

Construction of Differential Characteristics in ARX Designs

Application to Skein

Gaëtan Leurent

UCL Crypto Group**

Gaetan.Leurent@uclouvain.be

Abstract. In this paper, we study differential attacks against ARX schemes. We build upon the generalized characteristics of de Cannière and Rechberger and the multi-bit constraints of Leurent. Our main result is an algorithm to build complex non-linear differential characteristics for ARX constructions, that we applied to reduced versions of the hash function Skein. We present several characteristics for use in various attack scenarios: on the one hand we show attacks with a relatively low complexity, in relatively strong settings; and on the other hand weaker distinguishers reaching more rounds. Our most notable results are practical free-start and semi-free-start collision attacks for 20 rounds and 12 rounds of Skein-256, respectively. Since the full version of Skein-256 has 72 rounds, this result confirms the large security margin of the design.

These results are some of the first examples of complex differential trails built for pure ARX designs. We believe this is an important work to assess the security those functions against differential cryptanalysis. Our tools will be publicly available with the final version of this paper.

Keywords: Symmetric ciphers, Hash functions, ARX, Generalized characteristics, Differential attacks, Skein

1 Introduction

Most symmetric key cryptographic primitives can be classified as either S-Box based designs or ARX designs. The so-called ARX designs use only Additions ($a \boxplus b$), Rotations ($a \ggg i$), and Xors ($a \oplus b$). These operations are very simple and can be implemented efficiently in software or in hardware, but when mixed together, they interact in complex and non-linear ways. ARX designs have been quite popular recently; in particular, two of the SHA-3 finalists, BLAKE and Skein, follow this design strategy. This strategy has also been used for stream ciphers such as Salsa20 and ChaCha, and block ciphers, such as TEA, XTEA or HIGHT (RC5 uses additions and data-dependant rotations, but we only consider construction with fixed rotations). Recently, a dedicated short-input PRF, SipHash [AB12], has been built following the ARX design. We note that Salsa20 is in the eStream portfolio, while SipHash is already deployed as the default hash table implementation of the Perl and Ruby languages. More generally, functions of the MD/SHA family are built using Additions, Rotations, Xors, but also bitwise Boolean functions, and logical shifts; they are sometimes also referred to as ARX.

The ARX design philosophy is opposed to S-Box based designs such as the AES. Analysis of S-Box based designs usually happen at the word-level; differential characteristics are relatively easy to build, but efficient attacks often need novel techniques, such as the rebound attack against hash functions [MRST09]. For ARX designs, the analysis is done on a bit-level; finding good differential characteristics remains an important challenge. In particular, the seminal attacks on the MD/SHA-family by the team of X. Wang are based on differential characteristics built by hand [WLF⁺05,WY05,WYY05], and a significant effort has been dedicated to building tools to construct automatically such characteristics [DR06,SO06,FLN07a,MNS11,SLdW07,MNS12,MNS13]. This effort has been quite successful for functions of the MD/SHA family, and it has allowed new attacks based on specially designed characteristics: attacks against HMAC [FLN07b], the construction of a rogue MD5 CA certificate [SSA⁺09], and attacks against combiners [MRS09].

** Part of this work was done when the author was at the University of Luxembourg.

However, this body of work is mainly focused on MD/SHA designs, as opposed to pure ARX designs such as Skein, BLAKE or Salsa20. In MD/SHA-like functions, the Boolean functions play an important role, and the possibility to absorb differences gives a lot of freedom for the construction of differential characteristics. In pure ARX designs, the addition is the only source of non-linearity, and the freedom in the carry expansions is much harder to use than the absorption property of Boolean functions.

To this effect, Leurent introduced multi-bit constraints [Leu12] involving several consecutive bits of a variable (*i.e.* $x^{[i]}$ and $x^{[i-1]}$), instead of considering bits one by one. He describes reduced sets of 1.5-bit and 2.5-bit constraints, and explains how to propagate these constraints using S-systems and automata. This set of constraints is well suited to study ARX designs because it can extract a lot of information about the carry extensions in modular additions. A set of tools to propagate these constraints is given in [Leu12], and the main result is a negative result (for the cryptanalyst) showing that several previous attacks are invalid.

1.1 Our Results

In this paper, we study the problem of constructing differential characteristics for ARX schemes. This work is heavily inspired by the framework of generalized characteristics from de Cannière and Rechberger [DR06], and the multi-bit constraints of [Leu12]. As opposed to the results of [Leu12], we give positive results for cryptanalysts.

We first recall how to describe a differential characteristic, and the main ideas for constraint propagation in Section 2. Then, we describe a differential characteristic search algorithm in Section 3 using a constraint propagation tool, and we present our results on Skein in Section 4. Finally, we describe our technical improvements over the previous constraint propagation tools in Appendix A.

Construction of differential characteristics. We use a propagation tool to construct differential characteristics automatically. Using an efficient constraint propagation tool and some simple heuristics, we show that we can actually build complex non-linear characteristics. We obtain some of the first complex differential trails for ARX designs and we believe that this automated approach is an important step to assess the security of ARX designs against differential cryptanalysis.

Application to Skein. We apply this technique to reduced versions of the Skein hash function, where we build rebound-like characteristics by connecting two high-probability trails.

We compare our results with previous works in Table 1. Most previous works on Skein are either weak distinguishers (such as boomerang properties or free-tweak free-start partial-collisions), or attack with marginal improvement over brute-force (such as some biclique-based results). In this work, we present attacks in relatively strong settings (collisions and free-start collisions) with a relatively low complexity (several attacks are practical, and all our attack gain at least a factor 2^8).

Constraint propagation. Finally, we describe an alternative way to perform the constraint propagation for multi-bit constraints. Our approach is significantly more efficient than the technique of [Leu12], and uses the full set of 2^{32} constraints instead of a reduced set of 16 carefully chosen constraints. The reduced set is sufficient in most situations, but we show that the full set extracts some more information. This improvement was crucial to allow the characteristic search to work in practice.

In addition, our approach can also deal with larger systems than the previous technique with a reasonable complexity. In particular, we can deal with the 3-input modular sums, and 3-input Boolean functions used in functions of the MD/SHA family. We can also propagate 4

simultaneous trails in a boomerang configuration through an addition or an xor, with full 2-bit constraints.

Table 1. Comparison of attacks on reduced versions of Skein-256 (we omit attack on previous versions, and weak distinguishers). The full skein-256 has 72 rounds.

In order to compare various attack settings, we count the number of extra degrees of freedom used by the attack.

	Extra Degrees of freedom	Rounds	Time	Generic	Ref, notes
Collision	0	4	2^{96}	2^{128}	[KRS12], Biclique based
		8	2^{120}		
		9	2^{124}		
		12	$2^{126.5}$		
Free-start collision	8	22 [†]	$2^{253.8†}$	2^{256}	[LIS12], Biclique based
		37 [†]	$2^{255.7†}$		
Related-tweak [‡] partial q -multicollision	10	20	$q \cdot 2^{97}$	$2^{\frac{q-1}{q+1} \cdot 130}$	[SWWD10], 126 active bits
Free-tweak partial q -multicollision	12	32	$q \cdot 2^{105}$	$2^{\frac{q-1}{q+1} \cdot 205}$	[YCW13], 51 active bits
Collision	0	12	$\approx 2^{114*}$	2^{128}	4.4
Semi-free-start collision	4	12	$\approx 2^{40}$	2^{128}	4.4
Free-start collision	8	20	$\approx 2^{40}$	2^{128}	4.5
Free-start near-collision	8	24	$\approx 2^{40}$	$2^{88.4}$	4.5, 15 active bits
Related-tweak [‡] near-collision	10	24	$\approx 2^{40}$	$2^{117.3}$	4.6, 3 active bits
Related-tweak [‡] partial q -multicollision	10	32	$\approx q \cdot 2^{119*}$	$2^{\frac{q-1}{q+1} \cdot 205}$	4.6, 51 active bits
Free-tweak partial q -multicollision	12	32	$q \cdot 2^{105}$	$2^{\frac{q-1}{q+1} \cdot 205}$	4.6, 51 active bits
<i>Block cipher attacks</i>					
Key recovery (Threefish-512)		32	2^{181}	2^{512}	[YCW12], Boomerang
		33	2^{305}		
		34	2^{424}		

[†] Attacks on Skein-512. For Skein-256, fewer round will be attacked, with a complexity slightly below 2^{128} .

[‡] Using freedom degrees in the tweak *difference*, but the tweak *value* can be arbitrary.

* Using heuristic assumptions about the search for a large number of characteristics.

1.2 Related work

A recent result by Yu *et al.* achieves a similar result as our free-start free-tweak partial-collision on 32 rounds, and is also based on a complex non-linear trail for Skein-256. This work has been available on ePrint since April 2011 [YCJW11], but the characteristic given in that version of the paper was flawed [Leu12]. This has motivated our work on building such characteristics automatically.

More recently, they managed to build a valid characteristic and their work will be presented at FSE [YCW13]; this result was achieved simultaneously and independently from our work. Building such a trail by hand is impressive, but this kind of result it is very challenging to replicate or to apply to another primitive. We hope that our automatic approach will be easier to adapt to new settings.

2 Analysis of Differential Characteristics

The first step for working with differential characteristics (or trails) is to choose a way to represent a characteristic, and to evaluate its probability. The main idea of differential cryptanalysis is to consider the computation of the function for a pair of inputs X, X' , and to specify the difference between x and x' for every internal state variable x . The difference can be the xor difference, the modular difference, or more generally, use any group operation. However, this approach is

Table 2. Generalized constraints used in [DR06].

	(x, x') :	(0, 0)	(0, 1)	(1, 0)	(1, 1)
?	<i>anything</i>	✓	✓	✓	✓
-	$x = x'$	✓	-	-	✓
x	$x \neq x'$	-	✓	✓	-
0	$x = x' = 0$	✓	-	-	-
u	$(x, x') = (0, 1)$	-	✓	-	-
n	$(x, x') = (1, 0)$	-	-	✓	-
1	$x = x' = 1$	-	-	-	✓
#	<i>incompatible</i>	-	-	-	-
3	$x = 0$	✓	✓	-	-
5	$x' = 0$	✓	-	✓	-
7		✓	✓	✓	-
A	$x' = 1$	-	✓	-	✓
B		✓	✓	-	✓
C	$x = 1$	-	-	✓	✓
D		✓	-	✓	✓
E		-	✓	✓	✓

not efficient for ARX design, because both the modular difference and the xor difference play an important role. Several works have proposed better way to represent a differential characteristic for ARX designs.

Signed bitwise difference. The groundbreaking results of Wang *et al.* [WLF⁺05,WY05,WYY05] are based on a bitwise signed difference. For each bit of the state, they specify whether the bit is inactive ($x = x'$), active with a positive sign ($x = 0, x' = 1$), or active with a negative sign ($x = 1, x' = 0$). This information express both the xor difference and the modular difference.

Generalized characteristics. This was later generalized by de Cannière and Rechberger [DR06]: for each bit of the state, they look at all possible values of the pair (x, x') , and they specify which values are allowed. This give a set of 16 constraints as shown in Table 2. The constraints -, u and n correspond to the bitwise signed difference of Wang. De Cannière and Rechberger also describe an algorithm to build differential characteristics using this set of constraints.

Multi-bit constraints. Recently, Leurent studied differential characteristics for ARX designs, and introduced multi-bit constraints [Leu12]. These constraints are applied to the values of consecutive bits of a state variable (*e.g.* $x^{[i]}$ and $x^{[i-1]}$) instead of being purely bitwise. Multi-bit constraints are quite efficient to study ARX designs because they can capture the behaviour of carries in the modular addition. Two set of constraints are introduced in [Leu12]:

- a set of 16 constraints involving $(x^{[i]}, x'^{[i]}, x^{[i-1]})$ called 1.5-bit constraints;
- a set of 16 constraints involving $(x^{[i]}, x'^{[i]}, x^{[i-1]}, x'^{[i-1]}, x^{[i-2]})$ called 2.5-bit constraints.

The full sets of 2^8 1.5-bit constraints and 2^{32} 2.5-bit constraint are not used because the propagation method of [Leu12] becomes impractical with such large sets.

2.1 Constraint Propagation and Probability Computation

In [Leu12], the constraints are studied using the theory of S-functions introduced in [MVCP10]. We use the following definitions:

T-function A T-function on n -bit words with k inputs and l outputs is a function from $(\{0, 1\}^n)^k$ to $(\{0, 1\}^n)^l$ with the following property:

For all t , the t least significant bits of the outputs can be computed from the t least significant bits of the inputs.

S-function An S-function on n -bit words is a function from $(\{0, 1\}^n)^k$ to $(\{0, 1\}^n)^l$, for which we can define a small set of *states* \mathcal{S} , and an initial state $S[-1] \in \mathcal{S}$ with the following property:

For all t , bit t of the outputs and the state $S[t] \in \mathcal{S}$ can be computed from bit t of the inputs, and the state $S[t - 1]$.

For instance, the modular addition is an S-function, with a 1-bit state corresponding to the carry. An S-function can also include bitwise functions, shifts to the left by a fixed number of bits, or multiplications by constants. A system of equation that can be written as a S-function is called an S-system.

2.2 Differential Characteristics

In order to describe a differential characteristics with this framework, we specify a difference for each internal variable of a cipher, and we consider the operations that connect the variables. For a series a constraints Δ , we write $\delta x = \Delta$ to denote that the pair (x, x') follows the difference pattern Δ . For instance, $\delta x = \mathbf{x}--0$ is equivalent to $x \oplus x' = 1000$ and $x^{[0]} = 0$.

For each operation \odot , we can write a system:

$$\delta x = \Delta_x \quad \delta y = \Delta_y \quad \delta z = \Delta_z \quad z = x \odot y \quad z' = x' \odot y', \quad (1)$$

where x, y, z, x', y', z' are unknowns, and $\Delta_x, \Delta_y, \Delta_z$ are parameters. In an ARX design, all the operations except the rotations are S-function, and the difference operation δ can be written with bitwise operations and left-shifts; therefore system (1) is an S-system. Using tools to analyze this S-system, we can verify if the specified input and output patterns for each operation are compatible. We deal with the rotations $y = x \ggg i$ by just rotating the constraint pattern: if $\delta x = \Delta_x$ then we use $\delta y = \Delta_x \ggg i$.

We can also find new constraints that must be satisfied for any solution to the system. This allows to propagate constraints between the inputs and outputs of the operation \odot . When we consider a characteristic for a cipher, this process will be iterated for each operation, until no new constraints are found.

Moreover, we can compute the probability to reach the specified output pattern by counting the number of solutions. Assuming that the probabilities of each operations are independent, we can compute the probability of the full characteristic by multiplying the probabilities of each operations.

2.3 Tools for S-systems

In [Leu12], a set of constraints is represented by an S-system, and an automaton is built to compute the probability of each operation. To perform constraints propagation, each constraint is split into two disjoint subsets; if one of the subsets results in an incompatible system, the constraint can be restricted to the other subset without reducing the number of solutions.

This approach allows to achieve a good efficiency when the automaton is fully determinized: one can test whether a system is compatible with only n table accesses. However, the table becomes impractically large if the set of constraints is too large, or if the operation is too complex. In [Leu12], the automaton is fully determinized for 1.5-bit constraints, but could not be determinized for 2.5-bit constraints; this results in a quite inefficient propagation algorithm for 2.5-bit constraints.

In this work, we explore a different option using non-deterministic automata. This allows to deal with large set of constraints and more complex operations. We need to perform many operations to verify whether a system is compatible, but the automata are very sparse and can

be represented by tables small enough to fit in the cache (the tables of [Leu12] take hundreds of megabytes for an addition); this gives better results in practice. In addition, we show special properties of the automata allowing an efficient propagation algorithm without splitting the constraints into subsets. Due to space constraints, the technical details of our new approach are given in Appendix A.

2.4 Comparison

Table 3. Experiments with toy versions of Skein. We give the number of input/output differences accepted by each technique, and the ratio of false positive.

Method	2 rounds / 4 bits (total: 2^{32})		3 rounds / 6 bits (sparse*)		
	Accepted	F pos.	Accepted	F pos.	Time [†]
Exhaustive search	$2^{25.1}$ (35960536)	–	$2^{18.7}$ (427667)	–	
2.5-bit full set	$2^{25.3}$ (40597936)	0.13	$2^{19.2}$ (619492)	0.4	2.5 ms
2.5-bit reduced set [Leu12]	$2^{25.3}$ (40820032)	0.14	$2^{19.5}$ (746742)	0.7	50 ms
1.5-bit reduced set [Leu12]	$2^{25.3}$ (40820032)	0.14	$2^{20.4}$ (1372774)	2.2	0.5 ms
1-bit constraints [DR06]	$2^{25.4}$ (43564288)	0.21	$2^{20.7}$ (1762857)	3.1	0.5 ms
Check adds independently	$2^{25.8}$ (56484732)	0.57			

* Weight 4 differences. The total number of input/output differences is $(\sum_{i=0}^4 \binom{24}{i})^2 \approx 2^{26.75}$.

† Average time to verify one input/output difference (over the false positives of the 1.5-bit reduced set).

We show a comparison with previous methods in Table 3. We use the same settings as [Leu12]:

1. A reduced Skein with two rounds and 4 words of 4 bits each; In this setting the full 2.5-bit constraints offer a little advantage over the reduced set of 2.5-bit constraints.
2. A reduced Skein with three rounds and 4 words of 6 bits each. We only use sparse differences (less than 4 active bits in the input and output), because the full space is too large to be exhausted in practice. In this setting, the full 2.5-bit constraints give a significant improvement over the reduced set of 2.5-bit constraints.

These experiments show that using the full set of 2.5-bit constraints gives better results than using the reduced set of [Leu12]. We also give timing informations¹: our new approach for constraint propagation is one order of magnitude faster than the previous method with a reduced set of 2.5-bit constraints, and somewhat slower than the previous method with 1.5-bit constraints.

3 Automatic Construction of Differential Characteristics

In order to mount a differential attack for a hash function or a block cipher, an important task is to build a differential characteristic. For the analysis of ARX primitives (and MD/SHA-like designs), the characteristic is usually designed at the bit level. This turns out to be a very challenging task because of the complex interactions between the operations, and the large number of state elements to consider.

This problem has been heavily studied for attacks on the MD/SHA family of hash functions: a series of attacks by X. Wang and her team are based on differential characteristics built by hand [WLF⁺05,WY05,WYY05,YCW13], while later works gave algorithms to build such characteristics automatically [DR06,SO06,FLN07a,MNS11,SLdW07]. Unfortunately, most of those tools are not publicly available.

In this section, we show that the multi-bit constraints can be used to design a successful algorithm for this task on pure ARX designs. Our algorithm is heavily inspired by the pioneer

¹ The comparison is done with similar implementations.

work of de Cannière and Rechberger [DR06], and the more detailed explanation given in [Pey08] and [MDIP09].

3.1 Types of Trails

Differential trails can be classified in two categories: iterative and non-iterative. An iterative characteristic exploits the round-based nature of many cryptographic constructions: if a trail can be built over a few rounds with the same input and output difference Δ , then this characteristic can be repeated to reach a larger number of rounds. In practice very few iterative characteristics have been found for ARX constructions, because many designs use different rotation amounts or Boolean functions over the rounds, or a non-iterative key-schedule. Notable exceptions include the attacks of den Boer and Bosselaers against MD5 [dBB93], and the recent work of Dunkelman and Khovratovich on BLAKE [DK11]. In this work, we focus on non-iterative trails.

The main way to build non-iterated trails is to connect two simple and high-probability trails using a complex and low-probability section in between. The choice of the high-probability trails will depend on the attack setting, and should be done by the cryptanalyst using specific properties of the design, while the complex section will be build automatically by an algorithm (or by hand). When the characteristic is used in a hash-function attack, the cost of the low-probability section can usually be avoided.

For instance, the characteristics used for the attacks on SHA-1 use a linear section built using local collisions [CJ98,WYY05], and a non-linear section to connect a given input difference to the linear characteristic. This general idea is also the core of the rebound attack [MRST09]: it combines two high-probability trails using a low-probability transition through an S-box layer.

In our applications, we will use a rebound-like approach to connect two high-probability trails with a complex low-probability section.

3.2 Algorithm

Our algorithm takes as input a characteristic representing two high-probability trails $\Delta_1 \rightarrow \Delta_2$ and $\Delta_3 \rightarrow \Delta_4$. The middle section is initially unconstrained, *i.e.* filled with ?. The main part of the algorithm is a search phase which tries to fill the middle part with a valid characteristic. We follow the general idea of the algorithm of de Cannière and Rechberger, by repeating the following operations, as illustrated in Figure 1:

Propagation: deduce more information from the current characteristic by running the propagation algorithm on each operation.

Guessing: select an unconstrained state bit (*i.e.* a ? constraint), and reduce the set of allowed values (*e.g.* to a - or x constraint).

When a contradiction is found, we go back to the last guess, and make the opposite choice, leading to a backtracking algorithm. However, we abort after some number of trials and restart from scratch because mistakes in the early guesses would never be corrected.

Our algorithm is built from the idea that the constraint propagation is relatively efficient to check if a transition $\Delta \rightarrow \Delta'$ is possible. Therefore to connect the differences Δ_2 and Δ_3 from the high-probability trails, we essentially guess the middle difference Δ' and we check whether the transitions $\Delta_2 \rightarrow \Delta'$ and $\Delta' \rightarrow \Delta_3$ are possible.

This leads to the following difference with the algorithm of de Cannière and Rechberger:

- We specify in advance which words of the state will be restricted in the guessing phase, using state words in the middle of the unspecified section.
- We guess from the low bits to the high bits, and we can compare incomplete characteristics by counting how many bits have been guessed before aborting the search.
- Every time the backtracking process is aborted, we remember which guess was best and the random guesses of the next run are biased toward this choice.
- We only use signed differences, *i.e.* we use the constraints -, u, and n.

3.3 Finding pairs

The hardest part of our attacks is to build the differential trails. Finding conforming pairs for the middle section is relatively easy using the propagation algorithm: one just has to make random choices for the unconstrained bits in the middle and run the propagation algorithm after each choice. In practice the paths we found leave very few choices to make, and most of them lead to valid pairs. The degrees of freedom in the key can then be used to build many different pairs. This can be compared to the rebound attack on AES-like designs [MRST09]: in this attack the trails are easy to build, and finding pairs for the inbound phase has a small amortized cost.

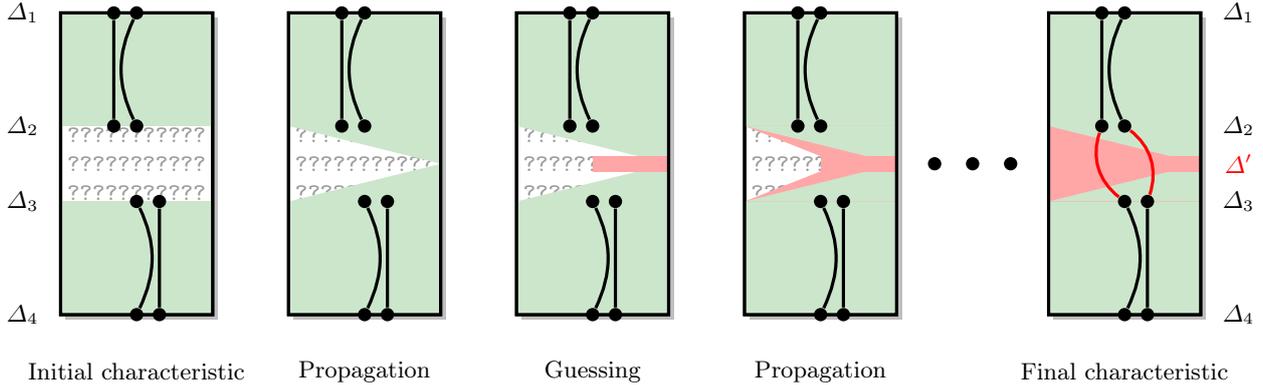


Fig. 1. Overview of the search algorithm. We start with high-probability trails $\Delta_1 \rightarrow \Delta_2$ and $\Delta_3 \rightarrow \Delta_4$, and we connect them through a difference Δ'

4 Application to Skein-256

In this section, we apply our algorithm to build characteristics for several attack scenarios on Skein-256.

4.1 Short Description of Threefish and Skein

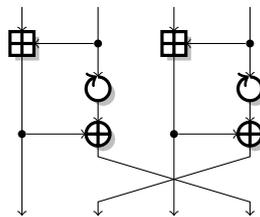


Fig. 2. Threefish-256 round

The compression function of Skein is based on the block cipher Threefish. In this paper we only study Threefish-256, which uses a 256-bit key (as 4 64-bit values), a 128-bit tweak (as 2 64-bit values), and a 256-bit state (as 4 64-bit values). The full version of Skein has 72 rounds. We denote the i th word of the state after r rounds as $e_{r,i}$. There is a key addition layer every 4 rounds:

$$e_{r,i} = \begin{cases} v_{r,i} + k_{r/4,i} & \text{if } r \bmod 4 = 0 \\ v_{r,i} & \text{otherwise} \end{cases}$$

where $k_{r/4,i}$ is the i th word of the round key at round $r/4$. The round function is shown by Figure 2. The state $v_{r+1,i}$ (for $i = 0, 1, \dots, n_w$) after round $r + 1$ is obtained from $e_{r,i}$ by applying a MIX transformation and a permutation of 4 words as following:

$$\begin{aligned} (f_{r,2j}, f_{r,2j+1}) &:= \text{MIX}_{r,j}(e_{r,2j}, e_{r,2j+1}) && \text{for } j = 0, 1, \dots, n_w/2 \\ v_{r+1,i} &:= f_{r,\sigma(i)} && \text{for } i = 0, 1, \dots, n_w \end{aligned}$$

where σ is the permutation (0 3 2 1) (specified in [FLS⁺10]) and $(c, d) = \text{MIX}_{r,j}(a, b)$ is defined as:

$$\begin{aligned} c &= a \boxplus b \\ d &= (b \lll R_{r \bmod 8, j}) \oplus c \end{aligned}$$

The rotations $R_{r \bmod 8, j}$ are specified in [FLS⁺10]. The key scheduling algorithm of Threefish produces the round keys from a tweak (t_0, t_1) and a key as following:

$$\begin{aligned} k_{l,0} &= k_{(l-1) \bmod 5} && k_{l,1} = k_{(l+1) \bmod 5} + t_{l \bmod 3} \\ k_{l,2} &= k_{(l+2) \bmod 5} + t_{(l+1) \bmod 3} && k_{l,3} = k_{(l+3) \bmod 5} + l, \end{aligned}$$

where $k_4 = C_{240} \oplus \bigoplus_{i=0}^4 k_i$ with C_{240} a constant specified in [FLS⁺10], and $t_2 = t_0 \oplus t_1$. The compression function F for Skein is given as $F(M, H, T) = E_{H,T}(M) \oplus M$, where H is the chaining value, M is the message, and T is a block counter. This follows the Matyas-Meyer-Oseas construction for the compression function, and the Haifa construction for the iteration.

In this work, we only consider attacks on multiples of four rounds, because the structure of Skein is built with 4-round blocks with key additions in between. We describe attacks in three different settings in Sections 4.4, 4.5, and 4.6. The attacks are based on different kinds of trails shown in Figures 4, 5, and 6, and examples of characteristics are given in Tables 11, 13, and 14, respectively. All the characteristics have been verified by building a conforming pair.

4.2 Building Characteristics

To describe a differential characteristic for Skein with our framework, we write constraints for each $e_{r,i}$ value, and for the $v_{r,i}$ values before a key addition (*i.e.* when $r \bmod 4 = 0$). For each round, we have 4 equations and 2 rotations, corresponding to two MIX functions. We also write the full key schedule as a system of equations.

We note that the variables $e_{r,2j}$ with $r \bmod 4 = 0$ are only involved in modular additions: $f_{r,2j} = e_{r,2j} \boxplus e_{r,2j+1}$ and $e_{r,2j} = v_{r,2j} \boxplus k_{r/4,2j}$. Therefore, we could remove these variables, and write $f_{r,2j} = v_{r,2j} \boxplus k_{r/4,2j} \boxplus e_{r,2j+1}$ using a three-input modular addition. In practice, the propagation algorithm for three-input modular addition takes significantly longer, so we keep the variables, but we try to avoid constraining them since the multi-bit constraints can propagate the modular difference.

Choosing the high-probability characteristics. In attack setting with differences in the key, we build the high-probability trails starting from a non-active state, with a low-weight key difference. When we go through the key addition, a difference is introduced in the state, and we propagate the difference by linearizing the function. If we have no difference in the key, we start with a single active bit in the state and we propagate the difference for a few rounds by linearizing the function. Most of our trails use the most significant bit as the active bit in order to increase their probabilities.

4.3 General Results

For the algorithm to work successfully, we need to find a delicate balance in the initial characteristic. If the unconstrained section is too short, there will not be enough degrees of freedom to connect the high-probability parts. On the other hand, if the unconstrained section is too long, the propagation algorithm will not filter bad characteristics efficiently.

In practice, we can only build characteristics when we have a key addition layer in the unconstrained part of the characteristic. This way, the algorithm can use degrees of freedom from the key to connect the initial characteristics. In general it seems hard to find enough degrees of freedom to build a valid trail without using degrees of freedom from the key: for a random function f and arbitrary differences Δ_2 and Δ_3 , we expect on average a single pair satisfying $f(x + \Delta_2) = f(x) + \Delta_3$. We can consider the intermediate differences for one such pair as a differential characteristics but a differential characteristic with a single valid pair is not very useful for a differential attack.

In order to let the algorithm use the degrees of freedom in the key efficiently, we use the registers before and after a key addition as guessing points: $v_{r,0}, v_{r,1}, v_{r,2}, v_{r,3}, e_{r,1}, e_{r,3}$ with $r \bmod 4 = 0$ (as discussed above we do not constrain $e_{r,0}$ and $e_{r,2}$).

We find that the characteristics built by the algorithm are rather dense, and use all the degrees of freedom in the state, and many degrees of freedom in the key. This is not a problem for attacks on the compression function, but the characteristics are harder to use in attacks against the full hash function, where fewer degrees of freedom are available to the attacker. We note that this problem is less acute for attack against functions of the MD/SHA family, where the message block is much larger than the state.

On the other hand, the trail built by hand by Yu *et al.* [YCW13] is somewhat sparser, and still leaves many degrees of freedom for the key and the middle state. More work will be needed to find such trails automatically.

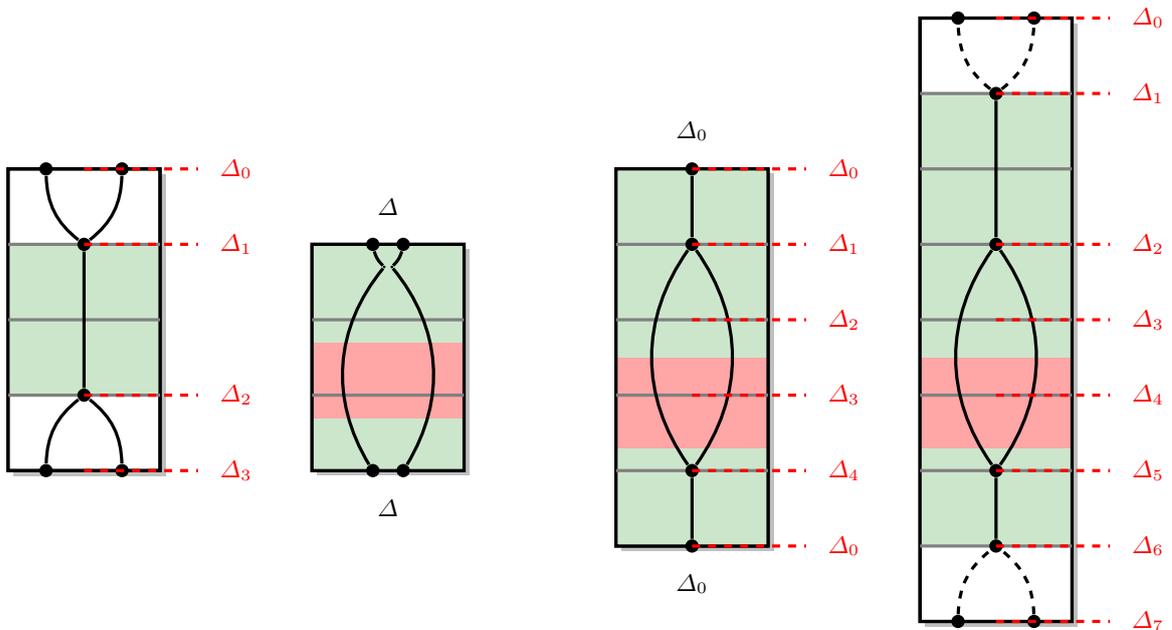


Fig. 3. Previous trails: rel-key, rel-tweak.

Fig. 4. Collision trails: fixed key.

Fig. 5. Collision trails: related-key.

Fig. 6. (Near-)Collision trails: rel-key, rel-tweak.

4.4 Collision Attacks

We first study attacks with no difference in the key (*i.e.* the chaining value) so that they can be applied to the full hash function. We try to build characteristics for a collision attack, therefore we use the same difference in the initial state and in the final state so that they can cancel out in the feed-forward². We start with a low-weight difference in one of the first rounds and we propagate by linearization through rounds 0–4 and backward through round 11.

We show an example of such characteristic in Table 11. This characteristic can be used for a practical semi-free-start collision attack on 12-round Skein, and we give an example in Table 6. We also built a similar characteristic for rounds 4–16 of Skein, in Table 12, and a collision example in Table 7.

Full collision attack. To build a collision attack on the full hash function, we have to deal with the fact that the characteristic is only valid for a small fraction of the keys (*i.e.* a small fraction of the chaining values). We use a large number of characteristics, and a large number of random chaining values, in a meet-in-the-middle fashion.

Our experiments indicate that we can build characteristics with about 2^{70} solutions for a cost of 2^{40} . If we extrapolate this experimental result, we expect that it is possible to build many such characteristics. Let’s assume that we can build N characteristics for a cost of $N \times 2^{40}$; each characteristic has 2^{70} solutions out of 2^{150} valid keys. In a second phase, we will hash M random message blocks and test if they can give a collision using one of the characteristics. Out of the M chaining values generated, we expect that $M \times N \times 2^{150-256}$ will be valid for one characteristic, and $M \times N \times 2^{70-256}$ values will actually lead to a collision after verification. An important step of the attack will be to find for which characteristic a given chaining value can be valid, but this can be done efficiently using a hash table indexed by the bits of the chaining value which are imposed by the characteristics.

The optimal complexity is achieved with $N = 2^{73}$ and $M = 2^{113}$. With these parameters we only have to verify 2^{80} valid chaining values, so the verification step is negligible. This gives a collision attack on 12-round Skein-256 with a time complexity of 2^{114} , using memory to store 2^{73} characteristics³.

The assumption that we can build so many good characteristics is a strong assumption, and it is hard to verify. However, we believe that this estimation is a safe upper bound, and that better characteristics would be found by running the search algorithm for longer times.

In our experiments, we tested a few different high probability trails as input to the algorithms, and we spend an effort equivalent to about 2^{43} hash computations on our best candidate (corresponding to the path of Table 11). We have found more than 200 different characteristics; the best characteristic allow 2^{75} solutions, and several of them allow more than 2^{70} solutions. We also checked that the subset of valid keys for the best characteristic are distinct. In order to build a large number of characteristics, we would also use several different starting points.

4.5 Free-start Collision Attack

For a collision attack on the compression function, *i.e.* a free-start attack on the hash function, we can use a difference in the key (*i.e.* the chaining value). We note that the key schedule of Skein-256 repeats itself every 20 rounds when there is no tweak difference. Therefore, we build trails with two inactive blocks as shown in Figure 5: the difference introduced in the initial state by k_0 cancels out with the difference introduced in the final state by k_5 .

² We could build characteristics for 20 rounds if we consider near-collisions, but this would not work on the full hash function because of the finalization step.

³ To store a characteristic, we just need to store masks defining the valid keys, and one state in the middle (if is not necessary to store all the intermediate constraints). Then, we can test a chaining value candidate by just computing all the intermediate states and checking if we reach a collision. This would take about 4×256 bits.

We give a characteristic built using this idea in Table 13, and a collision pair in Table 8.

We can also extend this path to a free-start near-collision attack against 24-round Skein, if we extend the trail to 4 more rounds at the end. A linearized trail gives near-collisions with 15 active bits, and the cost of finding a conforming pair is negligible before the cost of finding the trail.

4.6 Free-tweak Free-start Near-collision Attack

Finally, we can use degrees of freedom in the tweak to reach the maximum number of rounds possible. Previous works have shown that the key schedule allows to have one round without any active key words if we use a difference in the tweak in order to cancel a difference in the key. Using this property we can build a 24-round trail, and extend it to 32 round by propagating the external difference for four extra rounds in each direction, as shown in Figure 6. This is the approach used in [YCJW11].

We give a characteristic built using this idea in Table 14, and an example of pair following the characteristic in rounds 4 to 28 in Table 9. This results in a low weight difference for the input and output, with many zero bits in predetermined position. Moreover, we can follow the approach of [YCJW11] and also specify a fixed characteristic for round 0 to 4 and 28 to 32. It costs about 2^{40} to build a characteristic that allows 2^{20} solutions, so we can estimate that the amortized costs of building a valid pair for rounds 4 to 28 is about 2^{20} . Using the analysis of [YCJW11], we would build a conforming pair for rounds 0 to 32 for a cost of $2^{20+43+43} = 2^{119}$, assuming that we can find 2^{66} different characteristics.

Alternatively, if we can choose the value of the tweak, then we only need a single characteristic, and we follow the same attack as [YCW13].

Note that the complexity of these attack is higher than the generic complexity of a partial-collision attack on $256 - 51$ pre-specified bits, $2^{102.5}$. However, the generic complexity to reach the fixed 256-bit difference with 51 pre-specified active bits is still 2^{128} . Alternatively, this attack can be considered as a q -multicollision attack [BKN09].

Conclusion

In this paper we describe an algorithm to build differential characteristics for ARX designs, and we apply the algorithm to find characteristics for various attack scenarios on Skein. Our attacks do not threaten the security of Skein, but we achieve good results when compared to previous attacks; our main results are low-complexity attacks in relatively strong settings. In particular, we show practical free-start and semi-free-start collision attacks for 20 rounds and 12 rounds of Skein-256, respectively.

We obtain some of the first complex differential trails for pure ARX functions (as opposed to MD/SHA-like functions with Boolean functions). Since our approach is rather generic, we expect that our technique can be applied to other ARX designs, and will be used to evaluate the security of these designs against differential cryptanalysis.

Our improvements to the tools of [Leu12], and the code to build differential characteristics for Skein will be published together with the paper. A preliminary version of the code is available anonymously at the following url: <http://ubuntuone.com/1SWsnYUswpmfffVmySAox>. We hope that this will promote cooperation between researchers, and avoid the current situation for characteristic search for MD/SHA-like designs, where several teams had to develop their own implementation.

Acknowledgement

We would like to thank Pierre-Alain Fouque and Thomas Peyrin for fruitful discussions about differential characteristics and propagation of constraints.

Gaëtan Leurent was supported by the AFR grant PDR-10-022 of the FNR Luxembourg, and is supported by the CRASH ERC grant from the European Union.

References

- AB12. Jean-Philippe Aumasson and Daniel J. Bernstein. SipHash: a fast short-input PRF. In Steven Galbraith and Mridul Nandi, editors, *INDOCRYPT*, Lecture Notes in Computer Science. Springer, 2012.
- BKN09. Alex Biryukov, Dmitry Khovratovich, and Ivica Nikolic. Distinguisher and Related-Key Attack on the Full AES-256. In Halevi [Hal09], pages 231–249.
- Can12. Anne Canteaut, editor. *Fast Software Encryption - 19th International Workshop, FSE 2012, Washington, DC, USA, March 19-21, 2012. Revised Selected Papers*, volume 7549 of *Lecture Notes in Computer Science*. Springer, 2012.
- CJ98. Florent Chabaud and Antoine Joux. Differential Collisions in SHA-0. In Hugo Krawczyk, editor, *CRYPTO*, volume 1462 of *Lecture Notes in Computer Science*, pages 56–71. Springer, 1998.
- Cra05. Ronald Cramer, editor. *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*. Springer, 2005.
- dBB93. Bert den Boer and Antoon Bosselaers. Collisions for the Compressin Function of MD5. In Tor Helleseth, editor, *EUROCRYPT*, volume 765 of *Lecture Notes in Computer Science*, pages 293–304. Springer, 1993.
- DK11. Orr Dunkelman and Dmitry Khovratovich. Iterative differentials, symmetries, and message modification in BLAKE-256. In *ECRYPT2 Hash Workshop*, 2011.
- DR06. Christophe De Cannière and Christian Rechberger. Finding SHA-1 Characteristics: General Results and Applications. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT*, volume 4284 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2006.
- FLN07a. Pierre-Alain Fouque, Gaëtan Leurent, and Phong Nguyen. Automatic Search of Differential Path in MD4. In *ECRYPT Hash Workshop*, 2007. <http://eprint.iacr.org/2007/206>.
- FLN07b. Pierre-Alain Fouque, Gaëtan Leurent, and Phong Q. Nguyen. Full Key-Recovery Attacks on HMAC/NMAC-MD4 and NMAC-MD5. In Alfred Menezes, editor, *CRYPTO*, volume 4622 of *Lecture Notes in Computer Science*, pages 13–30. Springer, 2007.
- FLS⁺10. Niels Ferguson, Stefan Lucks, Bruce Schneier, Doug Whiting, Mihir Bellare, Tadayoshi Kohno, Jon Callas, and Jesse Walker. The Skein hash function family. Submission to NIST, 2008/2010.
- Hal09. Shai Halevi, editor. *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*. Springer, 2009.
- KRS12. Dmitry Khovratovich, Christian Rechberger, and Alexandra Savelieva. Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 Family. In Canteaut [Can12], pages 244–263.
- Leu12. Gaëtan Leurent. Analysis of Differential Attacks in ARX Constructions. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT*, volume 7658 of *Lecture Notes in Computer Science*, pages 226–243. Springer, 2012.
- LIS12. Ji Li, Takanori Isobe, and Kyoji Shibutani. Converting meet-in-the-middle preimage attack into pseudo collision attack: Application to sha-2. In Anne Canteaut, editor, *FSE*, volume 7549 of *Lecture Notes in Computer Science*, pages 264–286. Springer, 2012.
- MDIP09. Nicky Mouha, Christophe De Cannière, Sebastiaan Indestege, and Bart Preneel. Finding Collisions for a 45-Step Simplified HAS-V. In Heung Youl Youm and Moti Yung, editors, *WISA*, volume 5932 of *Lecture Notes in Computer Science*, pages 206–225. Springer, 2009.
- MNS11. Florian Mendel, Tomislav Nad, and Martin Schläffer. Finding SHA-2 Characteristics: Searching through a Minefield of Contradictions. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 288–307. Springer, 2011.
- MNS12. Florian Mendel, Tomislav Nad, and Martin Schläffer. Collision Attacks on the Reduced Dual-Stream Hash Function RIPEMD-128. In Canteaut [Can12], pages 226–243.
- MNS13. Florian Mendel, Tomislav Nad, and Martin Schläffer. Finding collisions for round-reduced sm3. In Ed Dawson, editor, *CT-RSA*, volume 7779 of *Lecture Notes in Computer Science*, pages 174–188. Springer, 2013.
- MRS09. Florian Mendel, Christian Rechberger, and Martin Schläffer. MD5 Is Weaker Than Weak: Attacks on Concatenated Combiners. In Mitsuru Matsui, editor, *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 144–161. Springer, 2009.
- MRST09. Florian Mendel, Christian Rechberger, Martin Schläffer, and Søren S. Thomsen. The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl. In Orr Dunkelman, editor, *FSE*, volume 5665 of *Lecture Notes in Computer Science*, pages 260–276. Springer, 2009.

- MVCP10. Nicky Mouha, Vesselin Velichkov, Christophe De Cannière, and Bart Preneel. The Differential Analysis of S-Functions. In Alex Biryukov, Guang Gong, and Douglas R. Stinson, editors, *Selected Areas in Cryptography*, volume 6544 of *Lecture Notes in Computer Science*, pages 36–56. Springer, 2010.
- Pey08. Thomas Peyrin. *Analyse de fonctions de hachage cryptographiques*. PhD thesis, University of Versailles, 2008.
- SLdW07. Marc Stevens, Arjen K. Lenstra, and Benne de Weger. Chosen-Prefix Collisions for MD5 and Colliding X.509 Certificates for Different Identities. In Moni Naor, editor, *EUROCRYPT*, volume 4515 of *Lecture Notes in Computer Science*, pages 1–22. Springer, 2007.
- SO06. Martin Schl affer and Elisabeth Oswald. Searching for Differential Paths in MD4. In Matthew J. B. Robshaw, editor, *FSE*, volume 4047 of *Lecture Notes in Computer Science*, pages 242–261. Springer, 2006.
- SSA⁺09. Marc Stevens, Alexander Sotirov, Jacob Appelbaum, Arjen K. Lenstra, David Molnar, Dag Arne Osvik, and Benne de Weger. Short Chosen-Prefix Collisions for MD5 and the Creation of a Rogue CA Certificate. In Halevi [Hal09], pages 55–69.
- SWWD10. Bozhan Su, Wenling Wu, Shuang Wu, and Le Dong. Near-Collisions on the Reduced-Round Compression Functions of Skein and BLAKE. In Swee-Huay Heng, Rebecca N. Wright, and Bok-Min Goi, editors, *CANS*, volume 6467 of *Lecture Notes in Computer Science*, pages 124–139. Springer, 2010.
- WLF⁺05. Xiaoyun Wang, Xuejia Lai, Dengguo Feng, Hui Chen, and Xiuyuan Yu. Cryptanalysis of the Hash Functions MD4 and RIPEMD. In Cramer [Cra05], pages 1–18.
- WY05. Xiaoyun Wang and Hongbo Yu. How to Break MD5 and Other Hash Functions. In Cramer [Cra05], pages 19–35.
- WYY05. Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding Collisions in the Full SHA-1. In Victor Shoup, editor, *CRYPTO*, volume 3621 of *Lecture Notes in Computer Science*, pages 17–36. Springer, 2005.
- YCJW11. Hongbo Yu, Jiazhe Chen, Keting Jia, and Xiaoyun Wang. Near-Collision Attack on the Step-Reduced Compression Function of Skein-256. IACR Cryptology ePrint Archive, Report 2011/148, 2011.
- YCW12. Hongbo Yu, Jiazhe Chen, and Xiaoyun Wang. The Boomerang Attacks on the Round-Reduced Skein-512. In Lars R. Knudsen and Huapeng Wu, editors, *SAC*. Springer, 2012.
- YCW13. Hongbo Yu, Jiazhe Chen, and Xiaoyun Wang. Partial-Collision Attack on the Round-Reduced Compression Function of Skein-256. In Shiho Moriai, editor, *FSE*, *Lecture Notes in Computer Science*. Springer, 2013.

A Improved Constraint Propagation

In this work we describe a method that is specific to systems of the following form:

$$\begin{aligned} u &= f(a, b, c, \dots) & u' &= f(a', b', c', \dots) \\ \delta(a, a') &= A & \delta(b, b') &= B & \delta(c, c') &= C & \dots \\ & & \delta(u, u') &= U, \end{aligned} \tag{2}$$

where f is an S-function, and the difference δ is given by a set of constraints which fully determines $x^{[i]}$, $x'^{[i]}$, $x^{[i-1]}$, and $x'^{[i-1]}$. We consider $a, a', b, b', \dots, u, u'$ as variables, and A, B, \dots, U as parameters.

Building the automaton. To deal with 2.5-bit constraints, we use a base alphabet \mathcal{B} of 32 constraints, each specifying one possible value for $x^{[i]}$, $x'^{[i]}$, $x^{[i-1]}$, $x'^{[i-1]}$, $x^{[i-2]}$ (for 2-bit constraints, the base alphabet has 16 constraints). Since the system given by (2) with the constraints in \mathcal{B} is an S-system, we can compute a set of *states* \mathcal{S} , and a transition function:

$$\begin{aligned} \tau : \quad \mathcal{S} \times (\mathcal{B} \times \{0, 1\} \times \{0, 1\})^{p-1} \times \mathcal{B} &\rightarrow \mathcal{S} \\ q, (\bar{A}, a, a'), (\bar{B}, b, b'), \dots, \bar{U} &\mapsto q' \end{aligned}$$

so that each solution to the system corresponds to a path in the automaton with transition function τ . More details about the construction of τ are given in [MVCP10, Leu12]. In our implementation, we use the tools of [Leu12] to compute the transition table.

When we describe a differential characteristic, we use an alphabet $\mathcal{A} = \mathcal{P}(\mathcal{B})$ consisting the 2^{32} subsets of the base alphabet \mathcal{A} (2^{16} subsets for 2-bit constraints). We transform an automaton on the alphabet \mathcal{B} to operate on the alphabet \mathcal{A} by changing the transition function into a non-deterministic transition function:

$$\begin{aligned} \tau' : \quad \mathcal{S} \times (\mathcal{A} \times \{0, 1\} \times \{0, 1\})^{p-1} \times \mathcal{A} &\rightarrow \mathcal{P}(\mathcal{S}) \\ q, (A, a, a'), (B, b, b'), \dots, U &\mapsto \bigcup_{\bar{A} \in A, \dots, \bar{U} \in U} \tau(q, (\bar{A}, a, a'), \dots, \bar{U}) \end{aligned}$$

This automaton can test whether the constraints are satisfied for given values of the parameters A, B, \dots, U , of the variables a, a', b, b', \dots , and with $u = f(a, b, c, \dots)$, $u' = f(a', b', c', \dots)$. We further transform the automaton by removing the information about a, a', \dots :

$$\begin{aligned} \tau'' : \quad \mathcal{S} \times \mathcal{A}^p &\rightarrow \mathcal{P}(\mathcal{S}) \\ q, A, B, \dots, U &\mapsto \bigcup_{a, a', b, b', \dots \in \{0, 1\}} \tau'(q, (A, a, a'), \dots, U) \end{aligned}$$

This new automaton can decide whether there exists solutions to System (2) for given parameters A, B, \dots, U . The transition function is highly non-deterministic, but we still use the original automaton by relabelling the transitions, and reading several transitions at each step.

Lemma 1. *The transition automaton of a system following (2) with p parameters, v variables, and s bits of state has the following properties:*

- i) Each state can be labelled with a $2v$ -bit value corresponding to $a, a', b, b', \dots, u, u'$. All the input transitions share this value for $a^{[i]}$, $a'^{[i]}$, $b^{[i]}$, $b'^{[i]}$ \dots , $x^{[i]}$, $x'^{[i]}$, while all the output transitions share this value for $a^{[i-1]}$, $a'^{[i-1]}$, $b^{[i-1]}$, $b'^{[i-1]}$ \dots , $x^{[i-1]}$, $x'^{[i-1]}$.*
- ii) No pair of states are linked by two different transitions;*
- iii) Each state has exactly 2^{2v} output transitions (the transition table is sparse);*

- Proof.* *i)* In order to reject incoherent constraints for bit $i - 1$ and i of a variable, the automaton must store the values of the previous bits that are used for the constraint on bit i in the state.
- ii)* Let's assume we have two transitions from a state q to a state q' . Since the two transition go to the same state, they must specify the same values of the parameters on bit i . Moreover, the two transition come from the same state, so they must also specify the same values on bits before i . Therefore the two transitions are the same.
- iii)* Because the system follows the form $x = f(a, b, c, \dots), x' = f(a', b', c', \dots)$, any choice of the variables a, a', b, b', \dots is valid with exactly one value of x, x' .

Propagation. We use the properties of Lemma 1 in order to build an efficient propagation algorithm. Thanks to property *ii)*, we have a one to one correspondence between the paths in the original automaton, and the paths in the relabelled automaton. Therefore we can easily identify the constraints corresponding to actual solutions of the system. To propagate constraints, we first build the set of paths allowed by the initial constraints, we look at which edges are actually used in paths, and we build the new constraints by identifying the constraints corresponding to the edges.

Notations. We use the symbols from [Leu12] to denote the most common constraints as shown in Table 5. When a characteristic uses a less common constraint, we use an hexadecimal mask to represent it. The less common constraints used in the characteristics given in Appendix D are given in Table 10.

When the constraints on the current bit and the constraints on previous bits are independent, we write the constraints involving previous bit in exponent (*e.g.* see Figure 7). For instance, we have can write the constraints $<$ as $\mathbf{u}^{\mathbf{u}} \cup \mathbf{n}^{\mathbf{n}}$.

A.1 Propagation for a Differential Characteristic

A differential characteristic is given by a set a constraints for each internal state variable. An ARX design (or a more general MD/SHA-like design) is built with two kinds of operations:

- Operations that are **S-functions**: additions, xors, and bitwise Boolean function. We build a system for each operation following (2), and we use them to propagate constraints between the inputs and the output of the operation (the propagation goes both ways). To propagate a full characteristic, we propagate every operation until no new constraints are found.
- **Rotations**: since the constraints are local and only involve consecutive bits, we deal with a rotation $y = x \ggg i$ by simply rotating the constraint pattern: if $\delta x = \Delta_x$ then we use $\delta y = \Delta_x \ggg i$. However, we have to relax some constraints if the multi-bit relations are broken by the rotation.

A.2 Propagation Example

Let us show how the propagation operates with a simple example. For this example, we use 2-bit constraints, and we consider the operation $u = a \vee (a \boxplus a)$. The leads to the following system:

$$\begin{aligned} u = a \vee (a \boxplus a) \quad u' = a' \vee (a' \boxplus a') \\ \delta(a, a') = A \quad \delta(u, u') = U. \end{aligned} \tag{3}$$

This system has 2 parameters, 2 variables and 4 bits of state (two for each δ operation; the state of $a \boxplus a$ is already included in the state of $\delta(a, a')$). The automaton corresponding to this system is given in Figure 7. Note that the automaton only needs 9 states out of the $2^4 = 16$ possible values for the state of the S-system. In our work we always minimize the automata, and this

usually results in a significant reduction of the number of states. We can verify that Lemma 1 is respected.

We will show how the propagation algorithm works with the following input:

$$\delta(a, a') = -\mathbf{x}-- \qquad \delta(u, u') = ----. \qquad (4)$$

This correspond to a situation where an input difference must be absorbed through the operation.

We first build a graph with a copy of the transitions for each bit. Then for each bit, we remove the transitions that are not acceptable according to the initial constraints (4). More precisely, we only keep constraints that are subsets of $-/-$ for the first and second bits, subsets of $\mathbf{x}/-$ for the third bit, and subsets of $-/-$ for the fourth bit. We get the graph of Figure 8, and we look for paths starting for state 0 in the initial layer, and ending in any state of the final layer. (Note that the least significant bit is on the left in the graph, but on the right when we write $\delta x = -\mathbf{x}--$). The nodes and edges involved in these paths are shown in black. We note that the constraints are compatible because such paths exists, and we can count the number of paths to compute the number of solutions: there are 4 different paths in the graph, so there are 4 different solutions to System (3) satisfying (4). We can read the solutions by following the paths:

$\delta(a, a')$:	$1\mathbf{n}10$	$1\mathbf{u}10$	$1\mathbf{n}11$	$1\mathbf{u}11$
$\delta(u, u')$:	1110	1110	1111	1111

Let us now do the constraint propagation. For each bit, we look at the active edges in Figure 8, and we list the corresponding constraints for a and u in Table 4. The new constraints will be the union of all the active constraints. We get the following output (we disregard restrictions on previous bits for bit 0):

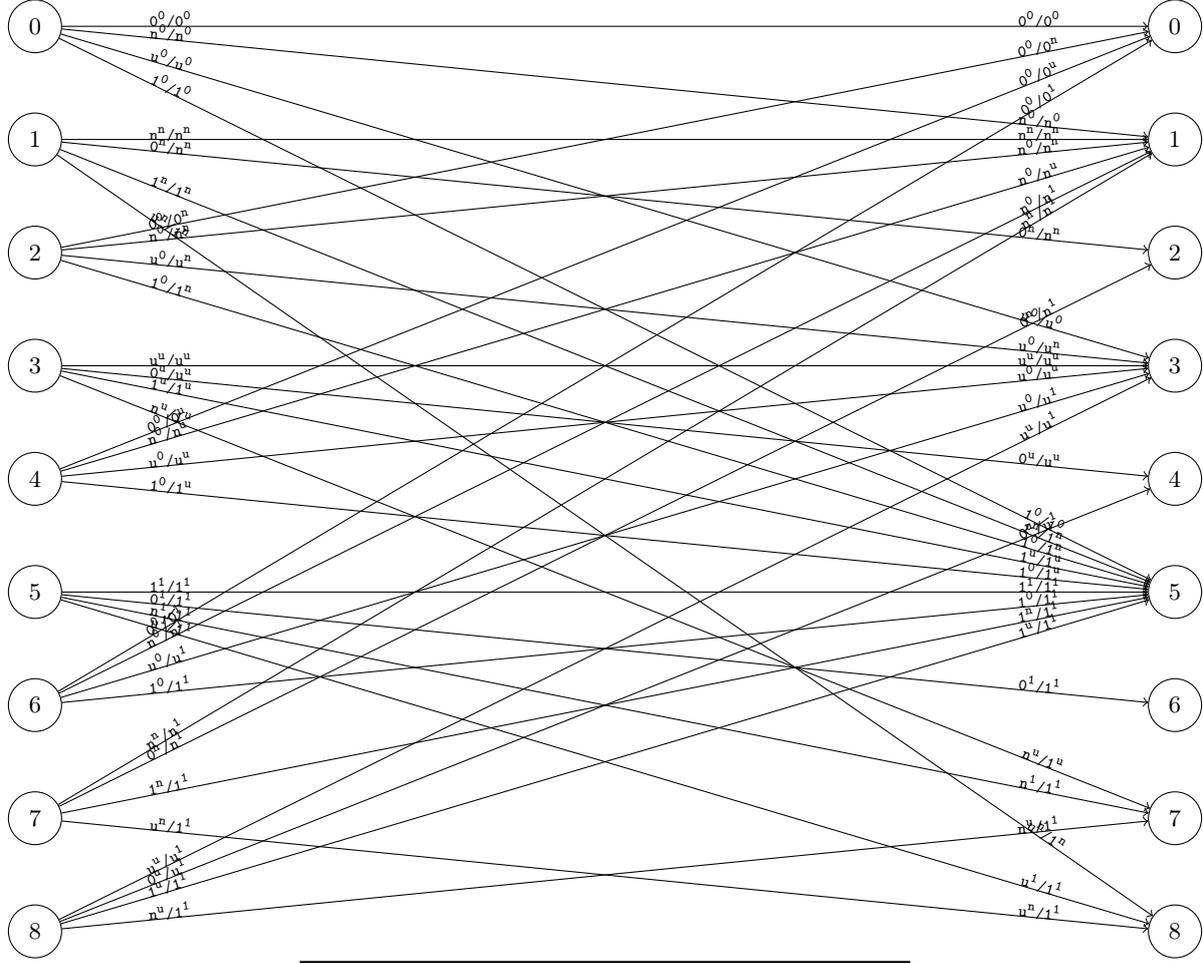
$$\delta(a, a') = 1^{\mathbf{x}}\mathbf{x}11^- \qquad \delta(u, u') = 1^11^11^-.$$

Here, the constraints on previous bits do not add any information, so we can omit them:

$$\delta(a, a') = 1\mathbf{x}1- \qquad \delta(u, u') = 111- \qquad (5)$$

It is easy to verify that any solution to the System (3) satisfying the initial constraints (4) also satisfies the deduced constraints.

B Figures



State	Transitions			
0	$0 \xrightarrow{0^0/0^0} 0$	$0 \xrightarrow{n^0/n^0} 1$	$0 \xrightarrow{u^0/u^0} 3$	$0 \xrightarrow{1^0/1^0} 5$
1	$1 \xrightarrow{n^n/n^n} 1$	$1 \xrightarrow{0^n/n^n} 2$	$1 \xrightarrow{1^n/1^n} 5$	$1 \xrightarrow{u^n/1^n} 8$
2	$2 \xrightarrow{0^0/0^n} 0$	$2 \xrightarrow{n^0/n^n} 1$	$2 \xrightarrow{u^0/u^n} 3$	$2 \xrightarrow{1^0/1^n} 5$
3	$3 \xrightarrow{u^u/u^u} 3$	$3 \xrightarrow{0^u/u^u} 4$	$3 \xrightarrow{1^u/1^u} 5$	$3 \xrightarrow{n^u/1^u} 7$
4	$4 \xrightarrow{0^0/0^u} 0$	$4 \xrightarrow{n^0/n^u} 1$	$4 \xrightarrow{u^0/u^u} 3$	$4 \xrightarrow{1^0/1^u} 5$
5	$5 \xrightarrow{1^1/1^1} 5$	$5 \xrightarrow{0^1/1^1} 6$	$5 \xrightarrow{n^1/1^1} 7$	$5 \xrightarrow{u^1/1^1} 8$
6	$6 \xrightarrow{0^0/0^1} 0$	$6 \xrightarrow{n^0/n^1} 1$	$6 \xrightarrow{u^0/u^1} 3$	$6 \xrightarrow{1^0/1^1} 5$
7	$7 \xrightarrow{n^n/n^1} 1$	$7 \xrightarrow{0^n/n^1} 2$	$7 \xrightarrow{1^n/1^1} 5$	$7 \xrightarrow{u^n/1^1} 8$
8	$8 \xrightarrow{u^u/u^1} 3$	$8 \xrightarrow{0^u/u^1} 4$	$8 \xrightarrow{1^u/1^1} 5$	$8 \xrightarrow{n^u/1^1} 7$

Fig. 7. Transitions for System (3)

Table 4. Active edges in figure 8, and new deduced constraints.

i	edges ($\delta a/\delta u$)	a constraints	u constraints
0	$0^0/0^0, 1^0/1^0$	$0^0 \cup 1^0 \equiv -^0$	$0^0 \cup 1^0 \equiv -^0$
1	$1^0/1^0, 1^1/1^1$	$1^0 \cup 1^1 \equiv 1^-$	$1^0 \cup 1^1 \equiv 1^-$
2	$n^1/1^1, u^1/1^1$	$n^1 \cup u^1 \equiv x^1$	$1^1 \cup 1^1 \equiv 1^1$
3	$1^n/1^1, 1^u/1^1$	$1^n \cup 1^u \equiv 1^x$	$1^1 \cup 1^1 \equiv 1^1$

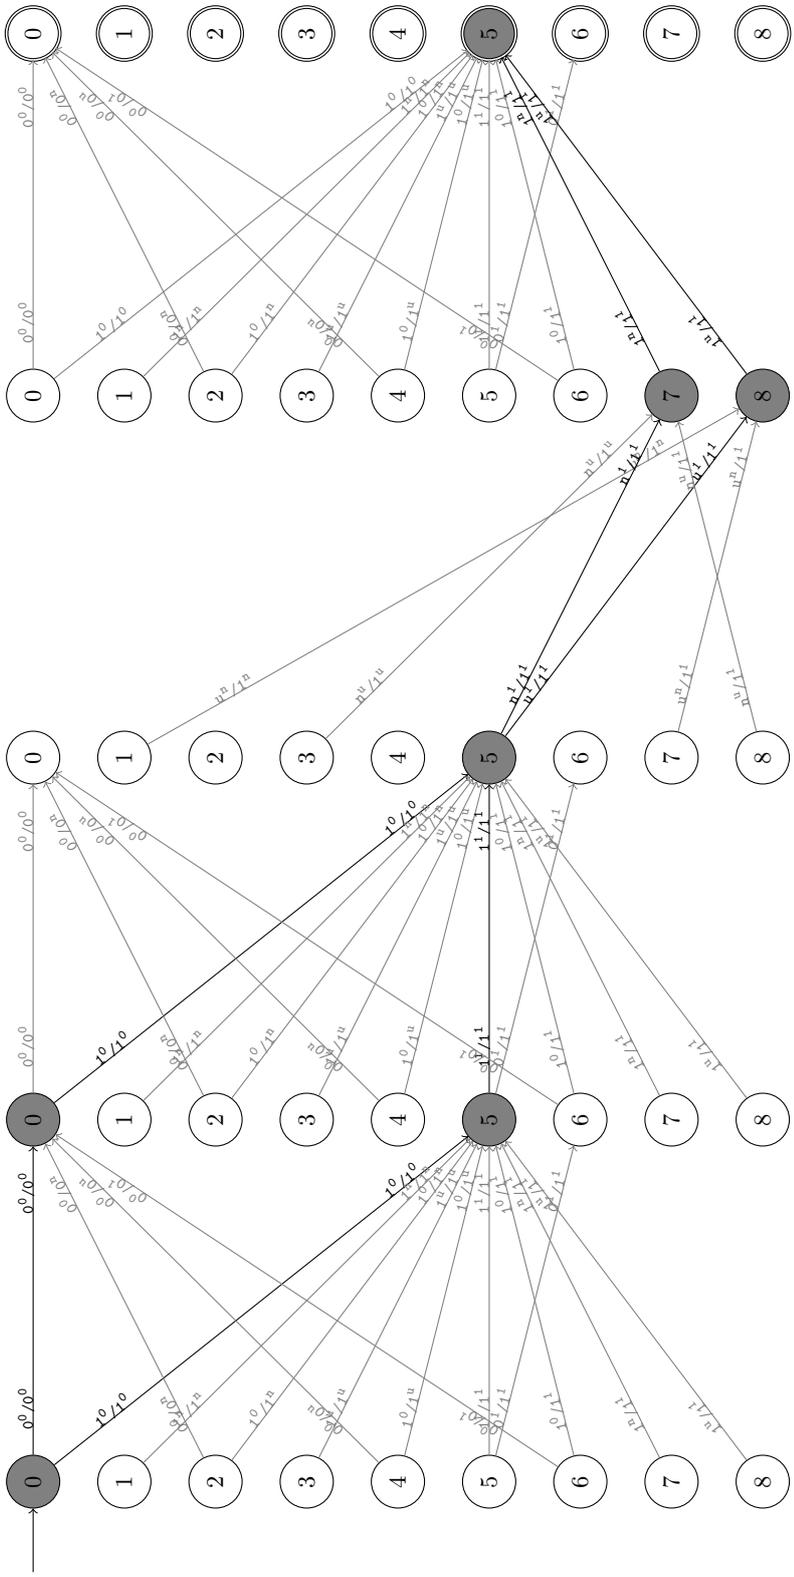


Fig. 8. Graph representation of System (3) with initial constraints (4)

$(x, x', 2x, 2x', 4x)$:

00000	00001	00010	00011	00100	00101	00110	00111	01000	01001	01010	01011	01100	01101	01110	01111	10000	10001	10010	10011	10100	10101	10110	10111	11000	11001	11010	11011	11100	11101	11110	11111						
?	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					
-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					
x	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					
u	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-					
n	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-				
1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-				
#	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-				
3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
7	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			
A	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-			
B	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
C	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
D	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
E	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
=	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
!	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
<	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	
>	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	
X	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
U	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
N	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
/	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
\	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	
V	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
M	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 5. Constraints identified in [Leu12] and written as full 2.5-bit constraints.

C (Near-)Collision Pairs for Skein-256

Table 6. Semi-free-start collision for 12-round Skein-256 (rounds 0–12). This pair is the same as given in Table 11.

	Input 1	Input 2	Difference	Weight
m_0	968cb2e66b0fb527	968cb2e66b0fb527	0000000000000000	0
m_1	37fce3361809b06a	37fce3361809b06a	0000000000000000	0
m_2	4bb032fb1894a60b	4bb032fb1894a60b	0000000000000000	0
m_3	d917aa4640682db6	d917aa4640682db6	0000000000000000	0
t_0	0000000000000000	0000000000000000	0000000000000000	0
t_1	0000000000000000	0000000000000000	0000000000000000	0
$v_{0,0}$	e7395021238d7d18	e3396021238d7d18	0400300000000000	3
$v_{0,1}$	7229b06628958c1a	7229a06628958c1a	0000100000000000	1
$v_{0,2}$	3ea410b0b8f1b533	3ea410b0b8f1b533	0000000000000000	0
$v_{0,3}$	fc0aa7147201f560	fc0aa7147201f560	0000000000000000	0
	Output 1	Output 2	Difference	Weight
$e_{12,0}$	2798a30c07459007	2398930c07459007	0400300000000000	3
$e_{12,1}$	2410f135e024aace	2410e135e024aace	0000100000000000	1
$e_{12,2}$	60490bbd9ddcb933	60490bbd9ddcb933	0000000000000000	0
$e_{12,3}$	7fd51384c7b528f3	7fd51384c7b528f3	0000000000000000	0
h_0	c0a1f32d24c8ed1f	c0a1f32d24c8ed1f	0000000000000000	0
h_1	56394153c8b126d4	56394153c8b126d4	0000000000000000	0
h_2	5eed1b0d252d0c00	5eed1b0d252d0c00	0000000000000000	0
h_3	83dfb490b5b4dd93	83dfb490b5b4dd93	0000000000000000	0

Table 7. Semi-free-start collision for 12-round Skein-256 (rounds 4–16). This pair is the same as given in Table 12.

	Input 1	Input 2	Difference	Weight
m_0	97c787b0252f1bef	97c787b0252f1bef	0000000000000000	0
m_1	9ba673bd9a918263	9ba673bd9a918263	0000000000000000	0
m_2	59f24b2909ae5223	59f24b2909ae5223	0000000000000000	0
m_3	963151773356523a	963151773356523a	0000000000000000	0
t_0	0000000000000000	0000000000000000	0000000000000000	0
t_1	0000000000000000	0000000000000000	0000000000000000	0
$v_{4,0}$	b9ded48b4e413597	39ded48b4e413597	8000000000000000	1
$v_{4,1}$	5a63d56d9481f1d6	5a63d56d9481f1d6	0000000000000000	0
$v_{4,2}$	0accb31ed067ae77	0accb31ed067ae77	0000000000000000	0
$v_{4,3}$	734e405bed9d64cc	734e405bed9d64cc	0000000000000000	0
	Output 1	Output 2	Difference	Weight
$e_{16,0}$	f3424f9d5f6d8c50	73424f9d5f6d8c50	8000000000000000	1
$e_{16,1}$	74a4ddb5e6e65d54	74a4ddb5e6e65d54	0000000000000000	0
$e_{16,2}$	bc4c51d904f3425d	bc4c51d904f3425d	0000000000000000	0
$e_{16,3}$	b511e49ca126be77	b511e49ca126be77	0000000000000000	0

Table 8. Free-start collision for 20-round Skein-256. This pair is the same as given in Table 13.

	Input 1	Input 2	Difference	Weight
m_0	5f977cfdd64d2f57	5f977cfdd64d2f57	0000000000000000	0
m_1	35839193022be6f4	b5839193022be6f4	8000000000000000	1
m_2	05e168930700458f	85e168930700458f	8000000000000000	1
m_3	6f47d57f8b6f9b78	6f47d57f8b6f9b78	0000000000000000	0
t_0	0000000000000000	0000000000000000	0000000000000000	0
t_1	0000000000000000	0000000000000000	0000000000000000	0
$v_{0,0}$	627f37f95152438c	627f37f95152438c	0000000000000000	0
$v_{0,1}$	0532b3fdf499d0d7	8532b3fdf499d0d7	8000000000000000	1
$v_{0,2}$	91c792ab31ba535c	11c792ab31ba535c	8000000000000000	1
$v_{0,3}$	72e80ac1aaee8118	72e80ac1aaee8118	0000000000000000	0

	Output 1	Output 2	Difference	Weight
$e_{20,0}$	6627a3d5c18e2057	6627a3d5c18e2057	0000000000000000	0
$e_{20,1}$	7a1e999992b2202d	fa1e999992b2202d	8000000000000000	1
$e_{20,2}$	2bf3a5067fac9218	abf3a5067fac9218	8000000000000000	1
$e_{20,3}$	b0ccc2f709dc2e35	b0ccc2f709dc2e35	0000000000000000	0

Table 9. Pair of input with low-weight difference for 32-round Skein-256. This pair is the same as given in Table 14; we don't specify how the differences are propagated in rounds 0 to 4 and 28 to 32.

	Input 1	Input 2	Difference	Weight
m_0	edb22ce30810011a	edb22ce30810011a	0000000000000000	0
m_1	08142e9044b0054a	08142e9044b0054a	0000000000000000	0
m_2	1e06bd5779535f97	1e06bd5779535f97	0000000000000000	0
m_3	82a5e785e5c5b836	02a5e785e5c5b836	8000000000000000	1
t_0	0000000000000000	8000000000000000	8000000000000000	1
t_1	0000000000000000	0000000000000000	0000000000000000	0
$v_{0,0}$	c0097c86ad089acd	c0dec29fae7a20d	00d7b7af57ef38c0	35
$v_{0,1}$	0eef94c587c9f8fc	91efc569f9eaf0fc	9f0051ac7e230800	23
$v_{0,2}$	a5333c6b7af97e18	a272f89740fdbae4	0741c4fc3a04c4fc	28
$v_{0,3}$	49df6d34f9ebc32f	cc9f6d0935eb8663	8540003dc00454c	19

	Output 1	Output 2	Difference	Weight
$e_{32,0}$	650f11ac87162f96	650f119c82f63796	0000003005e01800	9
$e_{32,1}$	22ed455a3e3dd26a	e5f12d8d8431cafa	c71c68d7ba0c1890	28
$e_{32,2}$	ef0d1179583e8671	ed0d118994327e51	020000f0cc0cf820	17
$e_{32,3}$	5de99dad57671f6a	5ec99dbd5347076a	0320001004201800	8

D Differential Characteristics for Skein-256

The characteristics given in Tables 11, 13, and 14 follow the general structure described in Figures 4, 5, and 6. For more details of the attacks, see Sections 4.4, 4.5, and 4.6, respectively.

We use the following colors in the characteristics:

- **red constraints** for active bits;
- **green constraints** for inactive bits;
- **orange constraints** for carry bits (inactive if the previous bit is inactive);
- **blue constraints** for other situations.

The most common characteristics are given in Table 5, while the unusual one are assigned a two digit code, and given in Table 10 in hexadecimal notation. The 32-bit hexadecimal values correspond to the columns of Table 5; for instance the constraint N would be represented by `f30c0cf3`. The two-digit codes are just used as shorthand so that all the information for the trails fit in the tables.

When using those characteristics, we start with the middle state given by the characteristic, we select a key satisfying the key constraints, and we check the remaining rounds. Therefore, the probabilities

given for the upper rounds are probability in the backward direction, while probabilities in the lower round are in the forward direction.

When the tweak is not given in the characteristic, it should be taken as zero.

Table 10. Description of the uncommon constraints used in the characteristics

Sym.	Mask	Sym.	Mask	Sym.	Mask	Sym.	Mask
00	00fff200	01	00f08f00	02	0001ff00		
10	0ff0f00f	11	f30c0000	12	cc30000f	13	f00f0000
14	0f0030cc	15	f00f00f0	16	cf0030cf	17	f00f0ff0
18	f0000c03	19	f2000cf3	1a	030c00f0	1b	f2000c01
1c	f3000cf3	1d	0f00f000	1e	002030cf	1f	c030000f
20	cf003000	21	000c0cf3	22	cf300000	23	00f0f00f
24	cf303000	25	8030004f	26	c330000f	27	0ff0100f
28	00f0000f	29	003030cf	2a	810c0cf3	2b	000030c0
2c	000f0ff0	2d	730c00f2	2e	cf3030cc	2f	f30c04e0
30	0000f00f	31	030c0000	32	0ff00000	33	0f102000
34	0400f008	35	0f00f00f	36	000030cf	37	cf3030c0
38	c0300000	39	00000cf3	3a	830c00f2	3b	030c0c03
3c	00000ff0	3d	0f0030c0	3e	c03030cf	3f	c03030c0
50	00cfc030	51	30cf0300	52	0c00f30c	53	0cf30004
54	0cf3000c	55	3001cf30	56	30c0c030	57	0cf3f300
58	00f3f30c	59	00f0030c	5a	000f00f0	5b	0cf30000
5c	30c04f00	5d	00f3000c	5e	0000c030	5f	30cf0000
60	0c03f200	61	00cf0030	62	00cf8030	63	0c030000
70	4f0000f2	71	8000004f	72	4f0000ff	73	4f0000f0
74	f00000ff	75	f200004f	76	010000ff	77	ff00004f
78	f20000ff	79	800000ff	7a	f2000001	7b	f200000f
7c	f000004f	7d	010000f2	7e	ff0000f0	7f	f10000ff
80	4f000080	81	8f0000ff	82	ff00000f	83	0f0000ff
84	4f0000f1	85	f000008f	86	81000042	87	020000ff
88	820000ff	89	8f000040	8a	8f0000f1	8b	ff0000f1
8c	810000ff	8d	ff00008f	8e	ff000080	8f	420000ff
90	40000080	91	ff000001	92	04000008	93	f200008f
94	ff000081	95	0f0000f1	96	ff000082	97	01000002
98	ff0000f2	99	ff000002	9a	0f0000f2	9b	f2000041
9c	ff000004	9d	ff000040	9e	4000008f	9f	080000ff
a0	ff000042	a1	f1000002	a2	410000ff	a3	280000ff
a4	f100000f	a5	82000041	a6	f100008f	a7	400000ff
a8	8f0000f2						

Table 11. Collision characteristic for rounds 0 to 12. $2^{158.1}$ valid keys and 2^{36} valid states, probability 2^{-119} .

	Constraints	Prob.	Example
k_0	1001-11!-0--- $\frac{7}{0}$ - $\frac{7}{1}$ - $\frac{7}{2}$ --- $\frac{0}{2}$ -0!-1001101!-!- $\frac{7}{3}$ - $\frac{7}{4}$ -1!- $\frac{7}{5}$ - $\frac{0}{6}$ - $\frac{7}{7}$ ---		968cb2e66b0fb527
k_1	----- $\frac{7}{8}$ -----!		37fce3361809b06a
k_2	---1! $\frac{7}{9}$ --- $\frac{7}{5}$ ---! $\frac{7}{a}$ - $\frac{7}{b}$ - $\frac{7}{c}$ - $\frac{7}{5}$ ---1!== $\frac{7}{3}$ - $\frac{7}{9}$ ---0-01 $\frac{7}{3}$ - $\frac{7}{d}$ -----1-		4bb032fb1894a60b
k_3	-1!-100--- $\frac{7}{3}$ ---1010!-0-0001---=000!=-01000-0101-0-1!-----!		d917aa4640682bdb6
k_4	--- $\frac{7}{e}$ - $\frac{7}{f}$ ---! $\frac{7}{6}$ - $\frac{7}{4}$ - $\frac{8}{8}$ ----- $\frac{8}{1}$ - $\frac{7}{f}$ - $\frac{7}{7}$ -----100 $\frac{7}{0}$ -0 $\frac{8}{8}$ ---		2806d2b7820694d2
$e_{0,0}$	---x-----x-----	2.0	7dc603078e9d323f
$e_{0,1}$	-----x-----	0.0	aa26939c409f3c84
$e_{0,2}$	-----	0.0	8a5443abd1865b3e
$e_{0,3}$	-----	0.0	d522515ab26a2316
$e_{1,0}$	---x-----	1.0	27ec96a3cf3c6ec3
$e_{1,1}$	-----	0.0	0e2c276ca0e6ab76
$e_{1,2}$	-----	0.0	5f76950683f07e54
$e_{1,3}$	-----	0.0	830b8684001d444a
$e_{2,0}$	---x-----	1.0	3618be1070231a39
$e_{2,1}$	-----	0.0	77840c878c0df816
$e_{2,2}$	-----	1.0	e2821b8a840dc29e
$e_{2,3}$	---x-----	0.0	81785cd206e91453
$e_{3,0}$	---x-----	1.0	ad9cca97fc31124f
$e_{3,1}$	---x-----n---1-----	0.0	8aee2bddf2aa04f7
$e_{3,2}$	---x-----	1.0	63fa785c8af6d6f1
$e_{3,3}$	---x-----	0.0	ee5acc6bf70ad049
$v_{4,0}$	-----0---n-----	1.8	388af675eedb1746
$v_{4,1}$	-----u----- $\frac{8}{2c}$ ----- $\frac{7}{e}$ ---	1.0	b30f4df549582a44
$v_{4,2}$	----- $\frac{8}{b}$ -----	0.0	525544c88201a73a
$v_{4,3}$	x-----n---u-----	2.0	654f8dcbbb989b7
$e_{4,0}$	----- $\frac{8}{d}$ -----n $\frac{8}{e}$ ----- $\frac{7}{6}$ $\frac{7}{7}$ $\frac{7}{7}$ ---	1.0	7087d9ac06e4c7b0
$e_{4,1}$	-----11--- $\frac{7}{6}$ -----0 $\frac{7}{e}$ ---u-----1!! $\frac{8}{5}$ $\frac{7}{6}$ -0-1---= $\frac{8}{f}$	0.0	febf80f061ecd04f
$e_{4,2}$	-----	0.0	2b6cef0ec269d4f0
$e_{4,3}$	x----- $\frac{7}{7}$ ---n- $\frac{8}{2}$ -u-----1 $\frac{8}{3}$ -1 $\frac{7}{7}$ --- $\frac{7}{7}$ ----- $\frac{8}{0}$ $\frac{7}{a}$ $\frac{8}{7}$ $\frac{4}{1}$ -1-	0.0	8d5660833da21e8a
$e_{5,0}$	$\frac{7}{e}$ ---1111-!- $\frac{7}{b}$ $\frac{7}{2}$ - $\frac{7}{c}$ --- $\frac{7}{1}$ -01001n $\frac{7}{e}$ -u---= $\frac{9}{1}$ -10!-01-001! $\frac{7}{1}$ - $\frac{7}{9}$ -----1	1.0	6f475a9c68d197ff
$e_{5,1}$	n---00111000 $\frac{8}{2}$ -11 $\frac{7}{7}$ ---00-0n0000n-n0---01-010-11-0110!-u-11-n-0	1.1	c38772871aa7327c
$e_{5,2}$	x--11 $\frac{7}{e}$ -0 $\frac{7}{2}$ ---0 $\frac{8}{0}$ ---1 $\frac{8}{e}$ --- $\frac{8}{2}$ n $\frac{8}{7}$ $\frac{8}{7}$ $\frac{4}{1}$ -u--- $\frac{9}{3}$ = $\frac{9}{1}$ - $\frac{7}{9}$ $\frac{4}{1}$ --- $\frac{7}{9}$ ---1! $\frac{7}{1}$ - $\frac{8}{1}$ ---	1.6	b8c34f92000bf37a
$e_{5,3}$	---01 $\frac{8}{3}$ -n $\frac{7}{6}$ ---01-01000-!-1= $\frac{9}{5}$ $\frac{8}{3}$ -n=n-11011!= $\frac{9}{8}$ -1---10-001111-1 $\frac{8}{9}$ ---	0.0	8f84833cf72ce8fe
$e_{6,0}$	u0-1001!-1001 $\frac{9}{1}$ -110!=-0nu0nuu0 $\frac{8}{3}$ -n $\frac{0}{8}$ ---011-11110001100101u0--1n011	2.0	32cecd238378ca7b
$e_{6,1}$	u0-0uuuu-111010-0u01n101n01111-0u0-100110110111001101001--u0000	0.0	00741dbc39b73480
$e_{6,2}$	u100nuuu0!-0!- $\frac{7}{a}$ -1!-0-10n10!== $\frac{7}{4}$ -n= $\frac{7}{1}$ -11-111000-1 $\frac{7}{3}$ ---001-1 $\frac{7}{6}$ ---	1.0	4847d2cef738dc78
$e_{6,3}$	n11111100n-100u--1n-1-0nn1uuu010u1011-1111u1100n-u0---u11-1u0--	1.0	fe51fdc25fd90cd2
$e_{7,0}$	0unn001n-10000-unn10n-nu1101n11-1011-101-01unnnnnnnnnnn0nnnn011	6.0	3342eadfb2dfefb
$e_{7,1}$	0un1un1uuuu011n-0unu0nn011-1u-1-0n10-0n-0101110un1n-1umu--n0101	3.0	360e26d263ae7d35
$e_{7,2}$	010001nnuu0110u11nnuuuu01u0nu0 $\frac{7}{0}$ ---101-111--u10001nn1010un010un-1-	1.4	4699d0915711e94a
$e_{7,3}$	-0nn00nn01uuuu1nuu11n0n010n0u-010011u1--0unuunuuuuuu0uuun00-	5.0	33433aa94dc92229
$v_{8,0}$	unnu1uun-nunuuu-0uu10-un1!-10-nuuunuuuu011011110un1n-1u0u--nu000		695111b220de7c30
$v_{8,1}$	-0nnunu0--0nu1000un0n0u1u001uun--0un-111nuu1nuu0uu110uu11nunnn-n-		34142913979831da
$v_{8,2}$	-1nnnu0111011nnuu00u1u110unn1un--010-1uunn0n10nn0000nu110n1n0-1-		79dd0b3aa4db0b73
$v_{8,3}$	u0uu-u1u-nnnnnn-u1nu--00n--u0-nnu00n01nunnu100u-u1u1-u10n--u0-1-		0aff6c8716d05ae2
$e_{8,0}$	NV $\frac{1}{0}$ xun $\frac{1}{0}$ UMVU $\frac{1}{2}$ UMNV $\frac{1}{6}$ x001 $\frac{1}{3}$ x1-1!-1 $\frac{1}{4}$ U $\frac{1}{0}$ x-1nuu-0--1---!-unu $\frac{1}{5}$ xn-u $\frac{1}{6}$ U $\frac{1}{0}$ x-1-	7.7	b50144ad3973223b
$e_{8,1}$	uuu-nnu1---n---unuun-u-unn0u---unnnuuuu-uuuuu-nunn-n-nu---uu0	6.3	0d2bd359d8005f90
$e_{8,2}$	UM $\frac{1}{7}$ $\frac{1}{0}$ x-1 $\frac{8}{3}$ -----n- $\frac{5}{2}$ NV $\frac{1}{5}$ VMN $\frac{1}{8}$ NVMN $\frac{1}{7}$ >-00 $\frac{1}{ab}$ N $\frac{1}{7}$ x $\frac{1}{ac}$ VMVUM $\frac{1}{7}$ x10 $\frac{1}{de}$ MV $\frac{1}{f}$ $\frac{1}{4}$ V $\frac{1}{0}$ x---	4.4	a1e3ddf226e1a045
$e_{8,3}$	nunu0uunn0---u-u--n--nu--u1nun1uuu00-nn-n--u-u0u1--uuu--u---1	9.3	a18c1f6d81e0100b
$e_{9,0}$	---uuuu0-n-nnu1u--1n0uuu00u-nnnuu-1uuu-0-n1---nnuuuu-unnn---n-	16.4	c22d1807117381cb
$e_{9,1}$	---n-!-u0-00u010u111nn-01 $\frac{9}{1}$ ---n-101n-u01-0 $\frac{7}{e}$ ---00001u0011n---u-	2.8	5c027cbfb8ca11dc
$e_{9,2}$	x--0--1u--unnnnnnnnnn-nu-u--1-n0nu-0un--u--1nunuuuuun--0	19.7	436ffd5fa8c1b050
$e_{9,3}$	---0110n-1-10110---1n100uu0-111000u011-n-0n---11u000-101u---0-	1.0	36fb6e0706978281
$e_{10,0}$	---n-0 $\frac{7}{e}$ -n-111n1--n01-0n $\frac{8}{3}$ =-u-!-1n-0-- $\frac{8}{0}$ ---n $\frac{7}{8}$ ---n10010-1110--- $\frac{7}{4}$	3.0	1e2f94c6ca3d93a7
$e_{10,1}$	--- $\frac{7}{7}$ --- $\frac{7}{4}$ -011u--- $\frac{8}{1}$!-1 $\frac{7}{8}$ ---n-!-u1---!---u-00---0111---	0.8	78069dbaa1541dd4
$e_{10,2}$	x----- $\frac{7}{7}$ ---10110--- $\frac{9}{1}$!-10-!-0-!-n01 $\frac{7}{3}$ ---0---100n10-10n1---	3.3	7a6b6b66af5932d1
$e_{10,3}$	x---0---1---0---10-01 $\frac{7}{4}$!-0-!-u0-0---0---0---u11-1 $\frac{7}{4}$ ---	0.0	03ea54e101c61f06
$e_{11,0}$	---n---n-01! $\frac{8}{2}$ ---n $\frac{7}{a}$ $\frac{9}{4}$ -n--- $\frac{7}{5}$ - $\frac{8}{2}$!- $\frac{7}{b}$ ---=0 $\frac{9}{a}$ -0101---	3.8	963632816b91b17b
$e_{11,1}$	---u-----u----- $\frac{8}{3}$ - $\frac{8}{6}$ ---1-----	1.0	b84ac6445b4b0d6
$e_{11,2}$	---0--- $\frac{7}{4}$ --- $\frac{7}{f}$ $\frac{7}{5}$ --- $\frac{9}{b}$ - $\frac{7}{4}$ - $\frac{7}{4}$ -1 $\frac{9}{a}$ ----- $\frac{7}{1}$ -0! $\frac{7}{7}$ - $\frac{9}{c}$ n---	1.6	7e55c047b11f51d7
$e_{11,3}$	---0---1-----1---1-----1 $\frac{7}{4}$ -----10u-----	0.0	4b66988f81adb235
$v_{12,0}$	---n--- $\frac{7}{8}$ ---n----- $\frac{8}{2}$ ---	2.0	4e80f8c5c6dd6251
$v_{12,1}$	-----x-----	0.0	fc0a1e7e5e1e15fc
$v_{12,2}$	----- $\frac{9}{4}$ ----- $\frac{8}{b}$ -----	0.0	c9bc58d732cd040c
$v_{12,3}$	-----1-----1-----	0.0	47d8304eafab7886

Table 13. Free-start collision characteristic for rounds 0 to 20. $2^{56.7}$ valid keys, probability 2^{-43}
 The characteristic shows rounds 4–16; rounds 0–4 and 16–20 are inactive.

	Constraints	Prob.	Example
k_0	-1-1-1111001 ⁸ _c -----111 ⁸ ₀ --1111!-- ⁸ ₀ -010-1001001-0 ⁷ _a ----111!--01!==		5f977cfdd64d2f57
k_1	x0-1010110000!-110010001100100110-000-1000101!- ⁹ ₁ ----011 ⁷ _d --110!=-		35839193022be6f4
k_2	x0-0-101-=-0-----010 ⁸ ₀ --0010-----000-== ⁷ ₈ ---0- ⁸ ₁ -----010---00 ⁷ ₂ ---		05e168930700458f
k_3	-11011110!-0!== ⁷ ₅ ---1010101111-1110001-110 ⁹ _a ---11111-0110110111 ⁷ ₂ --0		6f47d57f8b6f9b78
k_4	00011011!-000110100101101011-00111100!== ⁹ ₁ ---0101000011010111 ⁸ ₈ --0		1b634b58f1f50d76
$e_{4,0}$	x-----	0.0	dd5113862e4682f2
$e_{4,1}$	x-----	0.0	976b12a915df1438
$e_{4,2}$	-----	0.0	2682e7ab2b50853e
$e_{4,3}$	-----	0.0	c21caeeb08ac00af
$e_{5,0}$	-----	0.0	74bc262f4425972a
$e_{5,1}$	-----	0.0	f9c797c9b7c5d83b
$e_{5,2}$	-----	0.0	e89f969633fc85ed
$e_{5,3}$	-----n-----	0.0	26979807350b410f
$e_{6,0}$	-----	1.0	6e83bdf8fbeb6f65
$e_{6,1}$	----- ⁷ ₈ ⁸ ₃ ---11---u--- ⁹ _a ⁸ ₆ ---x--- ⁷ ₀ ---!	0.0	76b75dcddd173495
$e_{6,2}$	-----! ⁷ ₇ --- ⁹ _a -1---n-----	1.3	0f372e9d6907c6fc
$e_{6,3}$	-----== ⁸ _c -10-----!--- ⁸ _d --- ² _a --- ⁷ ₉ -!== ⁷ ₂ -----10 ⁷ _e ---	0.7	188d43891e190294
$e_{7,0}$	----- ⁹ ₄ ----- ⁷ ₂ ⁷ _e ---0u---1--- ³ _a ⁰ _{b-----⁷₀---=⁷₄---}	2.5	e53b1bc6d902a3fa
$e_{7,1}$	----- ⁹ ₃ ^a ₄ ⁹ ₇ ----- ² _a -----!! ⁷ ₇ ---!==---01 ⁰ ₁ - ⁸ ₉ ⁸ ₂ -----!! ⁸ _f ---1000---	0.0	c583f4662226eac0
$e_{7,2}$	0010 ⁸ ₀ -= ⁷ ₈ ---01000-11001000! ⁷ _c ---!-000011n ⁰ _e ---000 ⁷ ₇ ---01001100 ⁷ ₂ ---0	0.3	27c472268720c990
$e_{7,3}$	10110-00==00!==000110101n--1-n10111u01-10n1 ⁷ _e ---111110010-00	0.0	b0e1c6b1ee76ff28
$v_{8,0}$	-0-0-01 ⁹ _a - ⁸ ₉ -----!== ⁷ ₁ -01u1-----= ⁷ ₁ ⁹ ₇ ⁹ ₈ -----= ⁷ ₉ -----= ⁷ ₇ ---	2.6	aabf102cfb298eba
$v_{8,1}$	-u11011u11-10u ⁷ ₀ ---00111111n0-001n000-01! ⁹ ₁ -0n!==00111000u ⁷ ₂ --1	0.1	36d0c7f0c5760e09
$v_{8,2}$	1101100010-0011000111000110n!-000n110101! ⁷ _c -0n11110010001011 ⁷ ₂ --0	0.0	d8a638d87597c8b8
$v_{8,3}$	==!-00u10011 ⁸ ₈ -----110!--11u11-- ⁷ ₀ ====-1010101-1 ⁸ ₈ -----101 ⁷ _d --011!--	0.4	8899faec3eaa7adc
$e_{8,0}$	n01100001010000001111000110u00000000010001010011101010001001001		b0a078c00229d449
$e_{8,1}$	1u10011u00011u0010011101011n0000n01000011100n0110101001100u0001		a6189d7050e5a981
$e_{8,2}$	111101000000100110000100001n00001n10011110001n001101011000101110		f4098431678cd62e
$e_{8,3}$	1110100u0011000101110111111u101000010100111101111010101000110101		e83177ea14f7aa35
$e_{9,0}$	0n01011u10111u010001011000110000n01001100001n110111101110u1010		56b91630530f7dca
$e_{9,1}$	1010101n110n000011101000111011001n01011010110u010110100u01010010		abd0e8ecd6b16852
$e_{9,2}$	1101110u0011101011111100000110110n1111001000n00100000001100011		dc3afc1b7c848063
$e_{9,3}$	0n11000n11100101000u00100000100n0u11100101101111000n01000n0u110u		71e50209396f144c
$e_{10,0}$	uu000010nuuu1u0111111110001110nu0101001110000001110011u00un1100		0289ff1d29c0e61c
$e_{10,1}$	nnu101n0n1nnnn0uuunn01000n00uuu101001n1n00uuu01010u10101000unnn		d6fc3420a7814a87
$e_{10,2}$	0n00111000unnnnnnnnn111000100nuu10110101111nuu11100n010010n0nnnn		4e1ffe24b5f394af
$e_{10,3}$	nu0u0111nnuu0u11010u001u0001001nn010011100001n011000110n0uuu1010		87a34213a70d8d0a
$e_{11,0}$	11u1nuun10uuu11uuunn0011001nn10nn10nuuuuu10uuu100unn000u101u0unn		d986333dd14230a3
$e_{11,1}$	n1un1u00u1u1n1u0n0u101unu1nnn1n1nnn1n1uuuu0u01n00nn0010uu0111nu		d84e4abffe43321e
$e_{11,2}$	n1un0101n1uu0u101n0000u0011nu0uu10110nuuuuu00100nu000n1u1110un		d5c340385d0121b9
$e_{11,3}$	11uuuu01n1uuu101n1nnuun1100n100un00nu0101un0nu0n01u0nnnunun1nuun		c9d5f39892a94eb9
$v_{12,0}$	nu1n00u11nu1unu0u1n111un1nnnn1un110u1n1n1000un0n01n00un01n0u0001		b1d47dfdcf8562c1
$v_{12,1}$	nn00n010nu1nu0001nn00100n110n00nn10nunun00010100uu000011011u000u		cab0e4e9d5140360
$v_{12,2}$	nu011n1n1uunnu0nu01nuu1n1nu1uuu01110n11101u1unu011100u001n100n0		9f9933d0efaa7072
$v_{12,3}$	1u111u00uu011n01u01un010uu0u0u1uu0u0u11111n0u0nn00nuuuu100u11un0		b81d2a0207e3211a
$e_{12,0}$	xunu00u!==nn ¹ ₅ ⁵ _f V ¹ ₀ x10unn0nnnn-un01un ⁸ _x nu1 ⁷ ₆ -- ³ ₁ ³ _e ³ ₂ >=-1nnnn00nn ⁷ ₂ --1	0.0	211c537d5af4fe39
$e_{12,1}$	nnnuu1nuuu-nuuuuuuuuuuuuuuuu-nunnuu1nn0 ⁸ _e ---1001uu01000011un ⁹ ₈ --u	0.0	e6143042c70910d6
$e_{12,2}$	⁰ ₀ <-1-n1n0u1 ⁵ _a ² ₈ ⁵ _{j¹₀x--nu⁵_f³₃x-nuu¹₄³₄x-⁸₈xuuu-unnnnn¹₃³₈N¹₇xnnn³₆¹₁xu0²_b²₂>-}	0.5	ff30b0cec5f79fc9
$e_{12,3}$	-n-unnnunuuuu-ununnnunnnuunuuuu-uun-nuuuuun1nn----nuu0--u100u-	0.3	eda0bb950a0f0811
$e_{13,0}$	xuuu0nnnuuuuuu ⁷ _e ---000111nuuuuuu0100-unnnnnnnuu-uunnnuuuuunnnn	7.1	073083c021fe0f0f
$e_{13,1}$	xnnn1uuuuuuunnn ³ ₃ ---11100010000000-00101101000-nnnnnuuuuuuuuuu	5.4	f8cf7c400b47d0f0
$e_{13,2}$	x110-1001101u00---unn0--n100u--n-unuuuuuuuu-101010011111011010	11.2	ecd16c63d006a7da
$e_{13,3}$	-000-0101011n11---1000nnnnuuuuuuuu---000011u0100-10011010101111	0.0	82be91e18c32276f
$e_{14,0}$	xuuuuuuuuuuuuuuuuuuuuuuuuuuuuuu-!!--0101000-!!--01111111111111	3.6	00000002d45dff
$e_{14,1}$	-00-----0100100011110011---000== ⁰ ₃ ---100100=-1==11-111-1100010	0.0	8691e6867e4e3762
$e_{14,2}$	x11 ⁷ _e -111100unnnnnnnnn--1000!--u ⁷ ₄ !--000011!--0!--00111101001001	4.4	6f8ffe455c38cf49
$e_{14,3}$	-11101000011110000-111n0001100111 ⁷ ₂ ---1001010-!!--0-1101001011-	0.3	f43c3e33f255dd2e
$e_{15,0}$	-00---0100100011110011---000== ⁿ ₉ ⁶ ₆ ---11100 ⁷ ₇ ⁷ _e - ⁷ ₈ ⁸ ₃ -01-111-1100001	0.6	8691e686ab941761
$e_{15,1}$	-! ⁷ _e -111001 ⁷ ₈ ---10-0100---= ⁵ ₅ -u ⁷ ₈ ---010-1 ⁷ _e -1- ⁷ ₄ -!-00110111---	0.0	ef30a90e0533a378
$e_{15,2}$	x! ⁸ ₂ -0111100110000-1110-111!!--u ⁹ ₈ ---101000 ⁸ ₃ -1 ⁷ ₄ -!-11000111!!--	0.7	63cc3c794e8eac77
$e_{15,3}$	-00---01---0111010000---011---n ⁸ _c ---1 ⁷ _e !-- ⁷ ₄ - ⁸ ₂ ⁸ ₂ -!!--==0-111-0	0.0	0c8ba11cb26d2fbc
$v_{16,0}$	- ⁷ ₁ -----11! ⁷ _b -----!!--= ⁷ _b -----0! ⁷ ₉ ----- ⁸ ₅ ⁸ ₅ - ⁷ ₁ -01!--! ⁷ ₄ -00-	0.0	75c28f94b0c7bad9
$v_{16,1}$	----- ⁸ _a ----- ⁷ ₄ !------- ⁸ ₀ ----- ⁸ _a -----! ⁷ ₁ -----01---	0.0	c23af22a0c707d2f
$v_{16,2}$	x! ⁷ _f ---0 ⁹ _a ---1111!--111 ⁷ ₂ ---0 ⁷ _d ⁹ _d ----- ⁸ ₀ ⁸ ₁ ⁷ _b ----- ⁷ ₇ ⁷ _f ---! ⁷ _a -0-10 ⁷ _d ---	0.0	7057dd9600fbc33
$v_{16,3}$	x-----01-----!!--= ⁷ ₄ -----1= ⁷ _a -----01-----!!--	0.0	70f12cec5ff713d7

