^{r-1})$, $w_{\text{rev}}(a^1)$ is replaced by $w(b^0)$ which can be larger than $w_{\text{rev}}(a^1)$. If $w(\tilde{Q}_r) < T_r$ for a given T_r , most of the times we want $w(\tilde{Q}_r) - w_{\text{rev}}(a^1) + w(b^0)$ to be still smaller than T_r . So, during the search we perform the additional check $\text{score}_b < T_r - (w(\tilde{Q}_r) - w_{\text{rev}}(a^1))$.

5.3.2 Extension using the envelope space

This method aims at prioritizing the growth of the number of active columns in a^0 , so that $w_{\text{rev}}(a^0)$ grows as quickly as possible.

First, we build a space that contains the set of compatible states b^0 's, that we call *envelope space* and denote by \mathcal{E} . To do this, for each active column at a^1 we define the envelope space of its compatible column patterns as $0 + \langle \mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4 \rangle$, where $\mathbf{e}_i \in \mathbb{F}_2^5$ has a single active bit in position i . The envelope space \mathcal{E} is the union of all these envelope spaces and its dimension is five times the number of active columns in a^1 .

We scan \mathcal{E} in a tree-based fashion as done in Section 5.2, where the root of the tree is the offset (in this case the all zero state) and we iteratively add basis vectors. Since the envelope space is much larger than the actual space of compatible states, we must define a score function that is very efficient and allows to prune the tree as soon as possible. To this end, we try to make the number of stable bits in a^0 to grow as quickly as possible. A way to do it is to consider the envelope space before p_L^* and triangularize its basis.

Let $\mathcal{E} = \langle \mathbf{v}_1, \dots, \mathbf{v}_{5n} \rangle$, where n denotes the number of active columns in a^1 . Since p_L^* is linear, we can transpose the envelope space \mathcal{E} before p_L^* and get $\mathcal{E}^* = \langle \mathbf{v}_1^*, \dots, \mathbf{v}_{5n}^* \rangle$ with $\mathbf{v}_k^* = p_L^{*-1}(\mathbf{v}_k)$. We triangularize the basis of \mathcal{E}^* based on the lexicographic order relation on coordinates (i, j) and we modify the representation of \mathcal{E} accordingly. That is, when we add a vector \mathbf{v}_k^* to a vector \mathbf{v}_ℓ^* in \mathcal{E}^* , we add \mathbf{v}_k to \mathbf{v}_ℓ in \mathcal{E} . We obtain a new representation of \mathcal{E}^* as $\langle t_1^*, \dots, t_{5n}^* \rangle$. By construction, the triangularized basis contains first all basis vectors with active bits in row 0, then those with active bits in row 1, etc.

For each k , we define the stability mask \mathcal{M}_k as a state that is 1 in all positions smaller or equal than the pivot position of t_k^* , and 0 otherwise. We define score_a as twice the number of active columns of $a^0 \wedge \mathcal{M}_k$. Finally, we define score_b as twice the number of active columns in a^1 . In fact, the number of active columns in b^0 is the same of a^1 and each contributes at least 2 to the weight. On the other hand, since we are scanning the envelope space and not only the space of compatible states, we cannot use the weight of b^0 , because in this case the addition of a new unit can potentially decrease it.

6 Practical results and improved bounds for Ascon

In this section, we report on our practical results. The improved bounds are reported in Table 1. To scan the different spaces of trail cores, we follow the different strategies presented in [DV12, MDV17, DHP⁺20, DMA22]. We used parts of KECCAKTOOLS [DHVV13] and XOOTTOOLS [DHVV18a] for some routines for trail extension. All our tests are run on a server equipped with an AMD EPYC 7552 48-Core Processor @2.20GHz. We exploited the multicore architecture to run some of our tests in parallel, but execution times are reported as single core costs in the following. We round up the execution time to the closest integer.

In some cases, we compare our execution time to that reported in [EME22], which uses machines equipped with Intel Xeon E5-2669 and E5-4669 v4 @2.20GHz. Even if the machines are different, and thus execution times are not perfectly comparable, we can observe that our methods allow us to scan larger spaces than what was possible with the solvers-based method of [EME22].

In the following, we denote by \mathcal{D}_r^T the space of all r -round differential trail cores with weight $< T$, i.e. at most $T - 1$. Similarly, we denote by \mathcal{L}_r^T the space of all r -round differential trail cores with weight $< T$.

6.1 Results on 3 rounds: tight bound and all low-weight trails

Since the best known 3-round differential and linear trails have weight 40 [DEMS15] and 28 [DEM15] respectively, we scanned the spaces \mathcal{D}_3^{41} and \mathcal{L}_3^{30} to check whether they are the lightest trails¹. Our experimental results confirmed the results for differential trails in [EME22, MR22] and proved that 28 is the tight bound for linear trails. In fact, we found 2 differential trail cores of weight 40, 1 linear trail core of weight 28, and no trail cores with lower weight. The search took less than 3 minutes for differential trails and less than 4 seconds for linear trails.

¹Notice that to prove that they are the lightest trails, it is sufficient to scan the spaces \mathcal{D}_3^{40} and \mathcal{L}_3^{28} and prove that they are empty. To check how many differential trail cores of weight 40 and linear trail cores of weight 28 there exist, we chose to scan larger spaces.

Table 3: Details on the generation of canonical 3-round differential and linear trail cores below target weight 41 and 30, respectively.

search space	# cores	time	search details		
			step	# cores	time
\mathcal{D}_3^{41}	2	3m	$2w_{\text{rev}}(a_1) + w(b_1) < 40$	284,561	2m
			forward extension	2	4s
			$w(b_1) + 2w(b_2) \leq 40$	15,252	28s
			backward extension	0	2s
\mathcal{L}_3^{30}	1	4s	$2w_{\text{rev}}(a_1) + w(b_1) < 28$	1,935	1s
			forward extension	1	1s
			$w(b_1) + 2w(b_2) \leq 28$	972	1s
			backward extension	0	1s

818 To scan the above spaces, we followed the approach used in [DHVV18b], which is the
819 following. A 3-round trail core has weight $w_{\text{rev}}(a^1) + w(b^1) + w(b^2)$. We split all trail cores
820 in \mathcal{D}_3^{41} (resp. \mathcal{L}_3^{30}) into two sets based on whether $w_{\text{rev}}(a^1) < w(b^2)$ or $w_{\text{rev}}(a^1) \geq w(b^2)$.

- 821 • The former case implies that $2w_{\text{rev}}(a^1) + w(b^1) < 40$ (resp. < 28). Such trail cores
822 can be obtained by generating all 2-round trail cores (a^1, b^1) satisfying this inequality
823 and extending them in the forward direction by one round up to 40 (resp. 28).
- 824 • The latter case implies that $w(b^1) + 2w(b^2) \leq 40$ (resp. ≤ 28). Such trails can be
825 obtained by generating all 2-round trail cores (a^2, b^2) satisfying this inequality and
826 then extending them in the backward direction by one round up to 40 (resp. 28).

827 Detailed execution times are given in Table 3 together with the number of trail cores found
828 in each step of the search.

829 Beyond proving bounds for 3-round trails, we are also interested in the distribution of
830 low-weight 3-round trails in ASCON. To this end, we also scanned the space \mathcal{D}_3^{51} (resp.
831 \mathcal{L}_3^{52}), and counted all 3-round trails contained in such cores with weight below 51 (resp.
832 52). To count trails, we used the code for backward extension to build all patterns b^0
833 compatible with a^1 that satisfy $w(b^0) + w(b^1) + w(b^2) < 51$ (resp. < 52) and we count
834 each of them $w(b_2)$ times. Results are depicted in Fig. 9. We can notice that, per given
835 (even) weight ≥ 40 , the ratio between the number of linear trails and differential trails
836 ranges between 9.7 (for weight 46) and 66.5 (for weight 44). This is due to the fact that
837 the LAT of the ASCON S-box is more dense than its DDT.

838 6.2 Results on 4 rounds: improved (non-tight) bounds

839 The best known 4-round differential and linear trails in ASCON have weight 107 and
840 98 respectively [DEMS15, DEM15], while the previously proved lower bound is 72 for
841 both [EME22].

842 With our techniques, we scanned the spaces \mathcal{D}_4^{86} and \mathcal{L}_4^{88} . We found that both spaces
843 are empty, which implies that any 4-round differential trail has weight at least 86 and any
844 4-round linear trail has weight at least 88. This improves over known results, even if the
845 new bounds are still not tight.

846 Our search took around 13 days for differential trails and around 110 days for linear
847 trails. While in [EME22] the authors report a cost of 600 days each for differential and
848 linear trails to prove a bound of 72. Moreover, the authors in [EME22] estimate a cost of




Figure 9: Number of all canonical 3-round trails per weight.

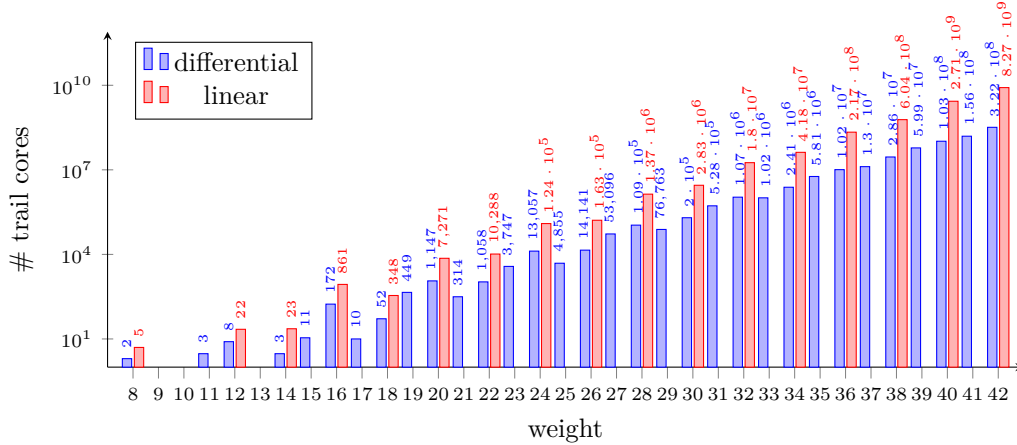


Figure 10: Number of all canonical 2-round trail cores per weight.

849 6688 days to prove a bound of 80 whereas in [MR22], they estimate 3898 days to prove
 850 this bound. Therefore, with our method we could reach higher bounds with significantly
 851 less computational cost.

852 To scan the above spaces, we followed [DHVV18b]. Any 4-round differential (resp.
 853 linear) trail core with weight $w_{\text{rev}}(a^1) + w(b^1) + w(b^2) + w(b^3) < 86$ (resp. < 88) has
 854 $w_{\text{rev}}(a^1) + w(b^1) < 43$ (resp. < 44) or $w(b^2) + w(b^3) < 43$ (resp. < 44). Otherwise, their
 855 sum would be at least 86 (resp. 88). We could thus generate all trail cores in \mathcal{D}_4^{86} (resp.
 856 \mathcal{L}_4^{88}) by generating all 2-round trail cores in \mathcal{D}_2^{43} (resp. \mathcal{L}_2^{44}) and extending them to 4
 857 rounds below 86 (resp. 88). To perform extension to 4 rounds, we first extended to 3
 858 rounds below 84 (resp. 86), since we know that the remaining round has weight at least 2.

859 Details on the number of trail cores found in each step of the search and the execution
 860 times are reported in Table 4. In Fig. 10, we report the number of all 2-round trail cores
 861 per given weight. Again, we can observe that (for even weights) the number of 2-round
 862 linear trail cores found is significantly higher than the number of 2-round differential
 863 trail cores. This difference of course reflects on the costs for extension.

Table 4: Details on the generation of canonical 4-round differential and linear trail cores with weight lower than 86 and 88, respectively. Timings are rounded to the closest integer.

search space	# cores	time	search details		
			step	# cores	time
\mathcal{D}_4^{86}	0	310h	generation of \mathcal{D}_2^{43}	704,744,005	100h
			forw.ext. to 3 rounds with $w < 84$	2,421,335	140h
			forw.ext. to 4 rounds with $w < 86$	0	1m
			back.ext. to 3 rounds with $w < 84$	2,424	66h
			back.ext. to 4 rounds with $w < 86$	0	3h
\mathcal{L}_4^{88}	0	2641h	generation of \mathcal{L}_2^{44}	11,866,934,404	397h
			forw.ext. to 3 rounds with $w < 86$	44,850,380	1411h
			forw.ext. to 4 rounds with $w < 88$	0	25m
			back.ext. to 3 rounds with $w < 86$	40,013	671h
			back.ext. to 4 rounds with $w < 88$	0	161h

6.3 Results on 5 rounds: new (non-tight) bounds

The best known differential trail over 5 rounds has weight 190 [DEMS15, GPT21], while the best known linear trail has weight 184 [MR22]. As far as we know, there are no proved lower bounds for 5-round trails, before this work. We can prove non-tight bounds of 100 for differential trails and 96 for linear trails. To this end, we scanned the spaces \mathcal{D}_5^{100} and \mathcal{L}_5^{96} , which resulted to be both empty. Our search took around 158 days for differential trails and around 127 days for linear trails.

To perform our search, we followed the approach of [DHP⁺20], to re-use the 2-round trail cores already built. We split the space \mathcal{D}_5^{100} (resp. \mathcal{L}_5^{96}) into two sets. The first contains all 5-round trail cores with $w_{\text{rev}}(a^1) + w(b^1) < 43$ (resp. < 44). To cover it, we extend all 2-round trail cores in \mathcal{D}_2^{43} (resp. \mathcal{L}_2^{44}), that we already have, by 3 rounds in the forward direction below weight 100 (resp. 96). The second set contains all 5-round trail cores with $w_{\text{rev}}(a^1) + w(b^1) \geq 43$ (resp. ≥ 44). This implies that $w(b^2) + w(b^3) + w(b^4) < 57$ (resp. < 52). Therefore, we generated all 3-round trail cores in \mathcal{D}_3^{37} (resp. \mathcal{L}_3^{52}) and extended them backwards below weight 100 (resp. 96).

Details on the different steps of our search are reported in Table 5. As we didn't need to regenerate the 2-round trail cores in \mathcal{D}_2^{43} and \mathcal{L}_2^{44} (because we already generated them for the search over 4 rounds), we report the corresponding time between parentheses and we don't consider it in the total cost of this search.

6.4 Results on 6 rounds: improved bounds beyond 2^{-128}

The previously proved lower bound on the weight of 6-round trails is 108, for both linear and differential trails [EME22]. With our techniques we can prove that the spaces \mathcal{D}_6^{129} and \mathcal{D}_6^{132} are both empty. It follows that any 6-round differential trail has weight at least 129 and any 6-round linear trail has weight at least 132. Even if our new bounds are still not tight, we are able to prove for the first time that 6-round trails in ASCON have differential probability or squared correlation lower than 2^{-128} .

Our search took around 6 days for differential trails and around 21 days for linear trails. While in [EME22], the authors report a cost of 2 months each for differential and linear trails. Both in this work and in [EME22], results for 6 rounds are built on top of results on 3 and 4 rounds, whose cost is not included in the figures for 6 rounds. Even if we include

Table 5: Results on the generation of canonical 5-round differential and linear trail cores with weight lower than 100 and 96, respectively. Timings between parentheses mean that we can reuse previous results and they are not counted in the total amount of time.

search space	# cores	time	search details		
			step	# cores	time
\mathcal{D}_5^{100}	0	3795h	generation of \mathcal{D}_2^{43}	704,744,005	(100h)
			forw.ext. by 3 rounds with $w < 100$	0	3683h
			generation of \mathcal{D}_3^{57}	437	112h
			back.ext. by 2 rounds with $w < 100$	0	3s
\mathcal{L}_5^{96}	0	3045h	generation of \mathcal{L}_2^{44}	11,866,934,404	(397h)
			forw.ext. by 3 rounds with $w < 96$		3037h
			generation of \mathcal{L}_3^{52}	309	8h
			back.ext. by 2 rounds with $w < 96$	0	1s

894 such costs in the total computational cost for 6 rounds, our technique still requires less
 895 time compared to [EME22] to reach better bounds.

896 To scan the space \mathcal{D}_6^{129} (resp. \mathcal{L}_6^{132}), we followed the approach of [DMA22]. First, we
 897 split the space in two subspaces that we denote \mathcal{S}_1 and \mathcal{S}_2 . The set \mathcal{S}_1 contains all 6-round
 898 trail cores with $w_{\text{rev}}(a^1) + w(b^1) + w(b^2) < 57$ (resp. < 52) or $w(b^3) + w(b^4) + w(b^5) < 57$
 899 (resp. < 52). The space \mathcal{S}_2 is the complement of \mathcal{S}_1 , that is the space of all 6-round trail
 900 cores with $w_{\text{rev}}(a^1) + w(b^1) + w(b^2) \geq 57$ (resp. ≥ 52) and $w(b^3) + w(b^4) + w(b^5) \geq 57$
 901 (resp. ≥ 52).

902 The details of our search are summarized here.

903 **Scanning \mathcal{S}_1 starting from \mathcal{D}_3^{57} (resp. \mathcal{L}_3^{52}).** The space \mathcal{S}_1 can be scanned by extending
 904 all 3-round trail cores in \mathcal{D}_3^{57} (resp. \mathcal{L}_3^{52}) by 3 rounds below weight 129 (resp. 132). We
 905 first extended all 3-round trails in the space by 3 rounds in the forward direction and then
 906 by 3 rounds in the backward direction. To extend to 6 rounds, we first extended to 4
 907 rounds below 121 (resp. 122) because we know that the two remaining rounds will weight
 908 at least 8. Then we extended to 5 rounds below 127 (resp. 130) because we know that the
 909 remaining round will weigh at least 2. For both differential and linear case, extension to
 910 5 rounds resulted in an empty set. Therefore, we didn't need to perform extension to 6
 911 rounds.

912 **Scanning \mathcal{S}_2 starting from \mathcal{D}_2^{43} (resp. \mathcal{L}_2^{44}).** The space \mathcal{S}_2 is further split into three
 913 subsets. In fact, any 6-round trail core with weight below 129 (resp. 132) can be generated
 914 by starting from a 2-round trail core of weight below 43 (resp. 44) placed at the beginning,
 915 or in the middle, or at the end of the trail. In the first case, the 2-round trail core is
 916 extended by four rounds in the forward direction. In the second case, it is extended by
 917 two rounds in the forward direction and two rounds in the backward direction. In the last
 918 case, it is extended by four rounds in the backward direction.

919 • **Starting from the beginning.** To extend 2-round trail cores to 6 rounds, we
 920 performed extension by one round at the time each time limiting the weight up to
 921 which we perform extension, considering the minimum contribution of the remaining
 922 rounds.

923 First, we extended 2-round trail cores to 3 rounds below $129 - 57 = 72$ (resp.
 924 $132 - 52 = 80$) because we are in the case where $w(b^3) + w(b^4) + w(b^5) \geq 57$ (resp.

Table 6: Results on the generation of canonical 6-round differential and linear trail cores with weight lower than 129 and 132, respectively. Timings between parentheses mean that we can reuse previous results and they are not counted in the total amount of time. - means that the step was not performed, because we know it leads to an empty space.

search space	# cores	time	search details		
			step	# cores	time
\mathcal{D}_6^{129}	0	135h	generation of \mathcal{D}_3^{57}	437	112h
			forw.ext. by 3 rounds with $w < 129$	0	9h
			backw.ext. by 3 rounds with $w < 129$	0	11h
			generation of \mathcal{D}_2^{43}	704,744,005	(100h)
			\mathcal{D}_2^{43} at the beginning		
			- f.e. to 3 rounds with $57 \leq w < 72$	43,465	(140h)
			- f.e. to 4 rounds with $w < 121$	0	3h
			\mathcal{D}_2^{44} in the middle	0	-
\mathcal{D}_2^{44} at the end	0	-			
\mathcal{L}_6^{132}	0	493h	generation of \mathcal{L}_3^{52}	309	(8h)
			forw.ext. by 3 rounds with $w < 132$	0	7h
			backw.ext. by 3 rounds with $w < 132$	0	450h
			generation of \mathcal{L}_2^{44}	11,866,934,404	(397h)
			\mathcal{L}_2^{44} at the beginning		
			- f.e. to 3 rounds with $52 \leq w < 80$	5,171,116	(1411h)
			- f.e. to 4 rounds with $w < 124$	14,082	36h
			- f.e. to 5 rounds with $w < 130$	0	1s
\mathcal{L}_2^{44} in the middle	0	-			
\mathcal{L}_2^{44} at the end	0	-			

925 ≥ 52). Among the obtained 3-round trail cores, we kept only those satisfying
926 $w_{\text{rev}}(a^1) + w(b^1) + w(b^2) \geq 57$ (resp. ≥ 52) because otherwise they belong to \mathcal{S}_1 .
927 Notice that the set of such trail cores is a subset of the set obtained during the search
928 over 4 rounds. In that case in fact, we extended all trail cores in \mathcal{D}_2^{43} (resp. \mathcal{L}_2^{44})
929 to 3 rounds below weight 84 (resp. 86). Therefore, we did not need to perform this
930 step but we just extracted the needed trail cores from such set.

931 Then, we extended the obtained 3-round trail cores to 4 rounds below $129 - 8 = 121$
932 (resp. $132 - 8 = 124$) because we know that $w(b^4) + w(b^5) \geq 8$, since any 2-round
933 trail has weight at least 8.

934 The obtained 4-round trail cores were then extended to 5 rounds below $129 - 2 = 127$
935 (resp. $132 - 2 = 130$) because we know that $w(b^5) \geq 2$.

936 Finally, we extended the obtained 5-round trail cores to 6 rounds below 129 (resp.
937 132).

938 Notice that, for differential trails, extension to 4 rounds already resulted in an empty
939 set. Therefore, extension to 5 and 6 rounds was not performed. For linear trails, it
940 is extension to 5 rounds that gave an empty set. Therefore, we could skip extension
941 to 6 rounds.

942 • **Starting from the middle.** We can assume that $w_{\text{rev}}(a^1) + w(b^1) \geq 43$ (resp.

943 ≥ 44) because the other case is covered in the previous step. First, we need to
 944 perform forward extension to 4 rounds below $129 - 43 = 86$ (resp. $132 - 44 = 88$)
 945 because $w_{\text{rev}}(a^1) + w(b^1) \geq 43$ (resp. ≥ 44). Notice that we already performed this
 946 search in Section 6.2. In fact, this was part of the search to build \mathcal{D}_4^{86} (resp. \mathcal{L}_4^{88}),
 947 which is empty. Therefore, we did not need to perform this step of the search.

948 • **Starting from the end.** We can assume that $w_{\text{rev}}(a^1) + w(b^1) \geq 43$ (resp. ≥ 44)
 949 and $w(b^2) + w(b^3) \geq 43$ (resp. ≥ 44), because the opposite is already covered in the
 950 two previous steps. First, we need to perform backward extension to 4 rounds below
 951 $129 - 43 = 86$ (resp. $132 - 44 = 88$) because $w_{\text{rev}}(a^1) + w(b^1) \geq 43$ (resp. ≥ 44).
 952 Again, we already performed this search in Section 6.2 to build \mathcal{D}_4^{86} (resp. \mathcal{L}_4^{88}). As
 953 we already know that this leads to an empty set, we can jump this step of the search.

954 Figures on the number of trail cores found in each step of the search and details on the
 955 execution time of each step are given in Table 6. When we can reuse trail cores generated
 956 in previous searches, we put the corresponding computational time between parentheses
 957 and we don't include it in the total cost. When a step is not performed because we know
 958 that it leads to an empty space, we put a dash.

959 6.5 Results on 8 rounds: improved (non-tight) bounds

960 Since \mathcal{D}_4^{86} and \mathcal{L}_4^{88} are empty, we can claim that also \mathcal{D}_8^{172} and \mathcal{L}_8^{176} are empty. In fact, if
 961 we split any 8-round differential (resp. linear) trail with weight < 172 (resp. < 176) in
 962 two 4-round trails, at least one of the two must have weight < 86 (resp. < 88). Otherwise,
 963 their sum would be ≥ 172 (resp. ≥ 176). Therefore, all 8-round differential (resp. linear)
 964 trails with weight below 172 (resp. 176) can be obtained by the extension of all 4-round
 965 trails with weight below 86 (resp. 88). But, we know that such 4-round trails do not exist.
 966 Therefore, also such 8-round trails do not exist. It follows that 172 is a lower bound on
 967 the weight of any 8-round differential trail and 176 is a lower bound on the weight of any
 968 8-round linear trail. Such bounds improve over previous known bound, which was 144 for
 969 both differential and linear trails. However, they are still non-tight.

970 6.6 Results on 12 rounds: improved bounds beyond 2^{-256}

971 With a reasoning similar to the one used for 8 rounds, we can prove that the spaces \mathcal{D}_{12}^{258}
 972 and \mathcal{L}_{12}^{264} are empty, given that the spaces \mathcal{D}_6^{129} and \mathcal{L}_6^{132} are empty. It follows that any
 973 12-round differential trail has weight at least 258 and any 12-round linear trail has weight
 974 at least 264. Such bounds improve over previous known bound, which was 216 for both
 975 differential and linear trails. Even if our new bounds are still non-tight, they allow us
 976 to prove for the first time that 12-round trails in ASCON have differential probability or
 977 squared correlation lower than 2^{-256} .

978 6.7 Results on 7, 9, 10, and 11 rounds: improved (non-tight) bounds

979 Based on the results obtained for 4, 5, and 6 rounds, we can derive bounds on 7, 9, 10,
 980 and 11 rounds. We explain how to do it for 10 rounds by combining the results for 4 and
 981 6 rounds. Then we show how to obtain bounds for the other numbers of rounds similarly.

982 We can cover the space \mathcal{D}_{10}^{215} (resp. \mathcal{L}_{10}^{220}) in the following way. We split the set in two
 983 subsets. The first contains all 10-round trail cores with $w_{\text{rev}}(a^1) + w(b^1) + w(b^2) + w(b^3) < 86$
 984 (resp. 88), while the second set is its complement. We can cover the first set by extending all
 985 4-round trail cores in \mathcal{D}_4^{86} (resp. \mathcal{L}_4^{88}) by 6 rounds in the forward direction below 215 (resp.
 986 220). The second set contains all 10-round trails with $w_{\text{rev}}(a^1) + w(b^1) + w(b^2) + w(b^3) \geq 86$
 987 (resp. 88), which implies that the other 6 rounds have weight below 129 (resp. 132).
 988 Therefore, we can cover it by extending all 6-round trail cores in \mathcal{D}_6^{129} (resp. \mathcal{L}_6^{132}) by 4

989 rounds in the backward direction below 215 (resp. 220). Since both \mathcal{D}_4^{86} and \mathcal{D}_6^{129} (resp.
 990 \mathcal{L}_4^{88} and \mathcal{L}_6^{132}) are empty, then also \mathcal{D}_{10}^{215} (resp. \mathcal{L}_{10}^{220}) is empty. Therefore, 215 and 220
 991 are lower bounds on the weight of 10-round differential and linear trails, respectively.

992 For the other numbers of rounds, we consider the combination that yields the best
 993 bounds. For 7 rounds, we can prove a bound of 131 for differential trails and 134 for linear
 994 trails, considering the results on 6 rounds and that 1 round weights at least 2. For 9 rounds,
 995 we combine the results for 4 and 5 rounds and obtain a bound of 186 for differential trails
 996 and 184 for linear trails. Finally, for 11 rounds we obtain a bound of 229 for differential
 997 trails and 228 for linear trails, by combining the results for 5 and 6 rounds.

998 For the sake of comparison, we can apply the same reasoning to the results presented
 999 in [EME22]. We can derive bounds for r rounds from the bounds on $r - 1$ rounds,
 1000 considering that one round has minimum weight 2. For differential and linear trails, this
 1001 gives a bound of 74 for 5 rounds, of 110 for 7 rounds, of 146 for 9 rounds, and 182 for 11
 1002 rounds.

1003 7 Conclusions

1004 In this work, we presented a dedicated tool for trail search in ASCON, based on the 2-round
 1005 trail core generation methods given in [MDV17] and improved methods for extension based
 1006 on the works done in [DV12, DHVV18b]. Using these techniques, we proved tight bound
 1007 for 3-rounds linear trails and improved the existing bounds for other number of rounds.
 1008 In particular, we prove for the first time bounds beyond 2^{-128} for 6 rounds, and for 12
 1009 rounds bounds beyond 2^{-256} . Our approach improves and proves bounds in a reasonable
 1010 amount of time and it confirms that dedicated tools can still outperform methods based
 1011 on general-purpose solvers.

1012 As a takeaway from this and previous works on KECCAK- p [MDV17], XOODOO
 1013 [DHVV18b], and SUBTERRANEAN [MMGD22] we highlight that:

- 1014 • For the 2-round trail search stage, the linear layers of ASCON and SUBTERRANEAN
 1015 allow a simpler definition of units compared to KECCAK- p and XOODOO where a
 1016 more complex linear layer is used.
- 1017 • A non-linear layer based on the parallel application of small S-boxes (as in KECCAK- p ,
 1018 XOODOO and ASCON) implies a simpler analysis of the propagation properties
 1019 compared to the non-linear layer of SUBTERRANEAN. In the latter case, the backward
 1020 extension is more complex, and the definition of the minimum reverse weight requires
 1021 a thorough proof which makes it more complicated.

1022 Acknowledgements

1023 Solane El Hirsch and Silvia Mella are supported by the Cryptography Research Center of
 1024 the Technology Innovation Institute (TII), Abu Dhabi (UAE), under the TII-Radboud
 1025 project with title *Evaluation and Implementation of Lightweight Cryptographic Primitives*
 1026 *and Protocols*.

1027 Alireza Mehrdad and Joan Daemen are supported by the European Research Council under
 1028 the ERC advanced grant agreement under grant ERC-2017-ADG Nr. 788980 ESCADA.

1029 References

- 1030 [BDPV07] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Sponge
 1031 functions. <https://keccak.team/files/SpongeFunctions.pdf>, 2007.

- 1032 [BDPV08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On
1033 the indifferentiability of the sponge construction. In Nigel P. Smart, editor,
1034 *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International
1035 Conference on the Theory and Applications of Cryptographic Techniques,
1036 Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture
1037 Notes in Computer Science*, pages 181–197. Springer, 2008.
- 1038 [BDPV11a] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Duplexing
1039 the sponge: Single-pass authenticated encryption and other applications.
1040 In Ali Miri and Serge Vaudenay, editors, *Selected Areas in Cryptography -
1041 18th International Workshop, SAC 2011, Toronto, ON, Canada, August 11-12,
1042 2011, Revised Selected Papers*, volume 7118 of *Lecture Notes in Computer
1043 Science*, pages 320–337. Springer, 2011.
- 1044 [BDPV11b] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. The
1045 keccak reference, January 2011.
- 1046 [BJK⁺16] Christof Beierle, Jérémy Jean, Stefan Kölbl, Gregor Leander, Amir Moradi,
1047 Thomas Peyrin, Yu Sasaki, Pascal Sasdrich, and Siang Meng Sim. The
1048 SKINNY Family of Block Ciphers and Its Low-Latency Variant MANTIS.
1049 In *Advances in Cryptology - CRYPTO 2016*, volume 9815 of *LNCS*, pages
1050 123–153. Springer, 2016.
- 1051 [BPP⁺17] Subhadeep Banik, Sumit Kumar Pandey, Thomas Peyrin, Yu Sasaki,
1052 Siang Meng Sim, and Yosuke Todo. GIFT: A Small Present - Towards
1053 Reaching the Limit of Lightweight Encryption. In *Cryptographic Hardware
1054 and Embedded Systems - CHES 2017*, volume 10529 of *LNCS*, pages 321–345.
1055 Springer, 2017.
- 1056 [com14] CAESAR committee. CAESAR: Competition for authenticated encryption:
1057 Security, applicability, and robustness, 2014.
- 1058 [Dae95] Joan Daemen. *Cipher and hash function design, strategies based on linear
1059 and differential cryptanalysis, PhD Thesis*. PhD thesis, K.U.Leuven, 1995.
- 1060 [DEM15] Christoph Dobraunig, Maria Eichlseder, and Florian Mendel. Heuristic tool for
1061 linear cryptanalysis with applications to CAESAR candidates. In Tetsu Iwata
1062 and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 -
1063 21st International Conference on the Theory and Application of Cryptology
1064 and Information Security, Auckland, New Zealand, November 29 - December
1065 3, 2015, Proceedings, Part II*, volume 9453 of *Lecture Notes in Computer
1066 Science*, pages 490–509. Springer, 2015.
- 1067 [DEM⁺20] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel,
1068 Bart Mennink, Robert Primas, and Thomas Unterluggauer. Isap v2.0. *IACR
1069 Trans. Symmetric Cryptol.*, 2020(S1):390–416, 2020.
- 1070 [DEMS15] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer.
1071 Cryptanalysis of ascon. In Kaisa Nyberg, editor, *Topics in Cryptology - CT-
1072 RSA 2015, The Cryptographer’s Track at the RSA Conference 2015, San
1073 Francisco, CA, USA, April 20-24, 2015. Proceedings*, volume 9048 of *Lecture
1074 Notes in Computer Science*, pages 371–387. Springer, 2015.
- 1075 [DEMS16] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer.
1076 Ascon v1.2. submission to caesar competition. Technical report, 2016.

- 1077 [DEMS21a] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer.
1078 Ascon v1.2: Lightweight authenticated encryption and hashing. *J. Cryptol.*,
1079 34(3):33, 2021.
- 1080 [DEMS21b] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer.
1081 Ascon v1.2. submission to nist. Technical report, 2021.
- 1082 [DHP⁺20] Joan Daemen, Seth Hoffert, Micha el Peeters, Gilles Van Assche, and Ronny
1083 Van Keer. Xoodyak, a lightweight cryptographic scheme. *IACR Trans.*
1084 *Symmetric Cryptol.*, 2020(S1):60–87, 2020.
- 1085 [DHVV13] J. Daemen, S. Hoffert, G. Van Assche, and R. Van Keer. KeccakTools software.
1086 <https://github.com/KeccakTeam/KeccakTools>, 2013.
- 1087 [DHVV18a] J. Daemen, S. Hoffert, G. Van Assche, and R. Van Keer. XooTools software.
1088 <https://github.com/XoodooTeam/Xoodoo>, 2018.
- 1089 [DHVV18b] Joan Daemen, Seth Hoffert, Gilles Van Assche, and Ronny Van Keer. The
1090 design of xoodoo and xoeff. *IACR Trans. Symmetric Cryptol.*, 2018(4):1–38,
1091 2018.
- 1092 [DMA22] Joan Daemen, Silvia Mella, and Gilles Van Assche. Tighter trail bounds
1093 for xoodoo. Cryptology ePrint Archive, Paper 2022/1088, 2022. <https://eprint.iacr.org/2022/1088>.
- 1095 [DPAR00] Joan Daemen, Micha el Peeters, Gilles Van Assche, and Vincent Rijmen.
1096 Nessie proposal: the block cipher NOEKEON. Nessie submission, 2000. <http://gro.noekeon.org/>.
- 1098 [DR20] Joan Daemen and Vincent Rijmen. *The Design of Rijndael - The Advanced*
1099 *Encryption Standard (AES), Second Edition*. Information Security and Crypt-
1100 ography. Springer, 2020.
- 1101 [DV12] Joan Daemen and Gilles Van Assche. Differential propagation analysis of kec-
1102 cak. In Anne Canteaut, editor, *Fast Software Encryption - 19th International*
1103 *Workshop, FSE 2012, Washington, DC, USA, March 19-21, 2012. Revised*
1104 *Selected Papers*, volume 7549 of *Lecture Notes in Computer Science*, pages
1105 422–441. Springer, 2012.
- 1106 [EME22] Johannes Erlacher, Florian Mendel, and Maria Eichlseder. Bounds for the
1107 security of ascon against differential and linear cryptanalysis. *IACR Trans.*
1108 *Symmetric Cryptol.*, 2022(1):64–87, 2022.
- 1109 [GPT21] David G erault, Thomas Peyrin, and Quan Quan Tan. Exploring differential-
1110 based distinguishers and forgeries for ASCON. *IACR Trans. Symmetric*
1111 *Cryptol.*, 2021(3):102–136, 2021.
- 1112 [MDV17] Silvia Mella, Joan Daemen, and Gilles Van Assche. New techniques for
1113 trail bounds and application to differential trails in Keccak. *IACR Trans.*
1114 *Symmetric Cryptol.*, 2017(1):329–357, 2017.
- 1115 [MMGD22] Alireza Mehrdad, Silvia Mella, Lorenzo Grassi, and Joan Daemen. Differential
1116 trail search in cryptographic primitives with big-circle chi: Application to
1117 subterranean. *IACR Trans. Symmetric Cryptol.*, 2022(2):253–288, 2022.
- 1118 [MP13] Nicky Mouha and Bart Preneel. Towards Finding Optimal Differential Charac-
1119 teristics for ARX: Application to Salsa20. Cryptology ePrint Archive, Report
1120 2013/328, 2013. <https://ia.cr/2013/328>.

- 1121 [MR22] Rusydi H. Makarim and Raghvendra Rohit. Towards tight differential bounds
1122 of ascon: A hybrid usage of smt and milp. *IACR Transactions on Symmetric*
1123 *Cryptology*, 2022(3):303–340, 2022.
- 1124 [SHW⁺14] Siwei Sun, Lei Hu, Peng Wang, Kexin Qiao, Xiaoshuang Ma, and Ling Song.
1125 Automatic Security Evaluation and (Related-key) Differential Characteristic
1126 Search: Application to SIMON, PRESENT, LBlock, DES(L) and Other
1127 Bit-Oriented Block Ciphers. In *Advances in Cryptology - ASIACRYPT 2014*,
1128 volume 8873 of *LNCS*, pages 158–178. Springer, 2014.
- 1129 [TMC⁺21] Meltem Sonmez Turan, Kerry McKay, Donghoon Chang, Cagdas Calik,
1130 Lawrence Bassham, Jinkeon Kang, and John Kelsey. Status report on the
1131 second round of the nist lightweight cryptography standardization process,
1132 2021.
- 1133 [WH19] Hongjun Wu and Tao Huang. TinyJAMBU: A Family of Lightweight Authen-
1134 ticated Encryption Algorithms, 2019.

1135 **A Representation of the affine spaces over \mathcal{S} .**

1136 In Table 7, for each possible column difference, we provide a representation of the affine
1137 space of compatible differences at the output of \mathcal{S} , the restriction weight, and the minimum
1138 reverse weight. In Table 8, for each possible column mask, we provide a representation of
1139 the affine space of compatible masks at the input of \mathcal{S} , the correlation weight, and the
1140 minimum reverse weight.

Table 7: Space of compatible differences at the output of p_S , restriction weight, and minimum reverse weight for all possible column differences.

difference	Affine space after \mathcal{S}	$w_r(\cdot)$	$w_{\text{rev}}(\cdot)$
00000	00000	0	0
00001	01001 + $\langle 00010, 00100, 10001 \rangle$	3	2
00010	10001 + $\langle 00010, 00100, 01000 \rangle$	3	2
00011	00001 + $\langle 00100, 01000, 10001 \rangle$	3	3
00100	00110 + $\langle 01000, 10000 \rangle$	2	2
00101	10001 + $\langle 00010, 01001, 01100 \rangle$	3	3
00110	00001 + $\langle 00010, 00100, 01000, 10000 \rangle$	4	2
00111	00010 + $\langle 00001, 00100, 01000 \rangle$	3	3
01000	00110 + $\langle 00001, 01000, 10000 \rangle$	3	2
01001	00001 + $\langle 00010, 10001, 10100, 11000 \rangle$	4	2
01010	00001 + $\langle 00101, 00110, 01000, 10000 \rangle$	4	2
01011	00010 + $\langle 00001, 00100, 01000, 10000 \rangle$	4	2
01100	00001 + $\langle 10001, 11000 \rangle$	2	2
01101	00001 + $\langle 00010, 00100, 10001, 11000 \rangle$	4	3
01110	00001 + $\langle 00101, 00110, 10000 \rangle$	3	2
01111	01000 + $\langle 00001, 00100, 10000 \rangle$	3	3
10000	01001 + $\langle 00010, 10001 \rangle$	2	2
10001	10001 + $\langle 00010, 00100 \rangle$	2	2
10010	00001 + $\langle 00010, 00100, 01000, 10001 \rangle$	4	3
10011	00010 + $\langle 00110, 01000 \rangle$	2	2
10100	00100 + $\langle 00001, 00010, 01000 \rangle$	3	3
10101	00101 + $\langle 00010, 10100, 11000 \rangle$	3	2
10110	10000 + $\langle 00001, 00010, 00100, 01000 \rangle$	4	2
10111	00010 + $\langle 00110, 01000, 10000 \rangle$	3	2
11000	00100 + $\langle 00001, 00010, 01000, 10000 \rangle$	4	2
11001	01000 + $\langle 00101, 10110, 11000 \rangle$	3	2
11010	00001 + $\langle 00100, 01001, 01010, 10000 \rangle$	4	2
11011	00010 + $\langle 00001, 00110, 01000, 10000 \rangle$	4	3
11100	00001 + $\langle 00010, 10001, 11000 \rangle$	3	3
11101	01000 + $\langle 00110, 10101, 11000 \rangle$	3	3
11110	01000 + $\langle 00001, 00010, 00100, 10000 \rangle$	4	2
11111	00010 + $\langle 00001, 00110, 10000 \rangle$	3	3

Table 8: Space of compatible masks at the input of p_S , correlation weight, and minimum reverse weight for all possible column masks.

mask	Affine space before \mathcal{S}	$w_c(\cdot)$	$w_{\text{rev}}(\cdot)$
00000	00000	0	0
00001	00011 + $\langle 01000, 10001 \rangle$	2	2
00010	01100 + $\langle 00011, 10000 \rangle$	2	2
00011	00100 + $\langle 00001, 00010, 01000, 10000 \rangle$	4	2
00100	01100 + $\langle 00001, 00010 \rangle$	2	2
00101	00100 + $\langle 00001, 00010, 01000, 10000 \rangle$	4	2
00110	00001 + $\langle 10001, 10010 \rangle$	2	2
00111	00010 + $\langle 10001, 11010 \rangle$	2	2
01000	10001 + $\langle 01010, 01100 \rangle$	2	2
01001	00001 + $\langle 00010, 00100, 01000, 10001 \rangle$	4	2
01010	00001 + $\langle 01001, 01010, 01100, 10000 \rangle$	4	2
01011	00010 + $\langle 00001, 00100, 10010, 11000 \rangle$	4	2
01100	10000 + $\langle 00001, 00010, 00100, 01000 \rangle$	4	2
01101	00001 + $\langle 00010, 00101, 01000, 10000 \rangle$	4	2
01110	00001 + $\langle 10001, 10010, 10100, 11000 \rangle$	4	2
01111	00010 + $\langle 00001, 01010, 01100, 10000 \rangle$	4	2
10000	00011 + $\langle 01000, 10101 \rangle$	2	2
10001	10001 + $\langle 00100, 01000 \rangle$	2	2
10010	00001 + $\langle 00101, 00110, 01000, 10000 \rangle$	4	2
10011	00001 + $\langle 00011, 00100, 01000, 10000 \rangle$	4	2
10100	00100 + $\langle 00001, 00010, 01000, 10100 \rangle$	4	4
10101	10000 + $\langle 00001, 00010, 00100, 01000 \rangle$	4	2
10110	00001 + $\langle 00100, 01000, 10001, 10010 \rangle$	4	2
10111	00001 + $\langle 00100, 01000, 10001, 10010 \rangle$	4	2
11000	00001 + $\langle 00010, 00100, 01000, 10001 \rangle$	4	2
11001	00100 + $\langle 00010, 01100 \rangle$	2	2
11010	00001 + $\langle 00010, 00100, 10001, 11000 \rangle$	4	2
11011	00100 + $\langle 00001, 00010, 01100, 10000 \rangle$	4	2
11100	00001 + $\langle 00010, 01000, 10001, 10100 \rangle$	4	2
11101	01000 + $\langle 00010, 01101 \rangle$	2	2
11110	00001 + $\langle 00010, 10001, 10100, 11000 \rangle$	4	2
11111	00100 + $\langle 00001, 00010, 01100, 10000 \rangle$	4	2