

Revisiting AES Related-Key Differential Attacks with Constraint Programming

David Gérardt¹, Pascal Lafourcade¹, Marine Minier^{b,*}, Christine Solnon^c

^aUniversité Clermont Auvergne, LIMOS, UMR 6158, F-63173, France

^bUniversité de Lorraine, LORIA, UMR 7503, F-54506, France

^cUniversité de Lyon, INSA-Lyon, F-69621, France, LIRIS, CNRS UMR5205

Abstract

The Advanced Encryption Standard (AES) is one of the most studied symmetric encryption schemes. During the last years, several attacks have been discovered in different adversarial models. In this paper, we focus on related-key differential attacks, where the adversary may introduce differences in plaintext pairs and also in keys. We show that Constraint Programming (CP) can be used to model these attacks, and that it allows us to efficiently find all optimal related-key differential characteristics for AES-128, AES-192 and AES-256. In particular, we improve the best related-key differential for the whole AES-256 and give the best related-key differential on 10 rounds of AES-192, which is the differential trail with the longest path. Those results allow us to improve existing related-key distinguishers, basic related-key attacks and q -multicollisions on AES-256.

Keywords: AES, Related-key attacks, Constraint Programming

Introduction

As attacking the Advanced Encryption Standard (AES) in the unknown key model seems to be out of reach at this time, many recent results focus on the so-called *related-key*, *known-key* or *chosen-key* models [1, 2, 3]. During the last decade, many results bring some grist to this research direction. In particular, the notion of *differential q -multicollisions* was introduced in [4]. A differential q -multicollision for a cipher $E_K(\cdot)$ is defined by a non zero key difference δK , a non zero plaintext difference δX and a set of q distinct pairs (X^i, K^i) with $i \in [1, q]$ such that all $E_{K^i}(X^i) \oplus E_{K^i \oplus \delta K}(X^i \oplus \delta X)$ are equal. Constructing such a q -multicollision for an ideal n -bit block cipher has a time complexity of $\mathcal{O}(q \cdot 2^{\frac{q-2}{q+2}n})$. However, for AES-256 the number of required AES encryptions has been shown to be equal to $q \cdot 2^{67}$ in [4].

Building such q -multicollisions requires finding optimal (in terms of probability) related-key differential characteristics. This challenging task was tackled for AES-128 with a graph

*Corresponding author

Email addresses: David.Gerault@uca.fr (David Gérardt), Pascal.Lafourcade@uca.fr (Pascal Lafourcade), Marine.Minier@loria.fr (Marine Minier), Christine.Solnon@insa-lyon.fr (Christine Solnon)

AES-192						
Attack	Nb rounds	Nb keys	Data	Time	Memory	Source
RK rectangle	10	64	2^{124}	2^{183}	N/A	[6]
RK amplified boomerang	12	4	2^{123}	2^{176}	2^{152}	[7]
RK distinguisher	10	2^{80}	2^{108} *	2^{108} *	-	Section 2.1
basic RK differential	10	2^{44}	2^{156}	2^{156}	2^{65}	Section 2.1
AES-256						
Attack	Nb rounds	Nb keys	Data	Time	Memory	Source
RK boomerang	14	4	$2^{99.5}$	$2^{99.5}$	2^{77}	[7]
RK distinguisher	14	2^{35}	2^{119} *	2^{119} *	-	[4]
basic RK differential	14	2^{35}	2^{131}	2^{131}	2^{65}	[4]
q -multicollisions	14	$2q$	$2q$	$q2^{67}$	-	[4]
RK distinguisher	14	2^{32}	2^{114} *	2^{114} *	-	Section 2.2
basic RK differential	14	2^{32}	2^{125}	2^{125}	2^{65}	Section 2.2
q -multicollisions	14	$2q$	$2q$	$q2^{66}$	-	Section 2.2

Table 1: Summary of existing attacks against AES-192 and AES-256 in the related-key and chosen-key models. RK stands for Related-Key, N/A means Not Available and * means for each key.

traversal approach in [3], and for AES-128, AES-192, and AES-256 with a depth-first search approach in [4]. However, the 4-round solution for AES-128 claimed to be optimal in [4, 3] has been shown to be sub-optimal in [5]. In [5], the authors used Constraint Programming (CP) to efficiently enumerate related-key differential characteristics on AES-128.

In this paper, we further investigate the interest of using CP for finding optimal related-key differential characteristics for AES-192 and AES-256 whereas [5] has only focused on AES-128. We give new optimal solutions found with our CP approach. Table 1 sums up our new results in different attack models.

In Section 1, we give a brief overview of how our CP models work. In Section 2, we show how to use the solutions found by our CP models to improve existing related-key differential attacks for AES-192 and AES-256.

1. CP models for finding AES related-key differential paths

Mounting related-key differential attacks requires finding a related-key differential characteristic [8, 1], *i.e.* a plaintext difference $\delta X = X \oplus X'$ and a key difference $\delta K = K \oplus K'$, such that δX becomes δX_r after r rounds with a probability as high as possible. The AES operations `ShiftRows` (SR), `MixColumns` (MC), `AddRoundKey` (ARK) are linear, *i.e.*, they propagate differences in a deterministic way (with probability 1). The only non-linear operation is `SubBytes` (SB) where the used S-box S transforms a given difference into another one in a probabilistic way. Even if the most important part of the AES `KeySchedule` (KS) is linear, it also makes regular calls to the S-box S .

To find optimal related-key differential characteristics and as done in [9] and [3], we use a two-step solving process. Step 1 works with a boolean representation of differences: We denote ΔA the boolean representation of the byte difference δA such that $\Delta A = 0 \Leftrightarrow \delta A = 0$ and $\Delta A = 1 \Leftrightarrow \delta A \in [1, 255]$. These boolean variables give difference positions. The goal of Step 1 is to find a *Boolean solution* that assigns values to Boolean variables such that the AES transformation rules are satisfied. During this first step, the `SubBytes` operation SB is not considered. Indeed, it does not introduce nor remove differences.

Then, Step 2 uses these positions to determine difference values at the byte level, *i.e.*, to find the actual value $\delta A \in [1, 255]$ for each boolean variable ΔA which is equal to 1. Note that some solutions at the boolean level (found during Step 1) cannot be transformed into solutions at the byte level during Step 2. These solutions are said to be *byte inconsistent*.

1.1. Basic CP Model for Step 1

A first CP model for Step 1 may be derived from the AES transformations in a rather straightforward way. We extend the model described in [5] for AES-128 to AES-192 and AES-256. A CP model is defined by a set of variables, such that each variable x has a domain $D(x)$, and a set of constraints, *i.e.*, relations that restrict the values that may be simultaneously assigned to the variables. For each differential byte δB , we define a Boolean variable ΔB whose domain is $D(\Delta B) = \{0, 1\}$: it is assigned to 0 if $\delta B = 0$, and to 1 otherwise.

The XOR constraint for ARK and KS. We first define a XOR constraint for *ARK* and *KS*. Let us consider three differential bytes δA , δB and δC such that $\delta A \oplus \delta B = \delta C$. If $\delta A = \delta B = 0$, then $\delta C = 0$. If $(\delta A = 0 \text{ and } \delta B \neq 0)$ or $(\delta A \neq 0 \text{ and } \delta B = 0)$ then $\delta C \neq 0$. However, if $\delta A \neq 0$ and $\delta B \neq 0$, then we cannot know if δC is equal to 0 or not: This depends on whether $\delta A = \delta B$ or not. When abstracting differential bytes δA , δB and δC with Boolean variables ΔA , ΔB and ΔC (which only model the fact that there is a difference or not), we obtain the following definition of the XOR constraint: $\text{XOR}(\Delta A, \Delta B, \Delta C) \Leftrightarrow \Delta A + \Delta B + \Delta C \neq 1$.

Both *ARK* and *KS* are directly modeled with XOR constraints. The definition of *KS* depends on the key length. The CP model defined in [5] only considers 128 bit key length. We have extended it to 192 and 256 bit lengths in a rather straightforward way.

ShiftRows and MixColumns. *SR* simply shifts variables. The MDS property of *MC* is ensured by posting a constraint on the sum of all variables on a same column before and after *MC*, which must belong to the set $\{0, 5, 6, 7, 8\}$.

Objective function. The goal is to minimize the number of S-boxes that must be crossed by Boolean differential paths. This is done in CP by introducing an integer variable obj_{Step1} which is constrained to be equal to the sum of all boolean variables associated with bytes on which a non linear transformation S is applied.

Limitations of the basic CP model. This basic CP model CP_{basic} is complete, *i.e.*, for any solution at the byte level (on δ variables), there exists a solution at the Boolean level (on Δ variables). However, preliminary experiments reported in [5] have shown us that there is a huge number of Boolean solutions which are byte inconsistent. For example, when the number of rounds is $r = 4$ for AES-128, the optimal cost is $obj_{Step1} = 11$, and there are more than 90 millions of Boolean solutions with $obj_{Step1} = 11$. However, none of these solutions is byte-consistent. In this case, most of the Step 1 solving time is spent at generating useless Boolean solutions which are discarded in Step 2.

1.2. Additional Constraints for Step 1

We have introduced in [5] a second model called CP_{EQ} for AES-128. This model removes many byte-inconsistent solutions by propagating equality constraints at the byte level. In this paper, we extend it to AES-192 and AES-256. A full version of this model could be found in [10] and in [11]¹.

New equality variables. For each couple of differential bytes $(\delta A, \delta B)$, we introduce a Boolean equality variable $EQ_{\delta A, \delta B}$ which is equal to 1 if $\delta A = \delta B$, and to 0 otherwise. These variables are constrained to define an equivalence relation by adding a symmetry constraint ($EQ_{\delta A, \delta B} = EQ_{\delta B, \delta A}$) and a transitivity constraint (if $EQ_{\delta A, \delta B} = EQ_{\delta B, \delta C} = 1$ then $EQ_{\delta A, \delta C} = 1$). Also, EQ variables are related to Δ variables by adding the constraints:

$$(EQ_{\delta A, \delta B} = 1) \Rightarrow (\Delta A = \Delta B) \text{ and } EQ_{\delta A, \delta B} + \Delta A + \Delta B \neq 0$$

Revisiting the XOR constraint. When defining the constraint $XOR(\Delta A, \Delta B, \Delta C)$, if $\Delta A = \Delta B = 1$, then we cannot know if ΔC is equal to 0 or 1. However, whenever $\Delta C = 0$ (resp. $\Delta C = 1$), we know for sure that the corresponding byte δC is equal to 0 (resp. different from 0), meaning that the two bytes δA and δB are equal (resp. different), *i.e.*, that $EQ_{\delta A, \delta B} = 1$ (resp. $EQ_{\delta A, \delta B} = 0$). The same reasoning may be done for ΔA and ΔB because $(\delta A \oplus \delta B = \delta C) \Leftrightarrow (\delta B \oplus \delta C = \delta A) \Leftrightarrow (\delta A \oplus \delta C = \delta B)$. Therefore, we redefine the XOR constraint as follows:

$$\begin{aligned} XOR(\Delta A, \Delta B, \Delta C) \Leftrightarrow & ((\Delta A + \Delta B + \Delta C \neq 1) \wedge (EQ_{\delta A, \delta B} = 1 - \Delta C) \\ & \wedge (EQ_{\delta A, \delta C} = 1 - \Delta B) \wedge (EQ_{\delta B, \delta C} = 1 - \Delta A)) \end{aligned}$$

Propagation of MDS at Byte Level. The MDS property ensures that, for each column, the total number of bytes which are different from 0, before and after applying MC , is either equal to 0 or strictly greater than 4. This property also holds for any xor difference between two different columns at different rounds. To propagate this property, for each pair of columns, we add a constraint on the sum of equality variables between bytes of these columns.

Constraints derived from KS. The `KeySchedule` mainly performs xor and S-box operations. As a consequence, each subkey byte $\delta K_i[j][k]$ at round i may be expressed as a xor between bytes of the original key difference $\delta K[j][k]$, and bytes that have passed through an S-box at round $i - 1$, denoted by $\delta S(K_{i-1}[j][k])$. Hence, for each byte $\delta K_i[j][k]$, we precompute the set $V(i, j, k)$ such that $V(i, j, k)$ only contains bytes of δK and $\delta S(K_{i-1})$ and $\delta K_i[j][k] = \bigoplus_{\delta A \in V(i, j, k)} \delta A$. For each set $V(i, j, k)$, we introduce a set variable $V_1(i, j, k)$ which is constrained to contain the subset of $V(i, j, k)$ corresponding to the Boolean variables equal to 1. We use these set variables to infer that two differential key bytes that have the same V_1 set are equal. Also, if $V_1(i, j, k)$ is empty (resp. contains one or two elements), we infer that $\Delta K_i[j][k]$ is equal to 0 (resp. a variable, or a xor between 2 variables).

¹The code of [11] is available through http://www.gerault.net/Doctoral_Program_CP17.zip.

1.3. CP Model for Step 2

Given a Boolean solution for Step 1, Step 2 aims at searching for the byte-consistent solution with the highest probability (or proving that there is no byte-consistent solution). Hence, for each Boolean variable ΔA of Step 1, we define an integer variable δA whose domain depends on the value of ΔA : If $\Delta A = 0$, then $D(\delta A) = \{0\}$ (i.e., δA is also assigned to 0); otherwise, $D(\delta A) = [1, 255]$. As we look for a byte-consistent solution with maximal probability, we also add an integer variable P_A for each byte A that passes through an S-box: This variable corresponds to the base 2 logarithm of the probability $\Pr(\delta A \rightarrow \delta S_A)$ of obtaining the output difference δS_A when the input difference is δA . If $\Delta A = 0$, then $\Pr(0 \rightarrow 0) = 1$ and therefore $D(P_A) = \{0\}$; otherwise, $\Pr(\delta A \rightarrow \delta S_A) \in \{\frac{2}{256}, \frac{4}{256}\}$ and $D(P_A) = \{-7, -6\}$.

At byte level, the `SubBytes` transformation, which has no effect at the Boolean level, must be modeled. This is done thanks to a ternary table constraint which extensively lists all triples (A, S_A, P_A) such that there exist two bytes B_1 and B_2 whose differences before and after passing through the S-box S are equal to A and S_A , respectively with a \log_2 probability equal to P_A . To find a byte-consistent solution with maximal differential probability, we maximize the sum of all P_A variables.

2. From related-key differentials to related-key attacks

In this Section, we summarize the new AES related-key differential paths that we have computed with the new CP models previously described, and give new basic related-key attacks, related-key distinguishers and q -multicollisions that we are able to mount by using them for AES-192 and AES-256.

2.1. AES-192

Summary of related-key differential paths computed with CP. Using our CP approach, we found that the best related-key differential trail is on 10 rounds with 29 active S-boxes and a highest probability equal to 2^{-176} : 2^{-37} coming from the keys and 2^{-139} from the ciphering part. The best differential characteristic is given in Table A.3 of Appendix A. We also give another trail with 30 active S-boxes where the differential characteristic has a probability equal to 2^{-188} : 2^{-80} coming from the keys and 2^{-108} from the ciphering part. This trail is optimal for the probability in the state and is given in Table A.4 of Appendix A.

The first differential characteristic, which has an optimal probability, allows us to mount a basic related-key differential attack as done in [4]. The second one allows us to build a related-key distinguisher as it minimizes the probability in the state. We also provide in Table A.5 of Appendix A, the best differential characteristic on 9 rounds of AES-192. This characteristic has a probability of 2^{-146} with 24 active S-boxes and is better than the one presented in [12].

Related-key distinguisher on 10 rounds. For this distinguisher, we use the related key differential characteristic given in Table A.4 in Appendix A which has a probability equal to 2^{-188} . Considering the related-key distinguisher model, the probability that the differences correctly propagate through the internal states is $2^{-108} = 2^{-18 \cdot 6}$ as we have 18 active S-boxes in the

internal states, all with probability 2^{-6} . It works for 1 out of $2^{80} = 2^{6 \cdot 4} \cdot 2^{7 \cdot 8}$ keys as we have 12 additional active S-boxes in the key schedule: 8 with probability 2^{-7} and the 4 others with probability 2^{-6} . Therefore, the related-key distinguisher works with complexity 2^{108} for 1 out of 2^{80} related-key pairs on average.

Basic related-key differential attack on 10 rounds. We first change the trail given in Table A.3 of Appendix A in the deciphering direction to get more active S-boxes in the two last rounds to recover the key bytes implied at the input of those S-boxes. The new trail has eight active S-boxes in the last round on two anti-diagonals and two active S-boxes in the penultimate round.

Thus, we must find the 10 key bytes $K_{10}^*[0][2]$, $K_{10}^*[1][3]$, $K_{10}^*[2][0]$, $K_{10}^*[3][1]$, $K_{10}^*[0][3]$, $K_{10}^*[1][0]$, $K_{10}^*[2][1]$, $K_{10}^*[3][2]$, $K_9^*[0][2]$ and $K_9^*[0][3]^2$. We use the following procedure from the ciphertexts for each of the $2^{37} \times 2^7$ key pairs³:

1. Repeat 2^{47} times:
 - (a) Compose two structures of 2^{64} ciphertexts with all possible values for the first and second anti-diagonals. Decrypt the first structure with K and the second one with K' .
 - (b) Sort the plaintexts and check for a pair with the correct input difference. Save these valid pairs if any.
2. For each of these pairs, derive 2^{16} variants for the 10 key bytes. There are 10 S-boxes in the two last rounds for which we know the input and output differences. Therefore, there are $2^{10} \cdot 8 \cdot 8 = 2^{16}$ possibilities for the 80 key bits per candidate pair without false alarms.
3. Pick the key candidate with the best occurrence.

The overall complexity of the whole procedure, which is repeated 2^{44} times, is $2^{47+65} = 2^{112}$ in data and time and 2^{65} in memory. We need to repeat 2^{47} times step 1.(a) and 1.(b) to keep on average $2^{47} \cdot 2^{64-109} = 2^4 = 16$ right pairs and 2^{47} wrong pairs and to completely discard false alarms. This gives us 80 key bits.

2.2. AES-256

Summary of related-key differentials computed with CP. The optimal byte solution for 14 rounds of AES-256 has a probability of 2^{-146} , and it is given in Table A.2 of Appendix A. Note that the one given in [4] has a probability of 2^{-154} to happen. We also obtain 43 solutions with a probability of 2^{-147} . We experimentally checked that the 7 bottom rounds of the AES conform to the expected probability by producing the wanted difference after 2^{30} pairs on average, as predicted by the trail.

²where * stands for the classical `InvMixColumns` transformation applied on the keys.

³the 2^7 term comes from all the possible unknown differences at the S-box output generated through the key schedule for the possible values of $\delta K_{10}^*[0][2]$, $\delta K_{10}^*[1][3]$, $\delta K_{10}^*[0][3]$ and $\delta K_{10}^*[3][2]$.

q-multicollisions. Using this optimized byte consistent differential solution, we are able to improve the attacks proposed in [4]. More precisely, the cost to compute a q -multicollision given in [4] depends on 11 active S-boxes with probability 2^{-67} . With the new differential characteristic, we gain a factor 2 on this complexity leading to a time complexity equal to $q \cdot 2^{66}$ encryptions. In the same way, the time complexity to find partial q -multicollisions becomes $q \cdot 2^{36}$ instead of $q \cdot 2^{37}$.

Related-key distinguisher. Considering the related-key distinguisher, the probability that the differences correctly propagate through the internal states is $2^{-19 \cdot 6} = 2^{-114}$ and it works for 1 out of $2^{32} = 2^{14} \cdot 2^{18}$ keys as we have 5 additional active S-boxes: 2 with probability 2^{-7} and 3 with probability 2^{-6} . Hence, the related-key distinguisher has a data/time complexity equal to $2^{146} = 2^{114} \cdot 2^{32}$.

Related-key attack. The related-key attack described in [4] may be directly applied using our new differential characteristic. The procedure is about to be the same than the one of [4] and for the basic related-key attack given for 10 AES-192 rounds except that the two modified rounds are rounds 0 and 1. In this case, the number of key bits to test in the two first rounds is equal to 80 and for each of the 2^{32} possible key pairs, we need to repeat 2^{28} times the process with two structures of 2^{64} plaintexts. Thus after Step 1, we have on average 4 right pairs, and for each pair we derive 2^{16} possible values for the ten key bytes without false alarms. Thus, the overall complexity of this attack becomes 2^{125} in data and time while testing 2^{32} keys. The required memory is 2^{65} .

References

- [1] E. Biham, New types of cryptanalytic attacks using related keys (extended abstract), in: Advances in Cryptology - EUROCRYPT '93, Vol. 765 of LNCS, Springer, 1993, pp. 398–409.
- [2] L. R. Knudsen, V. Rijmen, Known-key distinguishers for some block ciphers, in: Advances in Cryptology - ASIACRYPT 2007, Vol. 4833 of LNCS, Springer, 2007, pp. 315–324.
- [3] P. Fouque, J. Jean, T. Peyrin, Structural evaluation of AES and chosen-key distinguisher of 9-round AES-128, in: Advances in Cryptology - CRYPTO 2013 - Part I, Vol. 8042 of LNCS, Springer, 2013, pp. 183–203.
- [4] A. Biryukov, D. Khovratovich, I. Nikolic, Distinguisher and related-key attack on the full AES-256, in: Advances in Cryptology - CRYPTO 2009, Vol. 5677 of LNCS, Springer, 2009, pp. 231–249.
- [5] D. Gerault, M. Minier, C. Solnon, Constraint programming models for chosen key differential cryptanalysis, in: Principles and Practice of Constraint Programming - CP 2016, Vol. 9892 of LNCS, Springer, 2016, pp. 584–601.

- [6] J. Kim, S. Hong, B. Preneel, Related-key rectangle attacks on reduced AES-192 and AES-256, in: Fast Software Encryption - FSE 2007, Vol. 4593 of LNCS, Springer, 2007, pp. 225–241.
- [7] A. Biryukov, D. Khovratovich, Related-key cryptanalysis of the full AES-192 and AES-256, in: Advances in Cryptology - ASIACRYPT 2009, Vol. 5912 of LNCS, Springer, 2009, pp. 1–18.
- [8] E. Biham, A. Shamir, Differential cryptanalysis of feal and n-hash, in: Advances in Cryptology - EUROCRYPT '91, Vol. 547 of LNCS, Springer, 1991, pp. 1–16.
- [9] A. Biryukov, I. Nikolic, Automatic search for related-key differential characteristics in byte-oriented block ciphers: Application to aes, camellia, khazad and others, in: Advances in Cryptology - EUROCRYPT 2010, Vol. 6110 of LNCS, Springer, 2010, pp. 322–344.
- [10] D. Gérard, P. Lafourcade, M. Minier, C. Solnon, Revisiting aes related-key differential attacks with constraint programming, Cryptology ePrint Archive, Report 2017/139, <http://eprint.iacr.org/2017/139> (2017).
- [11] D. Gerault, P. Lafourcade, M. Minier, C. Solnon, Combining Solvers to Solve a Cryptanalytic Problem, CP/ICLP/SAT Doctoral Program - part of the conference CP 2017, available at http://www.gerault.net/CP17_DP.pdf (2017).
- [12] I. Nikolic, Cryptanalysis and design of symmetric primitives, Ph.D. thesis, University of Luxembourg, Luxembourg, Luxembourg (2011).

Appendix A. Our New Related-Key Differential Paths for AES-192 and AES-256

Round	$\delta X_i = X_i \oplus X'_i$	$\delta K_i = K_i \oplus K'_i$	Pr(States)	Pr(Key)
init.	addbdb76 addbdb76 addbdb76 addbdb76			
$i = 0$	69000000 00000000 69000000 00000000	c4dbdb76 addbdb76 c4dbdb76 addbdb76	$2^{-6 \cdot 2}$	—
1	9a000000 00000000 9a000000 00000000	b59a9ab5 00000000 b59a9ab5 00000000	$2^{-6 \cdot 2}$	—
2	69000000 69000000 00000000 00000000	c4dbdb76 69000000 addbdb76 00000000	$2^{-6 \cdot 2}$	—
3	9a000000 9a000000 00000000 00000000	b59a9ab5 b59a9ab5 00000000 00000000	$2^{-6 \cdot 2}$	—
4	69000000 00000000 00000000 00000000	c4dbdb76 addbdb76 00000000 00000000	2^{-6}	—
5	9a000000 00000000 00000000 00000000	b59a9ab5 00000000 00000000 00000000	2^{-6}	—
6	69000000 69000000 69000000 69000000	c4dbdb76 69000000 69000000 69000000	$2^{-6 \cdot 4}$	2^{-6}
7	00000000 00000000 00000000 00000000	2f9a9ab5 2f9a9ab5 2f9a9ab5 2f9a9ab5	—	$2^{-7 \cdot 2} \times 2^{-6 \cdot 2}$
8	69000000 00000000 69000000 00000000	69000000 00000000 69000000 00000000	$2^{-6 \cdot 2}$	—
9	00000000 00000000 00000000 00000000	2f9a9ab5 00000000 2f9a9ab5 00000000	—	—
10	69000000 69000000 00000000 00000000	69000000 69000000 00000000 00000000	$2^{-6 \cdot 2}$	—
11	00000000 00000000 00000000 00000000	2f9a9ab5 2f9a9ab5 00000000 00000000	—	—
12	69000000 00000000 00000000 00000000	69000000 00000000 00000000 00000000	2^{-6}	—
13	00000000 00000000 00000000 00000000	2f9a9ab5 00000000 00000000 00000000	—	—
End/14	69000000 69000000 69000000 69000000	69000000 69000000 69000000 69000000	—	—

Table A.2: Our own related-key differential on 14 AES-256 rounds that happens with a probability 2^{-146} .

Round	$\delta X_i = X_i \oplus X'_i$	$\delta K_i = K_i \oplus K'_i$	Pr(States)	Pr(Key)
init.	c816ad91 dc02027a d8000000 00000000			
$i = 0$	d800a300 d800007a 00000000 00000000	10160e91 04020200 d8000000 00000000	$2^{-6.4}$	—
1	00000000 00000000 d4000000 00000000	04020206 04020206 04020200 00000000	2^{-7}	$2^{-7} \times 2^{-6.3}$
2	d8000000 d8000000 d8000000 d8000000	d8000000 d8000000 dc020206 d8000000	$2^{-6.4}$	2^{-6}
3	00000000 00000000 d8000000 00000000	04020206 04020206 dc020206 04020206	2^{-6}	—
4	d8000000 00000000 00000000 00000000	d8000000 00000000 04020206 00000000	2^{-6}	—
5	d8000000 d8000000 00000000 00000000	dc020206 d8000000 00000000 00000000	$2^{-6.2}$	—
6	00000000 00000000 d8000000 00000000	04020206 04020206 d8000000 00000000	2^{-6}	—
7	00000000 00000000 00000000 00000000	00000000 00000000 04020206 00000000	—	—
8	d8000000 d8000000 d8000000 d8000000	d8000000 d8000000 d8000000 d8000000	$2^{-6.4}$	2^{-6}
9	00000002 00000002 d8000002 00000002	04020204 04020204 dc020204 04020204	$2^{-6.5}$	—
End/10	d8000400 06000400 ????????? ?????????	dc020204 04020204 ????????? ?????????	—	—

Table A.3: Our first related-key differential on 10 AES-192 rounds that happens with a probability 2^{-176} .

Round	$\delta X_i = X_i \oplus X'_i$	$\delta K_i = K_i \oplus K'_i$	Pr(States)	Pr(Key)
init.	e00411ef 00000000 140a0a1e 00000000			
$i = 0$	00000000 e4000000 00000000 00000000	e00411ef e4000000 140a0a1e 00000000	2^{-6}	—
1	e4000000 e4000000 e4000000 00000000	e4000000 f00a0a1e e4000000 00000000	$2^{-6.3}$	$2^{-7.2} \times 2^{-6.2}$
2	00000000 00000000 e4000000 00000000	140a0a1e 140a0a1e f00a0a1e 00000000	2^{-6}	—
3	e4000000 e4000000 e4000000 e4000000	e4000000 e4000000 f00a0a1e e4000000	$2^{-6.4}$	—
4	00000000 00000000 e4000000 00000000	140a0a1e 140a0a1e f00a0a1e 140a0a1e	2^{-6}	$2^{-7.3} \times 2^{-6}$
5	e4000000 00000000 00000000 00000000	e4000000 00000000 140a0a1e 00000000	2^{-6}	—
6	e4000000 e4000000 00000000 00000000	f00a0a1e e4000000 00000000 00000000	$2^{-6.2}$	$2^{-7.3} \times 2^{-6}$
7	00000000 00000000 e4000000 00000000	140a0a1e 140a0a1e e4000000 00000000	2^{-6}	—
8	00000000 00000000 00000000 00000000	00000000 00000000 00000000 00000000	—	—
9	e4000000 e4000000 e4000000 e4000000	e4000000 e4000000 e4000000 e4000000	$2^{-6.4}$	—
End/10	e4000000 e4000000 ????????? ?????????	f00a0a1e f00a0a1e ????????? ?????????	—	—

Table A.4: Our related-key differential on 10 AES-192 rounds with 30 active S-boxes and a probability of 2^{-188} .

Round	$\delta X_i = X_i \oplus X'_i$	$\delta K_i = K_i \oplus K'_i$	Pr(States)	Pr(Key)
init.	8e1b400c 9603039e 90000000 00000000			
$i = 0$	90004b00 9000009e 00000000 00000000	1e1b0b0c 06030300 90000000 00000000	$2^{-6.4}$	—
1	00000000 00000000 be000000 00000000	06030305 06030305 06030300 00000000	2^{-7}	$2^{-7} \times 2^{-6.3}$
2	90000000 90000000 90000000 90000000	90000000 90000000 96030305 90000000	$2^{-6.4}$	2^{-6}
3	00000000 00000000 90000000 00000000	06030305 06030305 96030305 06030305	2^{-6}	—
4	90000000 00000000 00000000 00000000	90000000 00000000 06030305 00000000	2^{-6}	—
5	90000000 90000000 00000000 00000000	96030305 90000000 00000000 00000000	$2^{-6.2}$	—
6	00000000 00000000 90000000 00000000	06030305 06030305 90000000 00000000	2^{-6}	—
7	00000000 00000000 00000000 00000000	00000000 00000000 06030305 00000000	—	—
8	90000000 90000000 90000000 90000000	90000000 90000000 90000000 90000000	$2^{-6.4}$	2^{-6}
End/9	06030305 06030305 06030305 06030305	00000000 00000000 00000000 00000000	—	—

Table A.5: Our related-key differential on 9 AES-192 rounds with 24 active S-boxes and a probability of 2^{-146} .