# Distinguisher and Related-Key Attack on the Full AES-256 

Alex Biryukov, Dmitry Khovratovich, and Ivica Nikolić<br>University of Luxembourg<br>\{alex.biryukov, dmitry.khovratovich,ivica.nikolic@uni.lu\}


#### Abstract

In this paper we construct a chosen-key distinguisher and a related-key attack on the full 256 -bit key AES. We define a notion of differential $q$-multicollision and show that for AES-256 $q$-multicollisions can be constructed in time $q \cdot 2^{67}$ and with negligible memory, while we prove that the same task for an ideal cipher of the same block size would require at least $O\left(q \cdot 2^{\frac{q-1}{q+1} 128}\right)$ time. Using similar approach and with the same complexity we can also construct $q$-pseudo collisions for AES-256 in Davies-Meyer mode, a scheme which is provably secure in the ideal-cipher model. We have also computed partial $q$-multicollisions in time $q \cdot 2^{37}$ on a PC to verify our results. These results show that AES-256 can not model an ideal cipher in theoretical constructions. Finally we extend our results to find the first publicly known attack on the full 14-round AES-256: a related-key distinguisher which works for one out of every $2^{35}$ keys with $2^{120}$ data and time complexity and negligible memory. This distinguisher is translated into a key-recovery attack with total complexity of $2^{131}$ time and $2^{65}$ memory.


Keywords: AES, related-key attack, chosen key distinguisher, DaviesMeyer, ideal cipher.

## 1 Introduction

The Advanced Encryption Standard (AES) is a block cipher which was chosen by NIST from a set of 15 candidate designs in a thorough evaluation process that lasted from September 1997 till October 2000. On November 26, 2001 Rijndael [5], a 128-bit block, 128/192/256-bit key block cipher has become a standard as U.S. FIPS 197 [12. In June 2003 the US government has approved the use of $128,192,256$ bit key AES for SECRET and 192, 256-bit key AES for TOP SECRET information [13]. In the last ten years AES has been subject to very intensive cryptanalytic effort, with best currently known attacks breaking 7, 10, 10 rounds for respective keysizes (128, 192, 256), with very high complexities.

In this paper we show for the first time in the open literature distinguishers and related-key attacks on the full 14-round 256 -bit key AES. Research presented in this paper follows the logic described in Fig. [1] First we identified slow diffusion and other differential weaknesses in the key schedule of AES-256 which match


Fig. 1. Outline of the research presented in this paper
nicely with the differential properties of the round function. This allows us to construct local collisions for AES, i.e. two round difference propagation patterns which result in low weight difference in the subkeys and zero difference in the 128 -bit block. We concatenate four such local collisions together and add another 6 -round trail on top in order to cover full 14 rounds of AES-256. The trai 1 has 41 active S-boxes ( 36 in the block and 5 in the key schedule), so we apply a special tool, a triangulation algorithm (designed for the purpose of finding collisions in hash functions), in order to find keys and plaintexts that conform to the trail.

From this point we go in two directions. First we show that for AES-256 one can construct a chosen-key distinguisher based on the new notion of a differential $q$-multicollision in time $q \cdot 2^{67}$ and with negligible memory. We prove that the same task for an ideal cipher of the same block size would require at least $q \cdot 2^{\frac{q-1}{q+1} 128}$ time for $q \leq 57$ and at least $q \cdot 2^{\frac{q-2}{q+2} 128}$ for $q>57$. I.e. for $q>3$ the differential multicollision for AES-256 can be constructed significantly faster than for an ideal cipher. Previously a known-key distinguisher for seven rounds of AES with $2^{56}$ texts was found in [11. To verify our results we found partial $q$-multicolisions in several hours on a PC using the publicly available implementation of AES-256.

As a direct application of this differential $q$-multicollision distinguisher we show that AES-256 when used in the Davies-Meyer mode allows to construct $q$ pseudo-collisions with fixed differences $\Delta I V, \Delta M$ in the IV and the message with complexity $q \cdot 2^{67}$. Again, such a result would require at least $q \cdot 2^{\frac{q-2}{q+2} 128}$ time for the ideal cipher in the Davies-Meyer mode. Results of this type try to enchance our definitions of block cipher security and to fill the gap between theoretical models like random oracle and ideal cipher and the real world of ciphers which have fixed description and are efficiently computable [24]. However a proper security definition which would capture the intuition behind chosen/known key attacks is still an open problem.

[^0]Table 1. Best attacks on AES-256

| Attack | \# rounds | \# keys | Data | Time | Memory | Source |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Known-key integral | 7 | 1 | $2^{56}$ | $2^{56}$ | $2^{56}$ | $[11]$ |
| Partial sums | 9 | 256 | $2^{85}$ | $2^{226}$ | $2^{32}$ | $[6$ |
| Related-key rectangle | 10 | 64 | $2^{114}$ | $2^{173}$ | $?$ | $[110]$ |
| $q$-multicollisions | 14 | $2 q$ | $2 q$ | $q \cdot 2^{67}$ | - | Sec. [2] |
| Partial $q$-multicollisions | 14 | $2 q$ | $2 q$ | $q \cdot 2^{37}$ | - | Sec. [2.3 |
| Related-key distinguisher | 14 | $2^{35}$ | $2^{119 *}$ | $2^{119 *}$ | - | Sec. 4.1 |
| Related-key key recovery | 14 | $2^{35}$ | $2^{96 *}$ | $2^{96 *}$ | $2^{65}$ | Sec. [4.2 |

*     - for each key.

The second direction that we studied was application of the trails that we have found to more standard attacks on a block cipher, for example related-key attacks. In particular we show that by changing the top two rounds of the trail that we used previously one obtains a differential trail with only 24 active Sboxes ( 19 in the round function and 5 in the key schedule). From this trail we can construct a differential distinguisher for AES-256 which works for one key out of $2^{35}$ and has complexity $2^{120}$ data and time, and negligible memory. This distinguisher can be used to mount a key-recovery attack on AES-256 with total complexity of $2^{35+96}=2^{131}$ time and $2^{65}$ memory. We summarize our findings in Table 1 .

This paper is organized as follows: in Section 2 we prove a lower bound on the complexity of finding differential $q$-multicollisions in the case of an ideal cipher and construct a distinguisher for the full AES-256. In Section 3 we show an application of these results to finding pseudo-collisions for the Davies-Meyer hashing mode instantiated with AES-256. In Section 4 we show a related-key attack on the full AES-256. In Section 5] we discuss new design criteria for block cipher key schedule as a consequence of our attack. Section 6 concludes the paper. In Appendix A we provide technical details about our differential trails.

Discussion. It is clear that the open key (chosen or known) security model is new and is still lacking a proper security definition. However we think that if one can support an open-key attack with a proof of security against such attack for the ideal cipher, this gives additional confidence that such property (in our case "differential multicollisions") should not be present in a good cipher. There are many constructions provably secure [37] in the ideal cipher model. This model assumes that both the key and the plaintext are accessible to the attacker. If a block cipher (e.g., AES) exhibits a property that should not appear in the ideal cipher then instantiation of a provably secure construction with this cipher could undesirably weaken the construction. Our Davies-Meyer example is exactly to show that a construction provably secure in the ICM can break
down if instantiated with AES-256 (such a hash function was never proposed for another reason - 128 -bit state is too short for a modern hash). The fact that this property does not automatically carry on to Davies-Meyer instantiated with all the other block-ciphers (hopefully), shows a non-trivial weakness of AES-256.

## 2 Multicollision Distinguisher

In this section we provide a chosen-key distinguisher for AES-256 which has practical complexity.

Definition 1. A set of two differences and $q$ pairs

$$
\left.\left\{\Delta_{K}, \Delta_{P} ;\left(P_{1}, K_{1}\right),\left(P_{2}, K_{2}\right), \ldots,\left(P_{q}, K_{q}\right)\right)\right\}
$$

is called a differential $q$-multicollision for a cipher $E_{K}(\cdot)$ if

$$
\begin{align*}
E_{K_{1}}\left(P_{1}\right) \oplus E_{K_{1} \oplus \Delta_{K}}\left(P_{1} \oplus \Delta_{P}\right)= & E_{K_{2}}\left(P_{2}\right) \oplus E_{K_{2} \oplus \Delta_{K}}\left(P_{2} \oplus \Delta_{P}\right)= \\
& =\cdots=E_{K_{q}}\left(P_{q}\right) \oplus E_{K_{q} \oplus \Delta_{K}}\left(P_{q} \oplus \Delta_{P}\right) \tag{1}
\end{align*}
$$

A differential $q$-multicollision can be also viewed as a set of $q$ right pairs with respect to the related-key differential, where the key is not fixed.

We compare the task of constructing a differential $q$-multicollision for an ideal cipher with that for AES-256. This task for an ideal cipher, i.e. a set of $2^{k}$ randomly chosen permutations, would require treating it as a black-box and making only encryption/decryption queries. We expect that for a good cipher with no (yet discovered) structural flaws, the task of constructing a differential $q$-multicollision would have the same complexity as for an ideal cipher.

Let us compute this complexity measured in the number of queries. Since the cipher is ideal, an adversary is only given an access to the encryption and decryption oracles, both having two inputs (a key and a plaintext/ciphertext) and one output. This is the same model of an adversary as in [2]3, p. 329].

In the beginning no triplet $\langle$ plaintext, key, ciphertext $\rangle$ is defined. Then, for each query of the adversary " $E_{K}(P)=$ ?" the encryption oracle takes a random value $C$ from a possible range (where $E_{K}(\cdot)$ is yet undefined) and thus defines $E_{K}(P)=C$. Also $E_{K}^{-1}(C)$ becomes defined. The same rule holds for a decryption query.

Lemma 1. To construct a differential $q$-multicollision for an ideal cipher with an n-bit block an adversary needs at least $O\left(q \cdot 2^{\frac{q-2}{q+2} n}\right)$ queries on the average.

Proof. See Sec. 2.1.
Remark 1. For small $q$, when the lower bound does not exceed $2^{n-1}$, a better estimate is obtained (see the proof of the lemma). In our case, for $n=128$, an adversary needs at least $q \cdot 2^{\frac{q-1}{q+1}} 128$ queries if $q \leq 57$.

Surprisingly, differential multicollisions for AES-256 can be constructed substantially faster. Furthermore, we can set $\Delta_{P}=0$ in a multicollision, so a stronger statement holds.

Theorem 1. A differential $q$-multicollision with $\Delta_{P}=0$ for AES-256 can be found with time complexity $q \cdot 2^{67}$.

Proof. See Sec. 2.2,
Thus for $q>3$ a differential $q$-multicollision for AES-256 can be constructed significantly faster than for an ideal cipher ${ }^{2}$ Therefore, AES-256 can not model an ideal cipher.

### 2.1 Proof of Lemma 1

Proof. Let $A$ be an adversary attacking the cipher, and assume that $A$ asks its oracles a total of $L$ queries, where $L<2^{n-1}$. Assume that a multicollision of the form (11) is found. Let us compute the probability of this event. First, we rewrite (11) as $U_{1}=U_{2}=\cdots=U_{q}$. With each term $U_{j}=E_{K_{j}}\left(P_{j}\right) \oplus E_{K_{j} \oplus \Delta_{K}}\left(P_{j} \oplus \Delta_{P}\right)$ we associate an integer $t_{j}$ such that $t_{j}$-th oracle query determines the value of $U_{j}$, i.e., computes the last (chronologically) element of the sum. Without loss of generality, assume that $t_{1}<t_{2}<\cdots<t_{q}$. Finally, define $t_{1}^{\prime}$ as the index of the query that determines the first element of the sum $U_{1}$.


Now compute for every $\left(t_{1}^{\prime}, t_{1}, t_{2}, t_{3}, \ldots, t_{q}\right)$ the probability that this set defines a differential $q$-multicollision. Before submitting $t_{i}$-th query, $i>1$, the following equation holds:

$$
U_{1}=U_{2}=\cdots=U_{i-1},
$$

where terms of $U_{1}, U_{2}, \ldots, U_{i-1}$ are completely determined by a tuple $\left(t_{1}^{\prime}, t_{1}\right.$, $\left.t_{2}, t_{3}, \ldots, t_{i-1}\right)$. Indeed, from $t_{1}^{\prime}$ and $t_{1}$ we define $K_{1}, \Delta_{K}, P_{1}, \Delta_{P}$; from $t_{j}$ we define $K_{j}$ and $P_{j}$.

Just before the moment $t_{i}$ only one term of $U_{i}$ is computed - w.l.o.g. let it be $E_{K_{i}}\left(P_{i}\right)$. Thus the equality $U_{i-1}=U_{i}$ should hold, i.e.

$$
U_{i-1}=E_{K_{i}}\left(P_{i}\right) \oplus \underbrace{E_{K_{i} \oplus \Delta_{K}}\left(P_{i} \oplus \Delta_{P}\right)}_{\text {queried at } t_{i}}
$$

[^1]By our definition, $t_{i}$ is the first moment when $E_{K_{i} \oplus \Delta_{K}}\left(P_{i} \oplus \Delta_{P}\right)$ is queried. Then either the decryption or the encryption oracle is called. In the first case the decryption oracle is called with a ciphertext $C$ and a key $K$, which for some $i$ should be equal to $K_{i} \oplus \Delta_{K}$. By the definition of $t_{i}$, the value $C$ is chosen from the set where $E_{K_{i} \oplus \Delta_{K}}(\cdot)$ is undefined. To become a part of a multicollision, there should exist $P_{i}$ such that $C=E_{K_{i}}\left(P_{i}\right) \oplus U_{i-1}$. On the other hand, after the decryption oracle is called, the following equation should hold:

$$
\begin{equation*}
E_{K_{i} \oplus \Delta_{K}}^{-1}(C)=P_{i} \oplus \Delta_{P} \tag{3}
\end{equation*}
$$

Since $L<2^{n-1}$, not more than $2^{n-1}$ texts were encrypted or decrypted with the key $K_{i} \oplus \Delta_{K}$. So the probability that (3) holds does not exceed $1 / 2^{n-1}$.

In the second case, let the encryption oracle be queried with a plaintext $P$ and a key $K$, which for some $i$ should be equal to $K_{i} \oplus \Delta_{K}$. For an answer $C$, a similar equation should hold:

$$
\begin{equation*}
C=U_{i-1} \oplus E_{K_{i}}\left(P_{i}\right) \tag{4}
\end{equation*}
$$

The same probability argument holds for this equation. Therefore, for every $i \geq 2$ we get a multiplier $2^{1-n}$ to the probability that a tuple $\left(t_{1}^{\prime}, t_{1}, t_{2}, t_{3}, \ldots, t_{q}\right)$ defines a differential $q$-multicollision. There are $\binom{L}{q+1}$ such tuples, each defining a differential $q$-multicollision with probability at max $2^{(q-1)(1-n)}$. We get the following equation for the number of queries required to get a $q$-multicollision with probability $1 / 2$ :

$$
\begin{equation*}
\binom{L}{q+1} \geq 2^{(q-1)(n-1)-1} \tag{5}
\end{equation*}
$$

Let us simplify the left part:

$$
\begin{align*}
\binom{L}{q+1}=\frac{L!}{(L-q-1)!(q+1)!} & =\frac{L(L-1) \cdots(L-q)}{(q+1)!} \leq \\
& \leq \frac{L^{q+1}}{(q+1)!} \leq \frac{L^{q+1}}{\frac{(q+1)^{q+1}}{e^{q+1}}}=\left(\frac{e L}{q+1}\right)^{q+1} \tag{6}
\end{align*}
$$

Substitute the result to (5):

$$
\begin{equation*}
\left(\frac{e L}{q+1}\right)^{q+1} \geq 2^{(q-1)(n-1)-1} \Rightarrow L \geq \frac{q+1}{e} 2^{\frac{q-1}{q+1}(n-1)-1}=O\left(q \cdot 2^{\frac{q-1}{q+1} n}\right) \tag{7}
\end{equation*}
$$

This is the bound for the number of queries needed to construct a multicollision with probability $1 / 2$. By Markov's inequality, the average number of queries exceeds this bound divided by two, so the right part of (7) is still a correct lower bound.

Now consider the case when $L \geq 2^{n-1}$. Let $\bar{K}$ be the set of keys such that there were more than $2^{n-1}$ encryption or decryption queries on each of these
keys. Define $l=|\bar{K}|$. If $l>q-2$ then $L$ exceeds $q \cdot 2^{n-2}$, which implies the statement of the lemma. If $l \leq q-2$ then there are at least $q-l$ sums $U_{i}$ in (2) that do not involve keys from $\bar{K}$. So if a $q$-multicollision has been found, then a $\frac{(q-l) \text {-multicollision has too been found such that it does not involve keys from }}{\underline{K}}$ $\bar{K}$. Then all the arguments on the probability of this event can be carried out from the first part of the proof.

Therefore, we gets the following inequality on $L$ :

$$
L \geq \frac{q-l+1}{e} 2^{\frac{q-l-1}{q-l+1}(n-1)-1}+l \cdot 2^{n-1} .
$$

For $l<q / 2$ we get the following:

$$
\begin{equation*}
L \geq \frac{q}{2 e} 2^{\frac{q-2}{q+2}(n-1)-1}=O\left(q \cdot 2^{\frac{q-2}{q+2} n}\right) \tag{8}
\end{equation*}
$$

For $l \geq q / 2$ we get

$$
\begin{equation*}
L \geq q \cdot 2^{n-2}=O\left(q \cdot 2^{n}\right) \tag{9}
\end{equation*}
$$

Equations (7), (8), and (9) complete the proof.

Remark 2. The function $F_{\Delta_{K}, \Delta_{P}}(K, P)=E_{K}(P) \oplus E_{K \oplus \Delta_{K}}\left(P \oplus \Delta_{P}\right)$ is a xor of two permutations. Patarin in [14 has shown that the xor of two random permutations can not be distinguished from a pseudo-random function with less than $2^{n}$ queries. In 15 it was proven that $q$-multicollision search for a random function requires at least $(q!)^{\frac{1}{q}} 2^{\frac{q-1}{q} n}$ effort. In our case we can not use this result since it assumes that $\Delta_{K}$ and $\Delta_{P}$ are fixed in advance while we allow an attacker to choose them during the attack.

### 2.2 Proof of Theorem 1

Here we construct a differential $q$-multicollision (11) with $\Delta_{P}=0$ in $q \cdot 2^{67}$ time. This is done in 5 steps:

1. Build a differential trail, which is efficient for the multicollision search.
2. Derive $\Delta_{K}$ from the trail.
3. Choose the active $S$-boxes, whose inputs will be fixed in the triangulation algorithm. Denote this set by $S$.
4. Run the triangulation algorithm and derive a set of free variables.
5. Produce $q$ pairs ( $P, K$ ) for (11) as follows:
(a) Assign inputs to S-boxes from $S$ with admissible values.
(b) Assign free variables randomly.
(c) Produce $(P, K)$.
(d) Check if $(P, K)$ and $\left(P, K \oplus \Delta_{K}\right)$ fit (11).

We expect that most of our readers are familiar with the description of AES and thus point out only main features of AES-256 that are crucial for our proof.

## Differential Trail

Notations. Differential trail (also called differential characteristic) is a sequence of differences in all the internal states of the cipher and all the subkeys. If we distinguish only between zero and non-zero (byte) differences, we call such a trail a trail with truncated differences.

We denote the subkey of round $i$ by $K^{i}$, i.e. the first (whitening) subkey is $K^{0}$, the subkey of round 1 is $K^{1}$, etc., the last subkey is $K^{14}$. The difference in $K^{i}$ is denoted by $\Delta K^{i}$. Bytes of a subkey are denoted by $K_{i, j}^{l}$, where $i$ stands for the row index, and $j$ stands for the column index in the standard matrix representation of AES. Bytes of the plaintext are denoted by $P_{i, j}$, and bytes of the internal state after the SubBytes transformation in round $r$ are denoted by $A_{i, j}^{r}$. Let also $B_{i, j}^{r}$ denote a byte in position $(i, j)$ after the $r$-th application of MixColumns.

Features of $A E S$-256. AES-256 has 14 rounds and a 256-bit key, which is two times larger than the internal state. Thus the key schedule consists of only 7 rounds. One key schedule round consists of the following transformations:

$$
\begin{array}{lrr}
K_{i, 0} \leftarrow S\left(K_{i+1,7}\right) \oplus K_{i, 0} \oplus C_{r}, & & 0 \leq i \leq 3 ; \\
K_{i, j} \leftarrow K_{i, j-1} \oplus K_{i, j}, & 0 \leq i \leq 3, & 1 \leq j \leq 3 ; \\
K_{i, 4} \leftarrow S\left(K_{i, 3} \oplus K_{i, 4},\right. & & 0 \leq i \leq 3 ;  \tag{10}\\
K_{i, j} \leftarrow K_{i, j-1} \oplus K_{i, j}, & 0 \leq i \leq 3, & 5 \leq j \leq 7,
\end{array}
$$

where $S()$ stands for S -box, and $C_{r}$ - for the round-dependant constant. Therefore, every round has 8 S-boxes.

Weakness in the key schedule. Two features of the key schedule help us to build a good differential trail. First, the key schedule has a slow diffusion in backward direction. It means that a difference in a single byte $K_{0,0}$ will propagate to only two bytes, $K_{0,0}$ and $K_{0,1}$, if we apply the inversion of the key schedule round. The next inverted round will affect only one more byte, etc. Thus we can build a trail with a low-weight difference in key schedule if we start with a low-weight difference in the last round and then step backwards.

The second feature is unique to AES256 due to its "key size" /"state size" ratio, $N_{k} / N_{b}=2$. We can inject "good"


Fig. 2. A local collision values with the first part of the key and then cancel them, after they pass the round, with the second part of the key. We call this a local collision (Fig. 2).

Constructing a trail. Step by step, we construct a differential trail from the last rounds to the first ones. The trail is described in details in Appendix A, both in a truncated form and with the actual differences given. The trail has 41 active S-boxes, and 5 of them are in the key schedule.

Search for a solution. After the trail has been defined, we produce a pair $(P, K)$ which with pair $\left(P, K+\Delta_{K}\right)$ fits the trail and thus is a part of a differential multicollision (1) since all such pairs have the same difference in the ciphertext. Now note that the trail explicitly states all the non-zero input $\delta_{I}$ and output $\delta_{O}$ differences of the S-boxes. According to the S-box properties, there are at most 4 solutions of the equation $S\left(x \oplus \delta_{I}\right) \oplus S(x)=\delta_{O}$. This set of solutions we call admissible inputs. Therefore, $(P, K)$ and $\left(P, K+\Delta_{K}\right)$ fit the trail if and only if in the execution $E_{K}(P)$ all active S-boxes get admissible inputs. In the next paragraphs we explain how to construct such executions efficiently.

## Triangulation Algorithm

Search for free variables. The triangulation algorithm was proposed in [9] as a tool for solving systems of non-linear equations, which appear in differential attacks. Given the constraints on the internal variables, the algorithm outputs a special set of variables, called free variables. These free variables can be assigned randomly; and this assignment together with pre-fixed variables completely and efficiently determines the whole execution. The fewer variables are fixed the better the algorithm works.

Our goal is to efficiently produce the cipher executions in which active S-boxes get admissible values. However, the algorithm can not process all the active Sboxes of our trail since they are positioned too far from one another.

We found that we can fix inputs to 30 out of 41 S-boxes: all the active S boxes in the internal states of rounds 1-4 and all the 5 active S-boxes in the key scheduling. The triangulation algorithm outputs 18 free variables, out of these 11 are in the key and provide freedom in the choice of the key for the distinguisher. These free variables are listed below.

| Key | Internal state |
| :---: | :---: |
| $K_{2,0}^{1}, K_{3,0}^{1}, K_{3,1}^{1}, K_{0,1}^{2}, K_{0,1}^{3}$ | $A_{0,1}^{2}, B_{1,1}^{2}, B_{1,2}^{2}$ |
| $K_{2,0}^{3}, K_{3,0}^{3}, K_{2,1}^{3}, K_{0,1}^{4}, K_{0,1}^{5}, K_{3,3}^{5}$ | $B_{2,3}^{2}, B_{3,3}^{2}, A_{1,1}^{3}, A_{1,2}^{3}$ |

Constructing a pair. Having assigned 18 free variables randomly and 30 S-box inputs with an admissible value, we substitute these values to the equations, which have been ordered by the triangulation algorithm. One by one, all the variables are determined and thus a pair $(P, K)$ is defined. It fits the trail if and only if the 11 S -boxes not covered by the triangulation algorithm get admissible values as inputs. For the S -box in round 6 only 2 values are admissible so the probability is $2^{-7}$ while for the other 10 S -boxes 4 values are admissible. This
results in the overall probability $2^{-(7+10 \cdot 6)}=2^{-67}$. Thus out of $2^{67}$ pairs one fits the trail on average. We wrote a program and checked that the distribution of the pairs is random enough so the probability estimates are likely to be correct. We also checked experimentally that bottom 7 rounds of AES produce expected difference after $2^{30}$ pairs on the average, exactly as predicted by the trail.

The complexity of the attack. Recall the scheme of the attack, which was given in the beginning of the proof. The first four steps are precomputations and actually have negligible cost. The triangulation algorithm works less than a second. The last step requires only to substitute the values into the equations one by one, which is computationally equivalent to a single encryption. Thus to get a right pair we need about $2^{67}$ operations each equivalent to one encryption. The attack needs negligible memory and is fully parallelizable.

### 2.3 Practical Distinguisher

The definition of differential $q$-multicollision can be further relaxed if we allow arbitrary difference at some byte positions of $\Delta_{P}, \Delta_{K}$ or $\Delta_{C}$. Although an attacker gets more freedom, finding such a construction for an ideal cipher becomes easier as well. To get a lower bound, only a slight modification of Lemma 1 is required.

For the 13 rounds of AES-256 the complexity of finding this type of differential 5 -multicollision, with fixed difference in 14 bytes of the plaintext and fixed ciphertext difference can be lower bounded by $2^{\frac{4.112}{6}}=2^{74.6}$ computations. For the full AES-256 a differential 10-multicollision with half of the plaintext difference fixed, and the fixed ciphertext difference the lower bound is $2^{\frac{9.64}{11}}=2^{52.3}$ computations. Note that these lower bounds are far from being tight. In practice we expect an efforts of $2^{112}$ and $2^{64}$ for finding each extra collision for 13 and 14 rounds of AES-256 respectively, since the differences are structured and fixed.

At the same time we can do it much faster, in just $q \cdot 2^{37}$ in both cases, which allowed us to compute these distinguishers in several hours on a PC. The actual values will be given in the extended version of this paper. The core of a practical distinguisher is a multicollision trail (Figure 5), where the behavior of S-boxes in the first two rounds is not restricted. Computing from the middle, we get 14 bytes with fixed difference before the second round, and 8 bytes with fixed difference before the first round. The triangulation algorithm covers all but six active S-boxes in rounds $3-14$ so that we find a (partial) $q$-multicollision with complexity $q \cdot 2^{37}$.

## 3 Pseudo-collisions for AES-Based Hashing

The Davies-Meyer mode of blockcipher-based hashing has been proven collisionresistant if instantiated by an ideal cipher [3]. In this section we show a similar proof in the ideal-cipher model for the $q$ pseudo-collision resistance, when differences in the IV and the message $\left(\Delta_{I}, \Delta_{M}\right)$ are fixed. We then show that it
is relatively easy to find $q$ pseudo-collisions for AES-256 in the Davies-Meyer mode. We also point out that we construct one-block pseudo-collisions and thus the technique of Joux [8] does not apply here.

Our goal is for fixed differences $\Delta_{I}, \Delta_{M}$ to find many pseudo-collisions for the $H^{E}(I, M) \stackrel{\text { def }}{=} E_{M}(I) \oplus I$ which is the Davies-Meyer compression function with AES-256 as the underlying cipher. Here $I$ is the 128 -bit IV, and $M$ is a 256 -bit message block. A pseudo-collision satisfies the following equality:

$$
H^{E}(I, M)=H^{E}\left(I \oplus \Delta_{I}, M \oplus \Delta_{M}\right)
$$

Let us rewrite it:

$$
\begin{align*}
& E_{M}(I) \oplus I=E_{M \oplus \Delta_{M}}\left(I \oplus \Delta_{I}\right) \oplus I \oplus \Delta_{I} \Leftrightarrow \\
& \tag{11}
\end{align*} \quad \Leftrightarrow E_{M}(I) \oplus E_{M \oplus \Delta_{M}}\left(I \oplus \Delta_{I}\right)=\Delta_{I} .
$$

While finding many pseudo-collisions with different $\Delta_{I}, \Delta_{M}$ can be done using the birthday paradox, the same task for fixed $\Delta_{I}, \Delta_{M}$ is hard. This problem can be expressed as finding a solution of

$$
\begin{align*}
E_{M_{1}}\left(P_{1}\right) \oplus E_{M_{1} \oplus \Delta_{M}}\left(P_{1} \oplus \Delta_{I}\right) & =E_{M_{2}}\left(P_{2}\right) \oplus E_{M_{2} \oplus \Delta_{M}}\left(P_{2} \oplus \Delta_{I}\right)=\cdots= \\
& =E_{M_{q}}\left(P_{q}\right) \oplus E_{M_{q} \oplus \Delta_{M}}\left(P_{q} \oplus \Delta_{I}\right)=\Delta_{I} \tag{12}
\end{align*}
$$

which is harder than finding a differential $q$-multicollision for $E_{K}(\cdot)$, because the ciphertext difference is unrestricted in (1). Therefore, Lemma 1 gives us a lower bound on the complexity of this attack.

Corollary 1. To construct $q$ pseudo-collisions (12) with fixed $\Delta_{I}, \Delta_{M}$ for an ideal cipher with an $n$-bit block an adversary needs at least $O\left(q \cdot 2^{\frac{q-2}{q+2} n}\right)$ queries on average.

Theorem 2. For AES-256 in the Davies-Meyer mode q pseudo-collisions (12) can be found in time $q \cdot 2^{67}$.

Proof. Pseudo-collision attack on the Davies-Meyer mode requires the difference in the plaintext $P$ to be equal to the difference in the ciphertext $C$. The DaviesMeyer feed-forward would then cancel this difference. Our differential trail needs only to be slightly modified for this purpose. The first round of the new trail is shown in Fig. 3 (the actual values are given in Appendix A); the other rounds are the same.

The resulting trail has 41 active S-boxes. The triangulation algorithm covers the same active S-boxes as in the proof of Theorem 1 and outputs 18 free variables. The complexity is thus the same.

Corollary 2. AES-256 can not be used to instantiate the Davies-Meyer hashing mode.

We expect that similar results can be shown for the other blockcipher-based constructions which are provably secure in the ideal-cipher model 3 .


Fig. 3. The first round of the differential trail for the attack on AES-256 in the DaviesMeyer mode

## 4 Related-Key Attack on AES-256

In this section we demonstrate two results: a related-key distinguisher and a key-recovery attack based on this distinguisher for the full AES-256.

### 4.1 Distinguisher

The first two rounds of the multicollision trail can easily be modified to build a related-key distinguisher with relatively few active S-boxes - see Appendix A for the details of the trail. The resulting trail has 19 active S-boxes in the internal states. The difference propagates through 14 S-boxes with probability $2^{-6 \cdot 14}$, and through the remaining five with prob. $2^{-7 \cdot 5}$. Therefore, we get a distinguisher with probability $2^{-(14 \cdot 6+5 \cdot 7)}=2^{-119}$. However, the trail has five additional active S-boxes (each with probability $2^{-7}$ ) in the key schedule. As a result, the distinguisher works for 1 out of $2^{35}$ related-key pairs on the average.

### 4.2 Key Recovery

There are several ways how our trail can be used for the full 256 -bit key recovery. Here we present one possibility. Steps of the attack are illustrated in Fig. 4] and are described below.

First step. We change the trail (see Fig. 4 or Appendix) to get more active S-boxes in the first two rounds which allows us to recover the key bytes at the entrance to these S-boxes. A new trail has eight active S-boxes in the first round in a "checkerboard pattern" and two in the second round. Our goal is to find ten key bytes $K_{0,0}^{0}, K_{0,2}^{0}, K_{1,1}^{0}, K_{1,3}^{0}, K_{2,0}^{0}, K_{2,2}^{0}, K_{3,1}^{0}, K_{3,3}^{0}, K_{0,0}^{1}, K_{0,2}^{1}$ so we execute the following procedure for each of the $2^{35}$ related-key pairs:

1. Repeat $2^{31}$ times:
(a) Compose two structures of $2^{64}$ plaintexts as specified below.
(b) Encrypt the 1st structure with $K$ and the 2nd with $K \oplus \Delta_{K}$.
(c) Sort the ciphertexts and check for a pair with the difference $\Delta_{C}$.
(d) Save a good pair if it is found.
2. For each candidate pair derive $2^{16}$ variants for the ten key bytes (see details below).
3. Pick the key candidate with the maximal number of votes.

The overall complexity of this procedure is $2^{31+65}=2^{96}$ data and time, and $2^{64}$ memory, and it finds 80 bits of the key.

Each structure has all the possible values in bytes $\left\{P_{i, j},(i+j) \bmod 2=0\right\}$ (these bytes line up in columns 0 and 2 after the ShiftRows of the 1st round). The other bytes are random constants with the constraint that the two structures are related by $\Delta_{P}$ as in the distinguisher. For a fixed key, each structure contains $2^{64}$ pairs with the proper differences after the S-boxes of the 2nd round (yellow differences in two bytes). The remaining active S-boxes in rounds $3-14$ require probability $2^{-93}$ for the trail to be fulfilled. Each structure produces a right pair of ciphertexts with probability $2^{64-93}=2^{-29}$. Thus $2^{31}$ structures produce on average 4 right pairs and $2^{31+128-128}=2^{31}$ wrong pairs.

Due to the uniform differential properties of AES S-boxes each active S-box for which we know input and output differences would suggest to us two candidates for the key byte of this S-box. For a candidate pair we guess two byte differences $\sqrt[3]{3}$ at $B_{0,0}^{1}$ and $B_{0,2}^{1}$. We know the difference in the remaining three bytes of each column, since they should cancel out with the differences coming from the key. Thus we can undo the MixColumns for each column which allows us to know the output differences of eight S-boxes of the 1st round. We know the input differences for these S-boxes from the plaintext. In addition the two guessed difference bytes serve as inputs to the two active S-boxes of the 2 nd round. The output differences for these S-boxes are known from the trail. Thus we have 10 S -boxes for which we know input and output differences which gives us $2^{10} \cdot 8 \cdot 8=2^{16}$ possibilities for 80 -bits of the key per candidate pair. The $2^{31}$ wrong pairs would suggest $2^{16} \cdot 2^{31}=2^{47}$ random keys, while the four good ones would all vote for the correct 80 -bit key and some random keys. No wrong key guesses survive this step and we get 80 key bits as a result.

Second step. We proceed with changing top rounds of the trail to derive other key bytes. We remove the $2^{-6}$ condition on the input to the active $S$-box $(0,0)$ of the 3 rd round. Then we get five active S-boxes in the second round and 16 active S-boxes in the first round. We prepare $2^{90}$ pairs with the ciphertext difference as in the trail and decrypt them. We will try to detect pairs $\left(2^{90-87}=8\right.$ on the average) that pass the conditions in rounds $3-14$. We partially encrypt all the resulting plaintext pairs using the known 80-bits and check whether the columns $\Delta B_{*, 0}^{1}$ and $\Delta B_{*, 2}^{1}$ follow the trail. This is a 48-bit filter on the pairs and thus we are left with $2^{42}$ candidate pairs. Then we guess the differences in bytes $B_{1,1}^{1}$ and $B_{3,3}^{1}$ (eight possibilities due to impossible input/output constraints in five

[^2]

Fig. 4. Key recovery steps
corresponding S-boxes of the 1 st and 2 nd rounds), and undo the MixColumns 4 . As a result, we have $2^{42+6}=2^{48}$ key candidates from all pairs counting on a $64+16=80$ bit key ( 64 from $K^{0}$ and 16 from bytes $K_{1,1}^{1}, K_{3,3}^{1}$ ). No wrong key guesses survive this step. We find the remaining 12 bytes of $K^{1}$ by exhaustive search in $2^{96}$ steps. The total complexity of the attack is thus dominated by the $2^{96+35}=2^{131}$ complexity of the first step.

Second step with TA. If we want to avoid the chosen-ciphertext framework, there is a way to combine the knowledge of 80 key bits on top and 35 key bits in the middle with a bit higher complexity. The problem is that the fixed bits are positioned far from one another (3 key schedule rounds apart), so it seems hard to make an efficient exhaustive search on the remaining part of the key.

We solve this problem by running the triangulation algorithm on the key schedule only, where 15 bytes are marked as fixed. The algorithm outputs 17 bytes in different subkeys as free variables. Then we assign these variables randomly, choose admissible values for the remaining ones, and thus define the key guess. There are $2^{17 \cdot 8+5}=2^{141}$ possible assignments, which would determine the complexity of the key recovery.

## 5 New Design Criteria for Block Ciphers

Our results imply new design criteria for the key schedules of the block ciphers. First of all, local collisions should be prevented. Although it is usually easy to

[^3]arrange a local collision in one round, there should be no good patterns for several rounds. A slow diffusion in the AES-256 key schedule helped us to concatenate many local collisions into a differential trail for the whole cipher. This should not be possible in a good cipher.

The key schedule should be also desynchronized with respect to the internal state. The active injection bytes in our AES trails are always located in the first row, which is not rotated. Therefore, the differences to be cancelled in a local collision should be located in the same column, or exactly 4 columns to the right in the subkey. This shift is preserved by the key schedule round, which should be certainly avoided.

Slow diffusion in the key schedule makes it also vulnerable to the triangulation algorithm or similar tools. Our preliminary analysis shows that the triangulation algorithm can cover up to two times the number of rounds needed for the full diffusion. If the key schedule in AES was ideal, we would be able to solve systems of equations on at most three-four rounds, while now we can attack five.

Finally, if one considers known or chosen key attacks as a threat he needs to add two extra full diffusions on top of the number of rounds secure against standard statistical attacks.

## 6 Conclusions

In this paper we show a chosen key distinguisher for the 256-bit key AES with almost practical complexity of $q \cdot 2^{67}$ queries and negligible memory. It was verified by computing partial $q$ multicollisions in time $q \cdot 2^{37}$ which takes several hours on a PC. We also show the first related-key attack on the full AES-256 with $2^{96}$ data and time complexity and $2^{65}$ memory which works for 1 out of every $2^{35}$ keys on average.

Acknowledgements. The authors wish to thank anonymous reviewers for their valuable comments. We also wish to thank Christian Rechberger for useful discussions. Dmitry Khovratovich is supported by PRP "Security \& Trust" grant of the University of Luxembourg. Ivica Nikolić is supported by the Fonds National de la Recherche Luxembourg grant TR-PHD-BFR07-031.

## References

1. Biham, E., Dunkelman, O., Keller, N.: Related-key boomerang and rectangle attacks. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 507-525. Springer, Heidelberg (2005)
2. Black, J.: The ideal-cipher model, revisited: An uninstantiable blockcipher-based hash function. In: Robshaw, M.J.B. (ed.) FSE 2006. LNCS, vol. 4047, pp. 328-340. Springer, Heidelberg (2006)
3. Black, J., Rogaway, P., Shrimpton, T.: Black-box analysis of the block-cipher-based hash-function constructions from PGV. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 320-335. Springer, Heidelberg (2002)
4. Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. J. ACM 51(4), 557-594 (2004)
5. Daemen, J., Rijmen, V.: The Design of Rijndael. AES - the Advanced Encryption Standard. Springer, Heidelberg (2002)
6. Ferguson, N., Kelsey, J., Lucks, S., Schneier, B., Stay, M., Wagner, D., Whiting, D.: Improved cryptanalysis of Rijndael. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 213-230. Springer, Heidelberg (2000)
7. Jaulmes, E., Joux, A., Valette, F.: On the security of randomized cbc-mac beyond the birthday paradox limit: A new construction. In: Daemen, J., Rijmen, V. (eds.) FSE 2002. LNCS, vol. 2365, pp. 237-251. Springer, Heidelberg (2002)
8. Joux, A.: Multicollisions in iterated hash functions. Application to cascaded constructions. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 306-316. Springer, Heidelberg (2004)
9. Khovratiovich, D., Biryukov, A., Nikolić, I.: Speeding up collision search for byteoriented hash functions. In: CT-RSA 2009. LNCS, vol. 5473, pp. 164-181. Springer, Heidelberg (2009)
10. Kim, J., Hong, S., Preneel, B.: Related-key rectangle attacks on reduced AES-192 and AES-256. In: Biryukov, A. (ed.) FSE 2007. LNCS, vol. 4593, pp. 225-241. Springer, Heidelberg (2007)
11. Knudsen, L.R., Rijmen, V.: Known-key distinguishers for some block ciphers. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 315-324. Springer, Heidelberg (2007)
12. National Institute of Standards and Technology (NIST). FIPS-197: Advanced Encryption Standard (November 2001), http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf
13. National Security Agency (NSA). National Policy on the Use of the Advanced Encryption Standard (AES) to Protect National Security Systems and National Security Information (June 2003), http://www.cnss.gov/Assets/pdf/cnssp_15_fs.pdf
14. Patarin, J.: A proof of security in $O\left(2^{n}\right)$ for the xor of two random permutations. In: Safavi-Naini, R. (ed.) ICITS 2008. LNCS, vol. 5155, pp. 232-248. Springer, Heidelberg (2008)
15. Suzuki, K., Tonien, D., Kurosawa, K., Toyota, K.: Birthday paradox for multicollisions. IEICE Transactions 91-A(1), 39-45 (2008)

## A Details on Trails

The actual differences. The actual difference values should satisfy conditions imposed by the S-box properties, the presence of MC-columns in the subkeys, etc. We also require that differences in rounds 5 and $7-13$ propagate through S-boxes with maximal probability - $2^{-6}$. We programmed the search for the actual values and validated that those conditions do not lead to contradictions. The search has a negligible cost. We thus have defined the differences $\Delta_{K}$ (determined by subkeys $K^{0}$ and $K^{1}$ ) and $\Delta_{C}$ (determined by the subkey $K^{14}$ ). In the following subsections we list the actual differences in the trails.

## Davies-Meyer trail (the first two rounds)

|  | Plaintext |  | Subkey | Ciphertext |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 01010101 |  | Of 0e 0f 0e | 01010101 |  |  |  |
| 0 | 00000000 |  | 07070707 | 00000000 |  |  |  |
| 0 | 00000000 |  | 07070707 | 00000000 |  |  |  |
|  | 00000000 |  | 09090909 | 00000000 |  |  |  |
| i | After SB | After MC | Subkey | i | After SB | After MC | Subkey |
|  | 30 Of 44 b 0 | 6500 2b 00 | 37003700 |  | 1b 000700 | 0c 000 e 00 | Of 010 e 00 |
|  | 7c b5 9308 | 1f 2c 1f 00 | 1f 00 1f 00 | 2 | 00120000 | 07000700 | 07000700 |
|  | 78 d6 c2 57 | 1f 00 e 200 | If 00 1f 00 |  | 0000 1a 00 | 07000700 | 07000700 |
|  | e7 c3 2903 | 21002133 | 21002100 |  | 00000016 | 09000900 | 09000900 |

Related-key distinguisher (the first two rounds)

| i | Plaintext |  | Subkey | Ciphertext |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0 e 0 e 0 e 0 e |  | Of 0e 0f 0e | 01010101 |  |  |
| , | 07070707 |  | 07070707 | 00000000 |  |  |
| - | 07070707 |  | 07070707 | 00000000 |  |  |
|  | 09090909 |  | 09090909 | 00000000 |  |  |
| i | After SB | After MC | Subkey | i After SB | After MC | Subkey |
|  | 1f 00 1f 00 | 3 e 003 e 00 | 37003700 | 07000700 | 0e 000 e 00 | Of 010 e 00 |
| 1 | 00000000 | 1f 00 if 00 | 1f 00 1f 00 | 00000000 | 07000700 | 07000700 |
| 1 | 00000000 | 1f 00 1f 00 | 1f 00 1f 00 | $2{ }^{2} 00000000$ | 07000700 | 07000700 |
|  | 00000000 | 21002100 | 21002100 | 00000000 | 09000900 | 09000900 |

Note on the colors. Differential trails are given in our figures in a truncated form, we marked distinct difference values with different colors. The reader does not need to care about the actual values in order to understand how the trail is constructed. However all the trail differences are provided in the tables of this Appendix.

The white cell stands for zero difference in a byte, the non-white cells stand for the non-zero differences. The same colors mean the same values except for the green, which denotes arbitrary differences. The exact relation between the colors and the values can be derived from the list of the actual differences. Grey and blue columns stand for MC-columns. In a spoiled MC-column one byte is marked with another color.

## Differential trail for finding multicollisions:

Table 2. Multicollision trail

|  | Plaintext |  | Subkey | Ciphertext |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | $\begin{aligned} & \hline 00 \quad 00 \quad 00 \quad 00 \\ & 00 \quad 0000 \quad 00 \\ & 00 \quad 0000 \quad 00 \\ & 00 \quad 00 \quad 00 \quad 00 \end{aligned}$ |  | Of 0 e 0f 0 e  <br> 07 07 07 <br> 07 07  <br> 07 07 07 <br> 09 07  <br> 09 09 09 | $\begin{aligned} & \hline 01010101 \\ & 00000000 \\ & 000000 \quad 00 \\ & 00 \quad 00 \quad 00 \quad 00 \end{aligned}$ |  |  |  |
|  | After SB | After | Subkey |  | After SB | After MC |  |
| 1 | 305 c e1 b0 7c b5 ed 72 a6 d6 c2 16 82 eb 2903 | 6500 02 00 <br> 1f 25 $1 f$ 00 <br> $1 f$ 00 e2 <br> 21 00  <br> 21 00 21 <br> 33   | 37003700 <br> 1f 00 1f 00 <br> 1f 00 1f 00 <br> 21002100 | 2 | $\begin{array}{llll} \hline 1 \mathrm{~b} & 00 & 07 & 00 \\ 00 & 12 & 00 & 00 \\ 00 & 00 & 1 \mathrm{a} & 00 \\ 0 & 00 & 00 & 16 \end{array}$ | Oc 00 0 e 00  <br> 07 00 07 00 <br> 07 00 07 00 <br> 09 00 09 00 | $\begin{array}{llll} \hline \text { Of } 010 \mathrm{e} & 00 \\ 07 & 00 & 07 & 00 \\ 07 & 00 & 07 & 00 \\ 09 & 00 & 09 & 00 \end{array}$ |
| 3 | $\begin{array}{llll} \hline \text { 1f 1f } & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | 3e 3 e 0000 <br> 1f 1f 00 <br> 1f <br> 1f 00 <br> 21 <br> 21 <br> 21 <br> 00 <br> 00 | 37 370000 <br> 1f 1f 0000 <br> 1f 1f 0000 <br> 21 <br> 21 <br> 1 0000 | 4 | $\begin{array}{llll} \hline 07 & 07 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 0 & 00 & 00 & 00 \end{array}$ | Oe 0 e 00 00 <br> 07 07 00 00 <br> 07 07 00 00 <br> 09 09 00 00 | Of 0 e 0000 07070000 07070000 09090000 |
| 5 | $\begin{array}{llll} \hline \text { 1f } & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ |  | 37 00 00 <br> 1f 00 00 <br> 1f 00  <br> 1f 00 00 <br> 21 00 00 <br> 00   | 6 | $\begin{array}{llll} \hline 07 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | 0 e 00 00 00 <br> 07 00 00 00 <br> 07 00 00 00 <br> 09 00 00 00 | $\begin{array}{lll} \hline \text { Of } 010101 \\ 07 & 00 & 00 \\ 00 \\ 07 & 00 & 00 \\ 00 & 00 & 00 \end{array}$ |
| 7 | $\begin{array}{llll} \hline \text { 1f } & \text { If } & \text { 1f } & \text { if } \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | 3 e 3 e 3 e 3 e 1f 1f 1f 1f 1f if 1f 1f 21 21 |  | 8 | $\begin{array}{llll} 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | $\left.\begin{array}{\|llll} \hline 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array} \right\rvert\,$ | $\begin{array}{llll} 01 & 00 & 01 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ |
| 9 | $\begin{array}{lllll} \hline \text { 1f } & 00 & \text { lf } & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ \hline \end{array}$ |  | 3 e 00 3 e 00 <br> 1f 00 $1 f$ 00  <br> 1f 00 $1 f$ 00 <br> 21 00 21 00 | 10 | $\begin{array}{lllll} \hline 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ \hline \end{array}$ | 00 00 00 00 <br> 00 00 00 00 <br> 00 00 00 00 <br> 00 00 00 00 | $\begin{array}{llll} 01 & 01 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ |
| 11 | $\begin{array}{llll} \hline \text { 1f } & \text { 1f } & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | 3e 3 e 0000 <br> 1f 1f 00 <br> 1f <br> 1f 00 <br> 21 <br> 21 <br> 21 <br> 000 <br> 00 | 3e 3 e 0000 <br> 1f 1f 00 <br> 1f 1f 00 <br> 1f <br> 21 <br> 21 <br> 21 <br> 000 | 12 | $\begin{array}{llll} \hline 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 0 & 00 & 00 & 00 \end{array}$ | $\left.\begin{array}{\|lllll} \hline 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array} \right\rvert\,$ | $\begin{array}{llll} 01 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ |
| 13 | $\begin{array}{\|cccc} \hline \text { If } 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | 3e 000000 1f 000000 1f 000000 21000000 | 3e 00 00 00  <br> 1f 00 00 00  <br> 1f 00 00 00  <br> 21 00 00 00 | 14 | $\begin{array}{lllll} \hline 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ | $\begin{array}{\|lllll} \hline 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ \hline \end{array}$ | $\begin{array}{llll} 01 & 01 & 01 & 01 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \\ 00 & 00 & 00 & 00 \end{array}$ |



Fig. 5. Multicollision trail. The actual values are given in Table 2


[^0]:    ${ }^{1}$ We use colors in the diagrams of the trails, so please refer also to the tables in the appendix if you print this paper on a black and white printer.

[^1]:    ${ }^{2}$ Moreover, even for $q=3$ we are not aware of any algorithm faster than $2^{2 n / 3}$.

[^2]:    ${ }^{3}$ For each byte difference there are only 8 possibilities that would not contradict with the five known differences in the 1st and 2 nd rounds.

[^3]:    ${ }^{4}$ We use the knowledge of the key byte $K_{0,0}^{1}$ to find the differences in the green diagonal at the 2 nd round.

