

New Limits for AES Known-Key Distinguishers

Lorenzo Grassi¹ and Christian Rechberger^{1,2}

¹ IAIK, Graz University of Technology, Austria

² DTU Compute, DTU, Denmark

`firstname.lastname@iaik.tugraz.at`

Abstract. Known-key distinguishers have been introduced to better understand the security of block ciphers in situations where the key can not be considered to be secret.

AES is often considered as a target of such analyses, simply because AES or its building blocks are used in many settings that go beyond classical encryption. The most recent known-key model of Gilbert (proposed at Asiacrypt 2014) allows to consider two 4-round distinguishers combined in an inside-out fashion (8 core rounds), and to extend it by one round in each direction (two extension rounds). The resulting 10-round distinguisher has a time complexity of 2^{64} . In that work, arguments were put forward suggesting that two extension rounds seems to be the limit in the known-key model, and that likely only a distinguisher that exploits the balance property can be extended in such way.

In this paper we disprove both these conjectures and arrive at the following results. We firstly show that the technique proposed by Gilbert can also be used to extend a known-key distinguisher based on truncated differential trails. This allows to improve all the known-key distinguishers currently present in literature for AES up to 10 rounds of AES. In particular, we are able to set up a 9-round known-key distinguisher for AES with a time complexity of 2^{23} and a 10-round known-key distinguisher with a time complexity of 2^{50} . Secondly we are also able to show that more than two extension rounds are possible. As a result of this, we describe the first known-key distinguishers on 12 rounds of AES, by extending an 8-round known-key distinguisher by two rounds in each direction (four extension rounds). The time complexity is 2^{82} .

We conclude with a discussion on why it seems not feasible to set up similar distinguishers on 14 rounds exploiting the same strategy.

Keywords: Block cipher, Permutation, AES, Known-Key Distinguisher

Table of Contents

1	Introduction.....	2
1.1	Known-Key Distinguishers for AES: the State of the Art	3
1.2	Our Main Results	3
2	Preliminary - Description of AES	5
3	Subspace trails	6
3.1	Subspace trails of AES	6
4	Known-Key Distinguishers for AES	8
4.1	Definition of Known-Key Distinguisher	8
4.2	7- and 8-Round Known-Key Distinguisher	11
4.3	Multiple Limited-Birthday 8-Round Known-Key Distinguisher ..	12
4.4	10-Round Gilbert's Known-Key Distinguisher	12
5	Key-Recovery Extensions using Truncated Differentials	15
5.1	Attack for the Case of 1-Round Extension	15
5.2	Attack for the Case of 2-Round Extension	16
6	9-Round Known-Key Distinguisher for AES	17
6.1	The Computational Cost of Generic Player is Not Negligible! ...	21
7	10-Round Distinguisher of AES - Full AES-128	21
7.1	Independent Subkeys: No Key Schedule	22
7.2	The Key Schedule Case	23
8	12-Round Distinguisher of AES	25
9	Infeasibility of a 14-round Known-Key Distinguisher.....	28
10	Conclusion, Discussion, and Open Problems	28
A	Considerations on Gilbert's 10-round Distinguisher	30
B	A possible Variant of Gilbert's Distinguisher - Details	31
C	The Rebound Attack - Details	32
D	Known-Key Distinguishers for 7- and 8-Round of AES based on the Square Property	34
D.1	The 8-round Known-Key Distinguisher based on the Balance Property - A Formal Description.....	36
E	Details of Known-Key Distinguisher when the Computational Cost of the Generic Player is Considered.....	38
E.1	Known-Key Distinguisher on 9-Round AES	38
E.2	Known-Key Distinguisher on 10-Round AES with Key Schedule ..	40
F	New 7-, 8- and 9-round Known-Key Distinguishers for AES	40
F.1	7-Round Known-Key Distinguisher	41
F.2	8-Round Known-Key Distinguisher	41
F.3	9-Round Known-Key Distinguisher	42
F.4	Consideration and Comparison with Gilbert's Distinguisher.....	44
G	Proof of Proposition 1 - Sect. 7.2.....	44
H	Gilbert's Known-Key Distinguisher on 12 Rounds	45

H.1	Considerations about the Computational Cost of the Verifier - Gilbert's 10-round Distinguisher	45
H.2	Known-Key Distinguisher on 12 Rounds Based on the Balance Property	47
I	The Herds Attack	49

1 Introduction

Block ciphers play an important role in symmetric cryptography providing the basic tool for encryption. They are the oldest and most scrutinized cryptographic tools. Consequently, they are the most trusted cryptographic algorithms that are often used as the underlying tool to construct other cryptographic algorithms, whose proofs of security are performed under the assumption that the underlying block cipher is ideal.

The concept of known-key distinguishers was introduced by Knudsen and Rijmen in [13]. In the classical single secret-key setting, the attacker does not know the randomly generated key and aims to recover it or builds a (secret-key) distinguisher that allows to distinguish the cipher from a random permutation. The security model in known-key attacks is quite different though: the attacker knows the randomly drawn key the block cipher operates with and aims to find a structural property for the cipher under the known key - a property which an ideal cipher (a permutation drawn at random) would not have. Only for completeness, a more relaxed version - called chosen-key distinguisher - can be considered, where the adversary is assumed to have a full control over the key. This model was introduced in [2], and has been extended to a related-key attack on the full-round AES-256.

Since their introductions, known-key attacks have been a major research topic in the symmetric-key community. Indeed, if known-key distinguishers could be considered less relevant than secret-key ones, they anyway allow to learn something about the security margin of a cipher. For example, if it is not possible to find distinguishers for a block cipher when the key is given, then one cannot find a distinguisher when the key is secret. Secondly and more important, block ciphers and hash functions are very close cryptographic primitives, as the latter can be built from the former and vice versa. For example, the Davies-Meyer construction or the Miyaguchi-Preneel construction can transform a secure block cipher into a secure compression function. In a hash setting, block cipher security models such as the known-key model (or the chosen-key model) make sense since in practice the attacker has full access and control over the internal computations. Moreover, an attack in these models depicts a structural flaw of the cipher, while it should be desired to work with a primitive that doesn't have any flaw, even in the most generous security model for the attacker. A classical example is the devastating effect on the compression function security of weak keys for a block cipher [18], which are usually considered as a minor flaw for a block cipher if the set of these weak-keys is small. Therefore, the security notions to consider for a block cipher will vary depending if this block cipher will be used in a hash function setting or not. Citing Knudsen and Rijmen [13], "*imagine a block cipher*" for which a known-key distinguisher exists, "*but where no efficient attacks are known in the traditional black-box model. Should we recommend the use of such a cipher? We don't think so!*"

The Rijndael block cipher [5] has been designed by Daemen and Rijmen in 1997 and accepted as the AES (Advanced Encryption Standard) since 2000 by NIST. Nowadays, it is probably the most used and studied block cipher. In this

paper, using the same strategy proposed by Gilbert at Asiacrypt 2014 [8], we improve all the known-key distinguishers currently present in the literature, and as major result we present the *first known-key distinguisher for 12-round AES*, in turn applicable to *full AES-192*.

1.1 Known-Key Distinguishers for AES: the State of the Art

In the known-key model, a full access to an instance of the encryption function associated with a *known* random key and its inverse is given. The purpose is to simultaneously control the inputs and the outputs of the primitive, i.e. to achieve input-output correlations that one could not efficiently achieve with inputs and outputs of a perfect random permutation to which would have an oracle access. A formal definition of known-key distinguisher is provided in Sect. 4, where we propose and describe in details a generic scenario for known-key distinguisher. We emphasize that *all the known-key distinguishers currently present in the literature - including the one presented in this paper - implicitly exploit (and can be described in) the scenario proposed in Sect. 4*.

AES served as a benchmark for cryptanalytic techniques since the very introduction of this model by Knudsen and Rijmen [13] with a 7-round result. Subsequently, 8-round results were obtained using truncated differentials [9], which were later on improved in [11]. Currently, this last one - which exploits the rebound technique [14] and the so called “multiple limited-birthday problem” - is the best 8-round known-key distinguisher in literature. Recently, Gilbert [8] found a way to extend an 8-round known-key distinguisher (using a novel representation of AES) into a much more intricate 10-round distinguisher and hence presented for the first time a known-key distinguisher for full AES-128. For the sake of completeness, it should be mentioned that even if the strategy proposed by Gilbert allows to set up efficient known-key distinguishers, its “*impact on the security of [...] AES when used as a known key primitive, e.g. in a hash function construction, is questionable*” (see abstract of [8]).

All these known-key distinguishers currently present in literature are briefly recalled in Sect. 4 using the “subspace trail notation”¹, recently introduced at FSE 2017. In Table 1 we list the known-key distinguishers for AES, including our main results (we refer to Table 2 in Sect. 4 for a complete list of our results).

1.2 Our Main Results

In the conclusion of his paper, Gilbert claims that it seems technically difficult to use a stronger property of the balanced one to extend an 8-round known-key distinguisher to a 10-round one. In particular, he left “*the investigation of improved 10-round known-key distinguishers and associated proofs - or even plausible heuristic arguments if rigorous proofs turn out to be too difficult to obtain - as an open issue.*”

¹ Our choice to use the subspace trail notation is due to the fact that it allows in some cases an easier and more formal description than the original notation.

Table 1. *AES known-key distinguishers.* The computation cost is the sum of the computational cost to generate N -tuples of plaintexts/ciphertexts and of the verification cost. The word “Extended” refers to a distinguisher which exploits the technique introduced by Gilbert [8] (in this case we also highlight which distinguisher is extended), while “MultDT” refers to Multiple Differential Trail. A detailed table with all the distinguishers presented in this paper is given in Sect. 4.

Rounds	Computations	Memory	Property	Reference
7	2^{56}	2^{56}	Zero-Sum	[13]
7	2^{24}	2^{16}	Differential Trail	[15]
7	2^{20}	2^{16}	Multiple Diff. Trail	App. F.1
8	2^{64}	2^{64}	Zero-Sum	[8] - App. D
8	2^{48}	2^{32}	Differential Trail	[9]
8	2^{44}	2^{32}	Multiple Diff. Trail	[11]
8	2^{23}	2^{16}	Extended 7-Rounds MultDT	App. F.2
9	2^{50}	2^{32}	Extended 8-Rounds MultDT	Sect. 6
9	2^{23}	2^{16}	Extended 7-Rounds MultDT	App. F.3
10	2^{64}	2^{64}	Extended 8-Rounds Zero-Sum	[8]
10	2^{50}	2^{32}	Extended 8-Rounds MultDT	Sect. 7.1
12	2^{82}	2^{32}	Extended 8-Rounds MultDT	Sect. 8

In this paper, we have accepted the challenge of Gilbert, and using a strategy similar to the one proposed by Gilbert in [8], we show how to construct more efficient 8-, 9- and 10-round distinguishers. To achieve this result, we exploit the known-key distinguishers based on truncated differential trails. In particular, we use as starting point the 8-round known-key distinguisher presented in [11], and we extend it at the end or/and at the beginning using the same strategy proposed by Gilbert. This allows to set up a 9-round known-key distinguisher (see Sect. 6) and a 10-round known-key distinguisher for AES (see Sect. 7.1) with time complexity approximately of 2^{50} . This also provides a counter-example to the claim “*the transposition of our technique to the 8-round distinguisher of [9] does not allow to derive a valid 10-round distinguisher*” made in [8]. Moreover, starting from the 7-round known-key distinguisher presented in [15] - improved in App. F.1 using the “multiple limited-birthday problem” proposed in [11] - and using exactly the same technique presented for the previous cited distinguishers, we are able to set up 8- and 9-round known-key distinguisher for AES (see App. F.2 and F.3), both with complexity approximately of 2^{23} .

As a major result, in Sect. 8 we show that it is possible to extend our 10-round distinguisher up to 12 rounds, which results to be the *first known-key distinguisher for full AES-192*. This also provides a counter-example of the claim made in [8] about the (im)possibility to use Gilbert’s technique to extend a 8-round distinguisher more than 2 rounds: “*The reader might wonder whether the technique we used to derive a known-key distinguisher for the 10-round AES from a known-key distinguisher for the 8-round AES does not allow to extend this 8-*

round known distinguisher by an arbitrary number of rounds. It is easy however to see that the argument showing that 10-round relation R is efficiently checkable does not transpose for showing that the relations over $r > 10$ rounds one could derive from the 8-round relation by expressing that the r -round inputs and outputs are related by $r - 8 > 2$ outer rounds to intermediate blocks that satisfy the 8-round relation are efficiently checkable.” Moreover, in App. H we also show that the same strategy can be (theoretically) used to extend the Gilbert’s 10-round distinguisher up to 12 rounds.

Finally, we discuss why our results no longer exclude known-key distinguishers up to 14 rounds, but at the same time why this seems currently not feasible. Using our results presented in the paper as starting point, we show that one of the main problem (but not the only one) about the possibility to extend a known-key distinguisher exploiting the technique proposed by Gilbert is related to the existence of key-recovery attack on AES with more than a single extension at the end and a computational complexity lower than 2^{128} computations². We refer to Sect. 9 for a complete discussion.

2 Preliminary - Description of AES

The Advanced Encryption Standard [5] is a *Substitution-Permutation network* that supports key size of 128, 192 and 256 bits. The 128-bit plaintext initializes the internal state as a 4×4 matrix of bytes as values in the finite fields \mathbb{F}_{256} , defined using the irreducible polynomial $x^8 + x^4 + x^3 + x + 1$. Depending on the version of AES, N_r round are applied to the state: $N_r = 10$ for AES-128, $N_r = 12$ for AES-192 and $N_r = 14$ for AES-256. An AES round applies four operations to the state matrix:

- *SubBytes* (S-Box) - applying the same 8-bit to 8-bit invertible S-Box 16 times in parallel on each byte of the state (it provides non-linearity in the cipher);
- *ShiftRows* (SR) - cyclic shift of each row to the left;
- *MixColumns* (MC) - multiplication of each column by a constant 4×4 invertible matrix M_{MC} (MC and SR provide diffusion in the cipher³);
- *AddRoundKey* (ARK) - XORing the state with a 128-bit subkey.

One round of AES can be described as $R(x) = K \oplus MC \circ SR \circ \text{S-Box}(x)$. In the first round an additional AddRoundKey operation (using a whitening key) is applied, and in the last round the MixColumns operation is omitted.

Finally, as we don’t use the details of the AES key schedule in this paper, we refer to [5] for a complete description.

The Notation Used in the Paper. Let x denote a plaintext, a ciphertext, an intermediate state or a key. Then $x_{i,j}$ with $i, j \in \{0, \dots, 3\}$ denotes the byte

² Note that given an attack on r rounds with a complexity lower than 2^{128} , one can attack $r + 1$ rounds of AES-256 by guessing the entire first/last secret subkey.

³ SR makes sure column values are spread, MC makes sure each column is mixed.

in the row i and in the column j . We denote by k^r the key of the r -th round, where k^0 is the secret key. If only the key of the final round is used, then we denote it by k to simplify the notation. Finally, we denote by R one round of AES, while we denote r rounds of AES by R^r . We sometimes use the notation R_K instead of R to highlight the round key K . As last thing, in the paper we often use the term “partial collision” (or “*collision*”) when two texts belong to the same coset of a given subspace X .

3 Subspace trails

Invariant subspace cryptanalysis can be a powerful cryptanalytic tool, and subspace trails [10] - introduced at FSE 2017 - are a recent generalization of it.

Let F denote a round function in a iterative block cipher and let $V \oplus a$ denote a coset of a vector space V . Then if $F(V \oplus a) = V \oplus a$ we say that $V \oplus a$ is an *invariant coset* of the subspace V for the function F . This concept can be generalized to *trails of subspaces*.

Definition 1. *Let $(V_1, V_2, \dots, V_{r+1})$ denote a set of $r+1$ subspaces with $\dim(V_i) \leq \dim(V_{i+1})$. If for each $i = 1, \dots, r$ and for each $a_i \in V_i^\perp$, there exist (unique) $a_{i+1} \in V_{i+1}^\perp$ such that $F(V_i \oplus a_i) \subseteq V_{i+1} \oplus a_{i+1}$, then $(V_1, V_2, \dots, V_{r+1})$ is subspace trail of length r for the function F . If all the previous relations hold with equality, the trail is called a constant-dimensional subspace trail.*

This means that if F^t denotes the application of t rounds with fixed keys, then $F^t(V_1 \oplus a_1) = V_{t+1} \oplus a_{t+1}$. We refer to [10] for more details about the concept of subspace trails. Our treatment here is however meant to be self-contained.

3.1 Subspace trails of AES

In this section, we recall the subspace trails of AES presented in [10]. For the following, we only work with vectors and vector spaces over $\mathbb{F}_{2^8}^{4 \times 4}$, and we denote by $\{e_{0,0}, \dots, e_{3,3}\}$ the unit vectors of $\mathbb{F}_{2^8}^{4 \times 4}$ (e.g. $e_{i,j}$ has a single 1 in row i and column j). We also recall that given a subspace X , the cosets $X \oplus a$ and $X \oplus b$ (where $a \neq b$) are *equivalent* (that is $X \oplus a \sim X \oplus b$) if and only if $a \oplus b \in X$.

Definition 2. *The column spaces \mathcal{C}_i are defined as $\mathcal{C}_i = \langle e_{0,i}, e_{1,i}, e_{2,i}, e_{3,i} \rangle$.*

For instance, \mathcal{C}_0 corresponds to the symbolic matrix

$$\mathcal{C}_0 = \left\{ \begin{bmatrix} x_1 & 0 & 0 & 0 \\ x_2 & 0 & 0 & 0 \\ x_3 & 0 & 0 & 0 \\ x_4 & 0 & 0 & 0 \end{bmatrix} \mid \forall x_1, x_2, x_3, x_4 \in \mathbb{F}_{2^8} \right\} \equiv \begin{bmatrix} x_1 & 0 & 0 & 0 \\ x_2 & 0 & 0 & 0 \\ x_3 & 0 & 0 & 0 \\ x_4 & 0 & 0 & 0 \end{bmatrix}.$$

Definition 3. *The diagonal spaces \mathcal{D}_i and the inverse-diagonal spaces \mathcal{ID}_i are respectively defined as $\mathcal{D}_i = SR^{-1}(\mathcal{C}_i) \equiv \langle e_{0,i}, e_{1,i+1}, e_{2,i+2}, e_{3,i+3} \rangle$ and $\mathcal{ID}_i = SR(\mathcal{C}_i) \equiv \langle e_{0,i}, e_{1,i-1}, e_{2,i-2}, e_{3,i-3} \rangle$, where the indexes are taken modulo 4.*

For instance, \mathcal{D}_0 and \mathcal{ID}_0 correspond to symbolic matrix

$$\mathcal{D}_0 \equiv \begin{bmatrix} x_1 & 0 & 0 & 0 \\ 0 & x_2 & 0 & 0 \\ 0 & 0 & x_3 & 0 \\ 0 & 0 & 0 & x_4 \end{bmatrix}, \quad \mathcal{ID}_0 \equiv \begin{bmatrix} x_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 \\ 0 & 0 & x_3 & 0 \\ 0 & x_4 & 0 & 0 \end{bmatrix}.$$

Definition 4. The i -th mixed spaces \mathcal{M}_i are defined as $\mathcal{M}_i = MC(\mathcal{ID}_i)$.

For instance, \mathcal{M}_0 corresponds to symbolic matrix

$$\mathcal{M}_0 \equiv \begin{bmatrix} 0x02 \cdot x_1 & x_4 & x_3 & 0x03 \cdot x_2 \\ x_1 & x_4 & 0x03 \cdot x_3 & 0x02 \cdot x_2 \\ x_1 & 0x03 \cdot x_4 & 0x02 \cdot x_3 & x_2 \\ 0x03 \cdot x_1 & 0x02 \cdot x_4 & x_3 & x_2 \end{bmatrix}.$$

Definition 5. For $I \subseteq \{0, 1, 2, 3\}$, let \mathcal{C}_I , \mathcal{D}_I , \mathcal{ID}_I and \mathcal{M}_I defined as

$$\mathcal{C}_I = \bigoplus_{i \in I} \mathcal{C}_i, \quad \mathcal{D}_I = \bigoplus_{i \in I} \mathcal{D}_i, \quad \mathcal{ID}_I = \bigoplus_{i \in I} \mathcal{ID}_i, \quad \mathcal{M}_I = \bigoplus_{i \in I} \mathcal{M}_i.$$

As shown in detail in [10]:

- for any coset $\mathcal{D}_I \oplus a$ there exists unique $b \in \mathcal{C}_I^\perp$ such that $R(\mathcal{D}_I \oplus a) = \mathcal{C}_I \oplus b$;
- for any coset $\mathcal{C}_I \oplus a$ there exists unique $b \in \mathcal{M}_I^\perp$ such that $R(\mathcal{C}_I \oplus a) = \mathcal{M}_I \oplus b$.

This simply states that a coset of a sum of diagonal spaces \mathcal{D}_I encrypts to a coset of a corresponding sum of column spaces. Similarly, a coset of a sum of column spaces \mathcal{C}_I encrypts to a coset of the corresponding sum of mixed spaces.

Theorem 1. For each I and for each $a \in \mathcal{D}_I^\perp$, there exists one and only one $b \in \mathcal{M}_I^\perp$ such that

$$R^2(\mathcal{D}_I \oplus a) = \mathcal{M}_I \oplus b. \quad (1)$$

We refer to [10] for a complete proof of this theorem. Observe that b depends on a and on the secret key k , and that this theorem doesn't depend on the particular choice of the S-Box (i.e. it is independent of the details of the S-Box).

Observe that if X is a generic subspace, $X \oplus a$ is a coset of X and x and y are two elements of the (same) coset $X \oplus a$, then $x \oplus y \in X$. It follows that:

Lemma 1. For all x, y and for all $I \subseteq \{0, 1, 2, 3\}$:

$$\text{Prob}(R^2(x) \oplus R^2(y) \in \mathcal{M}_I \mid x \oplus y \in \mathcal{D}_I) = 1. \quad (2)$$

As demonstrated in [10], we finally recall that for each $I, J \subseteq \{0, 1, 2, 3\}$:

$$\mathcal{M}_I \cap \mathcal{D}_J = \{0\} \quad \text{if and only if} \quad |I| + |J| \leq 4, \quad (3)$$

Theorem 2. Let $I, J \subseteq \{0, 1, 2, 3\}$ such that $|I| + |J| \leq 4$. For all $x \neq y$:

$$\text{Prob}(R^4(x) \oplus R^4(y) \in \mathcal{M}_I \mid x \oplus y \in \mathcal{D}_J) = 0. \quad (4)$$

4 Known-Key Distinguishers for AES

Before we present our new known-key distinguishers for AES, we review the most relevant ones to our work. First, we give a formal definition of the known-key distinguisher scenario, using as starting point the one given in [8] by Gilbert.

4.1 Definition of Known-Key Distinguisher

Informally, a known-key distinguisher exploits the fact that it is in general harder for an adversary who doesn't know the key to derive an N -tuple of input blocks of the considered block cipher E that is "abnormally correlated" with the corresponding N -tuple of output blocks than for one who knows the secret key. This difficulty is well expressed by the T -intractable definition, first proposed in [3] and [1], and then re-expressed by Gilbert as follows:

Definition 6. *Let $E : (K, X) \in \{0, 1\}^k \times \{0, 1\}^n \rightarrow E_K(X) \in \{0, 1\}^n$ denote a block cipher of block size n bits. Let $N \geq 1$ and \mathcal{R} denote an integer and any relation over the set S of N -tuples of n -bit blocks. \mathcal{R} is said to be T -intractable relatively to E if, given any algorithm \mathcal{A}' that is given an oracle access to a perfect random permutation Π of $\{0, 1\}^n$ and its inverse, it is impossible for \mathcal{A}' to construct in time $T' \leq T$ two N -tuples $\mathcal{X}' = (X'_i)$ and $\mathcal{Y}' = (Y'_i)$ such that $Y'_i = \Pi(X'_i)$, $i = 1, \dots, N$ and $\mathcal{X}' \mathcal{R} \mathcal{Y}'$ with a success probability $p' \geq 1/2$ over Π and the random choices of \mathcal{A}' . The computing time T' of \mathcal{A}' is measured as an equivalent number of computations of E , with the convention that the time needed for one oracle query to Π or Π^{-1} is equal to 1. Thus if q' denotes the number of queries of \mathcal{A}' to Π or Π^{-1} , then $q' \leq T'$.*

Definition 7. *Let $E : (K, X) \in \{0, 1\}^k \times \{0, 1\}^n \rightarrow E_K(X) \in \{0, 1\}^n$ denote a block cipher of block size n bits. A known-key distinguisher $(\mathcal{R}, \mathcal{A})$ of order $N \geq 1$ consists of (1) a relation \mathcal{R} over the N -tuples of n -bit blocks (2) an algorithm \mathcal{A} that on input a k -bit key K produces in time $T_{\mathcal{A}}$, i.e. in time equivalent with $T_{\mathcal{A}}$ computations of E , an N -tuple $\mathcal{X} = (X_i)$ $i = 1, \dots, N$ of plaintext blocks and an N -tuple $\mathcal{Y} = (Y_i)$ $i = 1, \dots, N$ of ciphertext blocks related by $Y_i = E_K(X_i)$ and by $\mathcal{X} \mathcal{R} \mathcal{Y}$. The two following conditions must be met:*

- *The relation \mathcal{R} must be $T_{\mathcal{A}}$ -intractable relatively to E ;*
- *The validity of \mathcal{R} must be efficiently checkable: we formalize this requirement by incorporating the time for checking whether two N -tuples are related by \mathcal{R} in the computing time $T_{\mathcal{A}}$ of algorithm \mathcal{A} .*

We emphasize that while the algorithm \mathcal{A} takes a random key K as input, *the relation \mathcal{R} satisfied by the N -tuples of input and output blocks constructed by \mathcal{A} or \mathcal{A}' is the same for all values of K (in other words, it is independent of K) and must be efficiently checkable without knowing K .*

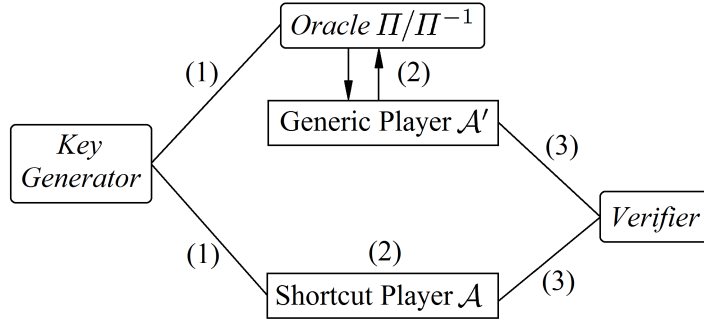


Fig. 1. A *Known-Key Distinguisher Scenario*. Step (0): a relationship \mathcal{R} is chosen. Step (1): the secret key is given to the Oracle Π/Π^{-1} and to the Shortcut Player \mathcal{A} . Step (2): the Shortcut Player \mathcal{A} and the Generic Player \mathcal{A}' generate the N -tuples that satisfy the required relationship \mathcal{R} . Step (3): the Verifier receives the N -tuple and checks if \mathcal{R} is satisfied or not. The faster player to generate the N -tuple wins the “game”.

The Known-Key Distinguisher Scenario. To better understand these definitions, we propose and describe in more details a generic scenario for a known-key distinguisher, which is depicted in Fig. 1. This scenario is composed of five characters, which are a key generator, an oracle, two players and a verifier. First of all - *step (0)*, we assume that a relation \mathcal{R} defined as in Def. 6 is chosen. At *step (1)*, the key generator generates a key, which is given to the oracle and to one of the two player. For the following, we call “*shortcut player*” the player that knows the key and “*generic player*” the player that doesn’t know it. Referring to the previous definitions by Gilbert, the generic player can be identified with the algorithm \mathcal{A}' , while the shortcut player can be identified with the algorithm \mathcal{A} . At *step (2)*, the two players generate the N -tuple of (plaintexts, ciphertexts) which satisfy the required relation \mathcal{R} . Since the generic player doesn’t know the key, he must ask the oracle (identified with Π and/or Π^{-1} in the previous definitions) for the encryption (resp. decryption) of *random* plaintexts (resp. ciphertexts). We stress that this step doesn’t consist only on the generation of (plaintext, ciphertext) pairs, but also includes any computational cost that the player must do in order to find the N -tuple with the required property. When a player finds the N -tuple which satisfies the required relation \mathcal{R} , he sends it to the verifier - *step (3)*. The verifier receives it and checks if the N -tuple satisfied this relation (remember that the verifier doesn’t know the key). The first/fastest player who sends the N -tuple with the required property wins the “game”.

Before going on, we emphasize that *the role of the verifier is only to prevent one or both of the two players from cheating*. In other words, in the case of honest players, the verifier can be skipped, and the winner of the game is simply the first/fastest player that claims to have found the N -tuple of (plaintexts, ciphertexts) which satisfy the required relation \mathcal{R} . We highlight that such a *verifier is implicitly present in all the distinguishers currently present in literature*.

Table 2. *Details of AES known-key distinguishers presented in this paper, obtained by extending distinguishers based on Multiple Differential Trails. “Rounds” denotes the number of rounds of the basic distinguisher + the number of rounds of the extensions (if even, the number of extension rounds is equal at the end and at the beginning). “Cost Case 1” denotes the cost of the shortcut player when the total cost of the generic player is approximated by the number of oracle-queries, while “Cost Case 2” refers to the case in which the total cost of the generic player is the sum of number of queries and of its computational cost. “Cost Verifier” denotes the cost of the verifier. A check-mark ✓ in the “KS” column denotes the case in which the key schedule holds, ✗ denotes the case in which the sub-keys are independent, while white-space/no-mark denotes the case in which the two previous cases are equivalent (for the distinguisher purposes).*

Rounds	KS	Cost Case 1	Cost Case 2	Cost Verifier	Memory	Reference
7 + 1		2^{23}	2^{21}	$2^{11.8}$	2^{16}	App. F.2
8 + 1		2^{50}	$2^{45.6}$	$2^{11.6}$	2^{32}	Sect. 6 - App. E.1
7 + 2	✗	2^{23}	2^{21}	$2^{12.6}$	2^{16}	App. F.3
7 + 2	✓	2^{21}	2^{21}	$2^{12.6}$	2^{16}	App. F.3
8 + 2	✗	2^{50}	$2^{45.6}$	$2^{12.5}$	2^{32}	Sect. 7.1
8 + 2	✓	2^{46}	2^{45}	$2^{12.5}$	2^{32}	Sect. 7.2 - App. E.2
8 + 4		2^{82}	2^{82}	$2^{71.1}$	2^{32}	Sect. 8

A distinguisher is meaningful if the cost of the generic player - we assume that the cost of one oracle-query is equal to the cost of one encryption - to generate the N -tuple is higher than the cost of the shortcut player, when the probability of success is equal for the two players. Note that in this scenario we are considering the computational costs of the two players to generate the N -tuples with a fixed probability of success (equal for both the players). We highlight that this scenario is completely equivalent to the one in which the computational cost to generate the N -tuple is fixed and equal for the two players, and one considers the probabilities of success of the two players (where it is requested that the probability of success is higher for the shortcut player than for the generic one).

Both for the distinguisher that we are going to present and for the Gilbert’s one, the computational cost of the verification step is not negligible. Thus, in order to compare our distinguishers to the others present in literature, we define *the cost of the distinguisher as the sum of the cost of the verification step (that is, the cost of the verifier) and of the cost to construct the set of plaintexts/ciphertexts with the required property (that is, the cost of the shortcut player - the cost of the other player is higher)*. For this reason, we assume for the following that *a relationship \mathcal{R} is efficiently checkable if and only if the computational cost of the verifier is negligible with respect to the player ones*. This implies that the cost of the distinguisher can be approximated with the computational cost of the shortcut player (the cost of the other player is always higher). Moreover, this assumption prevents the construction of meaningless known-key distinguishers, as discussed in Sect. 9.




Fig. 2. 7- and 8-round differential paths for AES-128.

Table 2 summarizes the main details of all the known-key distinguishers presented in this paper with respect to the above scenario. To better understand this table, some considerations must be done. Since the generic player depends by the oracle to generate the N -tuple (i.e. he cannot work alone to generate it), two possible settings can be analyzed. In the first one, only the number of oracle queries is considered to determine the computational cost of this player, that is the number of encryptions/decryptions required by the generic player to the oracle - this case is denoted by “Case 1” in Table 2. In the second one, both the number of oracle queries and any other computational cost of the generic player (which is in general not negligible) are considered - this case is denoted by “Case 2” in Table 2. Intuitively this second setting is weaker than the first one, in the sense that a known-key distinguisher in the first setting works also in the second one but not viceversa. In other words, one can expect that the required number N of tuples is higher in the first setting than in the second one (or equal in the best case). If the total cost of the generic player is well approximated by the number of queries, these two settings are completely equivalent.

In the following, we recall the known-key distinguishers present in literature in the above scenario. Before going on, we emphasize that *despite the claim made in [8], Gilbert’s 10 rounds distinguisher doesn’t satisfy the requirement that the verification cost is lower than the cost of the two players*, i.e. that the requirement \mathcal{R} is efficiently checkable (we refer to App. A for more details).

4.2 7- and 8-Round Known-Key Distinguisher

In the 7- and 8-round known-key distinguishers proposed in [15] and [9], the goal of the two players is to find two pairs of (plaintexts, ciphertexts) - i.e. (p^1, c^1) and (p^2, c^2) - with the following properties: the two plaintexts belong to the same coset of \mathcal{D}_i - i.e. $p^1 \oplus p^2 \in \mathcal{D}_i$ - and the two ciphertexts belong to the same coset of \mathcal{M}_i - i.e. $c^1 \oplus c^2 \in \mathcal{M}_i$ - for a fixed $i \in \{0, 1, 2, 3\}$.

In the above known-key distinguisher setting, the best technique that the shortcut player (i.e. the player who knows the key) can use to win the game is the *Rebound Attack*. The rebound attack is a differential attack and it was proposed in [16] for the cryptanalysis of AES-based hash functions. Since it is a differential attack, one needs a “good” (truncated) differential trail in order to exploit it. Examples of truncated differential trails used for 7- and 8-round AES are depicted in Fig. 2. The rebound attack consists of two phases, called inbound and outbound phase. In the first one, the attacker uses the knowledge




Fig. 3. 8-round differential characteristic for known-key distinguisher of AES-128.

of the key to find pairs of texts that satisfy the middle rounds of the truncated differential trail. In the second one, he propagates the solutions found in the first phase in the forward and in the backward directions, and checks if at least one of them satisfies the entire differential trail. A complete description of the rebound attack is given in App. C, with particular attention to the AES case.

As proved in [9], in the case of a perfect random permutation 2^{64} operations are required to find (plaintexts, ciphertexts) pairs (p_1, c_1) and (p_2, c_2) that have the required properties with good probability. Instead, for the AES case and using the rebound attack, 2^{48} computations are sufficient to find them with the same probability (besides a memory cost of $16 \times 2^{32} = 2^{36}$ bytes).

4.3 Multiple Limited-Birthday 8-Round Known-Key Distinguisher

An improvement of the previous known-key distinguisher on 8-round of AES was proposed in [11]. Using the subspace trail notation, in this modified version of the 8-round known-key distinguisher, the goal of the two players is to find two pairs of (plaintexts, ciphertexts) such that the two plaintexts belong to the same coset of \mathcal{D}_i for an arbitrary i and the two ciphertexts belong to the same coset of \mathcal{M}_j for an arbitrary j , where i and j are not fixed in advance and it is not required that they are equal (i.e. no condition is imposed on i and j). For arbitrary initial and final subspaces, the computational cost is reduced from 2^{48} to 2^{44} (note that there are 4 initial and final different subspaces \mathcal{D}_i and \mathcal{M}_j , for a total of $4^2 = 2^4$ possibilities) while the required memory is still 2^{32} , as shown in detail in [11]. In App. F.1 we show that the same technique can be used to improve the 7-round known-key distinguisher of AES presented in [15].

4.4 10-Round Gilbert’s Known-Key Distinguisher

Integral 8-Round Known-Key Distinguisher. Another 8-round known-key distinguisher for AES is based on the balance property and it was proposed by Gilbert in [8]. In this case, the goal of the two players is to find a set of 2^{64} (plaintext, ciphertext) pairs such that the sums over the plaintexts and over the ciphertexts are equal to zero (i.e. the plaintexts and the ciphertexts are

uniformly distributed). We review this distinguisher in details in App. D using the subspace trails notation instead of the *Super-SB* notation introduced by Gilbert ($Super-SB(\cdot) \equiv S\text{-Box} \circ ARK \circ MC \circ S\text{-Box}(\cdot)$).

We limit here to recall briefly the best strategy that the shortcut player can use to win the game. The idea is to start in the middle with a set S of texts defined as $S := \mathcal{D}_i \oplus \mathcal{M}_j \oplus c$ for a constant c , where $|S| = 2^{64}$. After encrypting S for 4 rounds, the texts are uniform distributed in each coset of \mathcal{M}_I of dimension 12 (i.e. $|I| = 3$ fixed). That is, after 4 rounds, each coset of \mathcal{M}_I for $|I| = 3$ contains exactly 2^{32} elements. The same happens if one decrypts S for 4 rounds. In this case, after decrypting S for 4 rounds, the texts are uniform distributed in each coset of \mathcal{D}_I of dimension 12 (i.e. $|I| = 3$ fixed), that is each coset of \mathcal{D}_I for $|I| = 3$ contains exactly 2^{32} elements. This implies the balance property both on the plaintexts and on the ciphertexts. We refer to App. D for details.

Extension to 10 Rounds. This distinguisher is the starting point used by Gilbert in order to set up the first 10-round known-key distinguisher for AES. The basic idea is to extend this 8-round distinguisher based on the balance property adding one round at the end and one at the beginning. In the known-key distinguisher scenario presented above, the players have to send to the verifier 2^{64} (plaintext, ciphertext) pairs, that is (p^i, c^i) for $i = 0, \dots, 2^{64} - 1$, with the following properties⁴:

1. there exists a key k^0 s.t. the texts $\{R_{k^0}(p^i)\}_i$ are uniform distributed among the cosets of \mathcal{D}_I with $|I| = 3$ fixed, i.e. 2^{32} texts for each coset of \mathcal{D}_I ;
2. there exists a key k^{10} s.t. the texts $\{R_{k^{10}}^{-1}(c^i)\}_i$ are uniform distributed among the cosets of \mathcal{M}_J with $|J| = 3$ fixed, i.e. 2^{32} texts for each coset of \mathcal{M}_J .

Equivalently, a key k^0 must exist such that the sum of the plaintexts after one round is equal to zero, i.e. $\bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p^i) = 0$, and a key k^{10} must exist such that the sum of the ciphertexts one round before is equal to zero, i.e. $\bigoplus_{i=0}^{2^{64}-1} R_{k^{10}}^{-1}(c^i) = 0$. We emphasize that even if this is a known-key distinguisher, the verifier must be able to check the previous properties *without* the knowledge of the key or the subkeys. Since the verifier has no information of the key, one must show that the above conditions are efficiently checkable.

The only way to verify these requirements is to find these two subkeys in an efficient way, which is not possible using a brute force attack (k^0 and k^{10} have 128 bits). Instead to check all the $2 \cdot 2^{128} = 2^{129}$ possible values of k^0 and k^{10} , the idea proposed by Gilbert is to use the integral attack [4]-[12] working on single columns of c^i and of $SR^{-1}(p^i)$. In this way, the verifier must guess only 32 bits instead of 128, and she has to repeat this operation 4 times (one for each column/diagonal) for each key. Thus, working independently on each column of the keys and of the texts, the verifier can check the zero-sum property. In App.

⁴ For this and the following distinguishers, we abuse the notation k^r to denote a key of a certain round r . We emphasize that k^r is not necessarily equal to the secret key, that is k^r can be different from the r -th subkey. Remember that it is only required that such a key exists, and not that it is equal to the real secret key.

H, we show that this procedure can be even improved working independently on each byte of k^0 and k^{10} instead of entire column/diagonal. This (theoretically) allows to extend the 10-round distinguisher of Gilbert up to 12 rounds using the technique presented in this paper.

In conclusion, the shortcut player (i.e. the one who knows the key) can construct these 2^{64} (plaintext, ciphertext) pairs using the same strategy proposed for the 8 rounds distinguisher (note that in this case the keys k^0 and k^{10} corresponds to the secret sub-keys). Instead, Gilbert showed that the probability that the generic player (i.e. the one who doesn't know the secret key) successfully outputs (input, output) pairs that satisfy the previous properties (both in the input and in the output) is upper bounded by $2^{-16.5}$ (see [8] - Prop. 6, for more details). Equivalently, the computational cost required by the generic player to construct such set with the same probability of the shortcut player is of $2^{16.5} \times 2^{64} = 2^{80.5}$, which is higher than the cost of the other player.

Finally, a consideration about the cost of the verifier (i.e. of the verification step) has to be done. In [8], it is claimed that⁵ “*the overall complexity of checking \mathcal{R} is strictly smaller than $N = 2^{64}$ AES₁₀^{*} operations*” (see Sect. 4.2 - page 218), that is the computational cost of the verifier is negligible with respect to the cost of the two players. In spite of this, even using a better strategy than the one proposed by Gilbert, we show in App. A that this claim is not true, that is the cost of the verifier (approximately $2^{69.4}$ ten-round encryptions) is higher than the cost of the shortcut player (approximately 2^{64} ten-round encryptions) using an integral attack - which is the strategy proposed in [8]. On the other hand, we show how to modify the distinguisher proposed by Gilbert in order to fix the problem and to fulfill this requirement (see App. H for all the details).

Generic Considerations. The previous 10-round distinguisher proposed by Gilbert is different from all the previous distinguishers up to 8 rounds present in literature. For all the distinguishers up to 8-round, the relation \mathcal{R} that the N -tuple of (plaintexts, ciphertexts) must satisfy doesn't involve any operation of the block cipher E . As a consequence, it allows the verifier to check whether the N -tuple of (plaintexts, ciphertexts) satisfy the required relation \mathcal{R} without knowing anything of the key. When \mathcal{R} doesn't re-use operations of E , this provides some heuristic evidence that it can be considered *meaningful*.

On the other hand, the previous 10-round distinguisher and the ones that we are going to propose don't satisfy this requirement, i.e. in these cases the relation \mathcal{R} involves and re-uses some operations of E . We refer to Sect. 3 of [8] for a detailed discussion on the reasons why such known-key distinguishers should not be systematically ruled out as if they were *artificial*.

A Variant of Gilbert's Distinguisher. Before we go on, we highlight a variant of the Gilbert's distinguisher - that also applies to all our proposed distinguishers present in the paper - which allows to better understand it. Consider the case in which the two players have to send to the verifier the N -tuple that

⁵ We emphasize that no proof or strong argumentation of this fact is given in [8].

verify the required relation \mathcal{R} together with the subkeys for which such relation is satisfied. As an example, in the 10-round integral distinguisher just presented, the players have to send 2^{64} (plaintexts, ciphertexts) pairs (p^i, c^i) and the two subkeys k^0 and k^{10} such that $\bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p^i) = \bigoplus_{i=0}^{2^{64}-1} R_{k^{10}}^{-1}(c^i) = 0$. Thus, since the task of the verifier is to check that the relation \mathcal{R} is satisfied only for the keys she received, it follows that her computational cost is negligible. On the other hand, we show in details in App. B that such variant of the distinguisher is meaningless, since it can be set up for any number of rounds of AES.

5 Key-Recovery Extensions using Truncated Differentials

Our known-key distinguishers exploit the same idea proposed for the first time by Gilbert. In particular, our idea is to extend the 8-round distinguishers recalled in Sect. 4.3 at the end or/and at the beginning, in the same way used by Gilbert to extend the 8-round distinguisher based on the balance property.

Since we are going to extend known-key distinguishers based on truncated differential properties, we need an efficient key-recovery attack that allows the verifier to check the required property on the N -tuple of (plaintexts, ciphertexts) that she receives by the players. For this reason, we re-propose the low-data complexity truncated differential attacks⁶ on 3- and 4-round AES-128 presented in [10]. The attacks that we present here are a little modified with respect to those presented in [10] due to different scope of this work. In particular, the attack on 3 rounds of [10] is described here as an attack on a single round, while the attack on 4 rounds is described here as an attack on 2 rounds (besides other changes for this second case, which are described in the following).

5.1 Attack for the Case of 1-Round Extension

Consider three plaintexts in the same coset of \mathcal{M}_i for $|i| = 1$ and the corresponding ciphertexts after one round⁷, that is (p^j, c^j) for $j = 1, 2, 3$. The goal of the attack is to find the key k such that the ciphertexts belong to the same coset of \mathcal{M}_i one round before, that is k has to satisfy the following condition⁸:

$$R_k^{-1}(c^1) \oplus R_k^{-1}(c^2) \in \mathcal{M}_i \quad \text{and} \quad R_k^{-1}(c^1) \oplus R_k^{-1}(c^3) \in \mathcal{M}_i.$$

For simplicity, we assume that the final MixColumns operation is omitted (otherwise one simply switches the final MixColumns and the final AddRoundKey operation, as usual in literature). Since each column of \mathcal{M}_i depends on different and independent variables, the idea of the attack is to work independently on each column of \mathcal{M}_i or equivalently of $SR^{-1}(k)$, and to exploit the relationships that hold among the bytes that lie in the same column of \mathcal{M}_i .

⁶ We emphasize that *both these attacks have been practical verified* (see [10] for details).

⁷ More generally, consider two couples of (plaintexts, ciphertexts) pairs, that is $\{(p_0^j, c_0^j), (p_1^j, c_1^j)\}$ for $j = 1, 2$ such that $p_0^j \oplus p_1^j \in \mathcal{M}_i$.

⁸ Note that if $R_k^{-1}(c^1) \oplus R_k^{-1}(c^2) \in \mathcal{M}_i$ and $R_k^{-1}(c^1) \oplus R_k^{-1}(c^3) \in \mathcal{M}_i$, it follows that also $R_k^{-1}(c^2) \oplus R_k^{-1}(c^3) \in \mathcal{M}_i$ since \mathcal{M}_i is a subspace.

Without loss of generality, we assume $I = \{0\}$ and we present the attack only for the first column of $SR^{-1}(k)$ (analogous for the others). The conditions that the bytes of the first column of $SR^{-1}(k)$ must satisfy are:

$$s_{0,0}^h = 0x02 \cdot s_{1,3}^h, \quad s_{2,2}^h = s_{1,3}^h, \quad s_{3,1}^h = 0x03 \cdot s_{1,3}^h, \quad (5)$$

where $s_{i,j}^h = \text{S-Box}^{-1}(c_{i,j}^1 \oplus k_{i,j}) \oplus \text{S-Box}^{-1}(c_{i,j}^h \oplus k_{i,j})$ for $h = 2, 3$. For each value of $k_{1,3}$ (2^8 possible values in total), the idea is to find the values of $k_{0,0}$, $k_{2,2}$ and $k_{3,1}$ that satisfy the previous relationships. On average, using a single pair of ciphertexts and working in this way, it is possible to find 2^8 combinations of these four bytes (i.e. one for each possible value of $k_{1,3}$). The idea is to test them using the second pair of ciphertexts: on average, only the right combination passes the test. The same procedure is used for the others columns.

The total computational cost of the attack is well approximated by the cost of the first phase, that is by the cost to find (on average) the 2^8 combinations of $k_{0,0}, \dots, k_{3,1}$ that satisfy (5) for the first column and similar for the others (the cost to check them with the second pair of texts is negligible). In particular, the computational cost of this attack using 3 chosen plaintexts can be approximated by $2^{17.1}$ S-Box look-ups (and negligible memory cost), or approximately $2^{11.6}$ table look-ups and a memory cost of $16 \times 2^{12} = 2^{16}$ using a precomputation phase. We refer to [10] for all the details.

For the following, we emphasize that the same attack works exactly in the same way also in the decryption direction (chosen ciphertexts attack) with the same complexity. In this case the idea is to consider three ciphertexts in the same coset of \mathcal{D}_i , and to look for a key such that the corresponding plaintexts belong to the same coset of \mathcal{D}_i after one round (see [10] for details).

5.2 Attack for the Case of 2-Round Extension

To set up the first 12-round known-key distinguisher of AES-128, we also need to recall (a modified version of) the low-data complexity truncated differential attack on 4-round of AES-128, which is obtained by extending the previous attack on 3 rounds at the end. We refer to [10] for a complete description of the attack, and for simplicity we assume that the final MixColumns is omitted.

Consider plaintexts in the same coset of \mathcal{M}_i for $|i| = 1$ and the corresponding ciphertexts after two rounds. The goal of the attack is to find the key such that the ciphertexts belong to the same coset of \mathcal{M}_i two rounds before. The idea of the attack is to guess two columns of $SR^{-1}(k^2)$, where k^2 is the final key. Given 5 plaintexts and the corresponding ciphertexts (p^j, c^j) for $j = 1, \dots, 5$, for each one of the 2^{64} possible values of these two columns of $SR^{-1}(k^2)$, the idea is to partially decrypt these 5 ciphertexts one round, that is to compute the eight bytes $s^j := R_{k^2}^{-1}(c^j)$ for each $i = 1, \dots, 5$. Due to the ShiftRows operation, these 8 bytes are distributed in two columns. Thus, the idea is to simply to repeat the previous attack on 3 rounds. However, due to the ShiftRows operation, the eight bytes of s^i are uniform distributed in the four columns, i.e. two byte for each column, that is for each column one can only exploit the relationship that holds among these two bytes (see [10] for details).

Using two pairs of ciphertexts (e.g. (c^1, c^2) and (c^1, c^3)), it is possible to find (on average) at most one combination of eight bytes of k^3 for each possible guess of the eight bytes of k^2 , for a total of 2^{64} possibilities. The idea is to test these found values against other pairs of ciphertexts, that is to check if the relationships among the bytes of the keys hold also for these other pairs of ciphertexts⁹. Since each relationship is satisfied with probability 2^{-32} (there are four relationships, each one satisfied with probability 2^{-8}), it is sufficient to test the found values of k^1 and k^2 against only other two pairs of ciphertexts, in order to eliminate all the wrong candidates with high probability. Thus, using 5 chosen plaintexts (i.e. 4 pairs with a common plaintext¹⁰), it is possible to recover 8 bytes of k^1 and of k^2 . To discover the complete key, the idea is essentially to repeat the same procedure on the last two columns of k^2 (we refer to [10] for details).

As shown in [10], the computational cost of this attack is well approximated by 2^{81} S-Box look-ups (with negligible cost of memory) or 2^{76} table look-ups and a memory cost of $16 \cdot 2^{12} = 2^{16}$ bytes. Moreover, the same attack works also in the decryption direction, with the same complexity. In particular, given ciphertexts in the same coset of \mathcal{D}_i for $|i| = 1$ and the corresponding plaintexts two rounds before, the idea is to look for the keys such that the plaintexts belong to the same coset of \mathcal{D}_i after two rounds.

6 9-Round Known-Key Distinguisher for AES

Exploiting the same idea proposed by Gilbert, we set up our known-key distinguisher for 9 rounds of AES by extending the 8-round distinguisher presented in [11] (and recalled in Sect. 4.3) at the end (or at the beginning).

In the above defined known-key scenario, the players have to send to the verifier n different tuples of (plaintext, ciphertext) pairs, that is $\{(p_i^1, c_i^1), (p_i^2, c_i^2)\}$ for $i = 0, \dots, n - 1$, with the following properties¹¹:

1. for each tuple, there exists j s.t. the two plaintexts belong to the same coset of \mathcal{D}_j , that is

$$\forall i = 0, \dots, n - 1, \quad \exists j \in \{0, \dots, 3\} \quad \text{s.t.} \quad p_i^1 \oplus p_i^2 \in \mathcal{D}_j;$$

2. there exists a key k s.t. for each tuple there exists l for which the two ciphertexts belong to the same coset of \mathcal{M}_l one round before, that is

$$\exists k \quad \text{s.t.} \quad \forall i = 0, \dots, n - 1, \quad \exists l \in \{0, \dots, 3\} \quad \text{s.t.} \quad R_k^{-1}(c_i^1) \oplus R_k^{-1}(c_i^2) \in \mathcal{M}_l.$$

⁹ This step is different from the one proposed in [10]. In that case, the idea is to find the right key by a brute force attack in order to keep the data complexity as low as possible. For our distinguisher, we propose to test the found key against other pairs of plaintexts and ciphertexts, since it is not possible to use a brute force attack.

¹⁰ Note that 4 different pairs can be obtained by 3 chosen plaintexts. However, such pairs are not useful for the attack, essentially for the same reason given in footnote 8. We refer to [10] for a complete and detailed explanation.

¹¹ We say that the tuples are different if $p_i^1 \neq p_i^2$ for each i and if $(p_i^1, p_i^2) \neq (p_j^1, p_j^2)$ and $(p_i^1, p_i^2) \neq (p_j^2, p_j^1)$ for each $i \neq j$.

The fastest player to construct these n -tuples wins the game.

We stress that the *key k must be equal for all the tuples*. In other words, if there exist two different tuples (c_0, c_1) and (c_2, c_3) such that $R_k^{-1}(c_0) \oplus R_k^{-1}(c_1) \in \mathcal{M}_l$ and $R_{\tilde{k}}^{-1}(c_2) \oplus R_{\tilde{k}}^{-1}(c_3) \in \mathcal{M}_l$ for two different keys $k \neq \tilde{k}$, then the above defined relationship \mathcal{R} is not satisfied. Note that without this request on the secret key k , it is extremely easy to construct tuples such that the two ciphertexts belong to the same coset of \mathcal{M}_l one round before. Indeed, as we have seen for the attack of Sect. 5, given two ciphertexts c^1 and c^2 , on average there exist $4 \cdot (2^8)^4 = 2^{34}$ different keys such that $R^{-1}(c^1) \oplus R^{-1}(c^2) \in \mathcal{M}_l$ for a certain l . Thus, it is straightforward to construct n different tuples with the above defined relationship \mathcal{R} but without any condition on the key¹² k .

Before we go on, it is also important to emphasize that no condition on the key k is imposed, except that it exists and it is equal for all the tuples. That is, it is not required that this key is equal to the real secret subkey. The same consideration holds also for the next distinguishers presented in this paper, and for the 10-round distinguisher presented by Gilbert in [8].

In the following, we present the distinguisher in details. To obtain a suitable value for n , we consider the best strategy that the generic player can adopt to win the game. A value of n is suitable when the computational cost of the generic player using this best strategy is worse than the one of the other player.

As we show in details in the following, to do this one has to consider the numbers of oracle-queries done by the two player and any further cost of the generic player. In particular, if only the number of oracle-queries is taken in account, then n must be equal or greater than 8, which implies that the computational cost for the shortcut player is of 2^{47} and for the generic player is of $2^{48.9}$. In order to make the advantage of the shortcut player more significant, we have chosen an (arbitrary) value of $n = 64$, which implies a cost for the shortcut player of 2^{50} computations and of $2^{65.6}$ computations for the generic player. Instead, if all the costs are considered (number of oracle-queries + cost of the generic player), then a suitable value of n is 3, the computational cost for the shortcut player is $2^{45.6}$ and for the generic player is approximately $2^{109.5}$. In both cases, the computational cost of the verifier is well approximated by $2^{11.6}$.

The Verifier. Given n tuples, for each one of them the verifier can easily check if the two plaintexts belong (or not) to the same cosets of \mathcal{D}_j for a certain j , by computing their XOR sum and checking that three diagonals are equal to zero.

More complicated is to check if there exists a (unique) key k for which the requirement on the ciphertexts is fulfilled. The idea is to find such key (if exists), using the attack described in Sect. 5.1. First of all, given a single tuple, there exist on average $4 \cdot (2^8)^4 = 2^{34}$ keys of the final round such that the two ciphertexts

¹² We observe that the claim “*the transposition of our technique to the 8-round distinguisher of [9] does not allow to derive a valid 10-round distinguisher*” made in [8] is justified only when no assumption on the key k is done. In other words, the above defined relationship \mathcal{R} together with the requirement of *uniqueness of the key k* allows to extend the 8-round distinguisher of [9] as in [8].

belong to the same coset of \mathcal{M}_l one round before for a certain l . Given two tuples, the probability that such key exists is only $(2^{34})^2 \cdot 2^{-128} = 2^{-60}$, while more generally, given n tuples, the probability that at least one key exists (for which the previous requirements are satisfied) is given by:

$$2^{34n} \cdot 2^{-128(n-1)} = 2^{-94 \cdot n + 128}.$$

This is due to the fact that for each tuple there are one average 2^{34} different keys and that the probability that two keys are equal is 2^{-128} . By this preliminary analysis, it is already possible to deduce that the number of tuples should be at least 2 (i.e. $n \geq 2$). Indeed, for $n = 1$ such a key always exists (which implies that using a random tuple it is possible to win the game), while for $n = 2$ the probability that such key exists for two random tuples is only 2^{-60} .

Thus, assume that the verifier receives $n \geq 2$ tuples. The idea is to use two tuples and the attack described in Sect. 4.3 to recover (if exists) the key that satisfies the required property. If $n > 2$, the verifier simply checks if the relation \mathcal{R} is satisfied by the found key for the other $n - 2$ tuples.

In more details, working independently on each column, the attacker uses the first tuple to find 2^8 combinations for each column of $SR^{-1}(k)$ and checks immediately them with the second tuple. Since she repeats this attack for each possible \mathcal{M}_i (i.e. 4 times), the cost of this step is of $4 \cdot 2^{17.1} = 2^{19.1}$ S-Box look-ups. In this way, the verifier finds on average only one key (if exists). If at least one possible key is found using two tuples, she simply checks if the other $n - 2$ tuples satisfy the relation \mathcal{R} for this found key (more generally, she repeats this step for all the keys found using the first two tuples). The cost of this operation is well approximated by $2 \cdot 16 = 2^5$ S-Box look-ups for each tuple (note that she must decrypt one round two ciphertexts).

In conclusion, given $n \geq 2$ tuples, the cost of the verifier is well approximated by $2^{19.1} + (n - 2) \cdot 2^5$ S-Box look-ups, that is approximately $2^{11.6}$ 9-round encryptions if $n \ll 2^{14}$.

The Shortcut Player. The shortcut player can simply use the same strategy described in [11] and in Sect. 4.3 for the known-key distinguisher on 8 rounds to find the n tuples that satisfy the above defined relation \mathcal{R} . Indeed, it is straightforward to prove that all the properties are satisfied, since for each tuple the two plaintexts belong to the same coset of \mathcal{D}_i (for a certain i) and the two ciphertexts belong to the same coset of \mathcal{M}_j (for a certain j) one round before with respect to the known key - by construction. Since the computational cost to build one tuple is of 2^{44} encryptions, the cost to construct n tuples is well approximated¹³ by $n \cdot 2^{44}$.

The Generic Player. Here we analyze and present the (intuitively) best strategy that the generic player can use to find n tuples with the required properties,

¹³ We don't exclude the possibility of some trade-offs that could allow to reduce the computational cost to construct n tuples, i.e. such that the total computational which increases less than linear. However, for our results, the "roughly" linear approximation is sufficient.

and the corresponding computational cost. Intuitively, the best strategy for this player is to choose tuples such that for each one of them the two plaintexts belong to the same coset of \mathcal{D}_j for a certain j . In this way, the required condition on the plaintexts is (obviously) satisfied. Then, the player asks the oracle for the corresponding ciphertexts. The idea is to check if there exists a key k and n tuples such that the two ciphertexts of each of these n tuples belong to the same coset of \mathcal{M}_l one round before. We remember that it is not necessary that the key for which this condition is satisfied is the real one.

As we have already seen, given a single tuple there exist on average 2^{34} keys such that the two ciphertexts belong to the same coset of \mathcal{M}_j one round before. To set up a meaningful distinguisher, a value of n is suitable if the number of oracle-queries of the generic player is higher than the cost of the shortcut player. By previous observations, given a set of n tuples, the probability that at least one common key exists for which the property on the ciphertexts is satisfied is $2^{-94n+128}$. Thus, the idea is to estimate the number of (plaintext, ciphertext) pairs that this player has to generate in order to win the game (that is, in order to find with high probability n tuples with the required property). If this number is higher than $2^{44} \cdot n$ for a fixed n , then the other player wins the game.

Since each coset of \mathcal{D}_j contains 2^{32} different plaintexts, it is possible to construct approximately 2^{63} different couples $\{(p^1, c^1), (p^2, c^2)\}$. Given t different cosets of \mathcal{D}_j , it is possible to construct $s = 2^{63} \cdot t$ different couples. It follows that one can construct approximately

$$\binom{s}{n} \approx \frac{s^n}{n!}$$

different sets of n different tuples (i.e. n different couples $\{(p^1, c^1), (p^2, c^2)\}$), where the approximation holds for $n \ll s$. Since the probability that a set of n tuples satisfy the above defined relation \mathcal{R} is $2^{-94n+128}$, the generic player must consider at least s different couples such that $s^n/n! \simeq 2^{94n-128}$ or equivalently

$$s \simeq 2^{94 - \frac{128}{n}} \cdot (n!)^{\frac{1}{n}}. \quad (6)$$

By this formula, for $n = 8$ this player has to consider approximately $2^{79.9}$ different tuples, or equivalently $2^{48.9}$ (plaintext, ciphertext) pairs (that is, $2^{16.9}$ initial different cosets of \mathcal{D}_j). Indeed, given $2^{16.9}$ initial different cosets of \mathcal{D}_j , it is possible to construct approximately $2^{16.9} \cdot 2^{63} = 2^{79.9}$ different couples, that is approximately $2^{62.4}$ different sets of 8 tuples. Since each of these sets satisfies the required properties with probability $2^{-94 \cdot 8 + 128} = 2^{-624}$, he has a good probability to find 8 different tuples with the required property. The cost to generate these $2^{48.9}$ (plaintexts, ciphertexts) pairs is of $2^{48.9}$ oracle-queries (with the assumption 1 oracle-query \simeq 1 encryption). On the other hand, the cost to generate these 8 tuples for the shortcut player is of $8 \cdot 2^{44} = 2^{47}$ (which is smaller). We emphasize that the cost of the generic player is higher than the cost of the shortcut player is satisfied for any value n with $n \geq 8$.

Finally, the same strategy can be used to extend the 7-round known-key distinguisher of App. F.1 in order to set up a 8-round known-key distinguisher with a time complexity of $2^{21.6}$. All the details are given in App. F.2.

6.1 The Computational Cost of Generic Player is Not Negligible!

Until now, we haven't considered the (further) cost of the generic player to identify the n tuples with the required relationship \mathcal{R} that he must send to the verifier. That is, we have only considered the cost (as number of oracle-queries) to generate a sufficient number of (plaintexts, ciphertexts) pairs to guarantee that n tuples with the required properties exist with a good probability. However, note that the player has to identify the n tuples with the required properties before to send them to the verifier. As we show in App. E.1, the computational cost of this step is not negligible. In particular, we propose a modified version of the attack presented in Sect. 5.1 that allows to find the required n -tuples and to minimize the total computational cost. As a final result, it follows that if the cost of this step is taken into account, then $n = 3$ tuples are sufficient to set up our distinguisher on 9 rounds of AES. We refer to App. E.1 for all the details.

7 10-Round Distinguisher of AES - Full AES-128

Using the same strategy proposed by Gilbert in [8], we set up our 10-round distinguisher by extending the 8-round one presented in [11] and in Sect. 4.3 both at the beginning and at the end, or equivalently by extending our 9-round distinguisher presented in the previous section at the beginning.

In the above defined known-key distinguisher scenario, the players have to send to the verifier n different tuples of (plaintext, ciphertext) pairs, that is $\{(p_i^1, c_i^1), (p_i^2, c_i^2)\}$ for $i = 0, \dots, n-1$, with the following properties:

1. there exists a key k^0 s.t. for each tuple there exists j for which the two plaintexts belong to the same coset of \mathcal{D}_j after one round, that is

$$\exists k^0 \text{ s.t. } \forall i = 0, \dots, n-1, \quad \exists j \in \{0, \dots, 3\} \text{ s.t. } R_{k^0}(p_i^1) \oplus R_{k^0}(p_i^2) \in \mathcal{D}_j;$$

2. there exists a key k^{10} s.t. for each tuple there exists l for which the two ciphertexts belong to the same coset of \mathcal{M}_l one round before, that is

$$\exists k^{10} \text{ s.t. } \forall i = 0, \dots, n-1, \quad \exists l \in \{0, \dots, 3\} \text{ s.t. } R_{k^{10}}^{-1}(c_i^1) \oplus R_{k^{10}}^{-1}(c_i^2) \in \mathcal{M}_l.$$

We stress that the keys k^0 and k^{10} must be equal for all the tuples, otherwise it is straightforward to generate tuples with the required properties (same argumentation of the 9-round case). However, a difference with the previous 9-round distinguisher arises. In the previous case, the verifier must verify the existence of a single key (by finding it, if exists), since the property on the plaintexts can be verified directly on them without guessing any secret-key material. For the 10-round case instead, the verifier has to check the existence of both k^0 and k^{10} . Thus, two possible scenarios can be considered and studied:

1. no key-schedule holds - k^0 and k^{10} are independent;
2. AES key-schedule among k^0 and k^{10} .

Intuitively, the second case (i.e. with key schedule) is harder than the first one (i.e. without key schedule) for the generic player, since a further property must be verified. In other words, the time required by this player to generate the tuples for the second scenario is not lower than for the first one, or in other words the probability of success in the second scenario is not higher than in the first one.

Before we present this distinguisher in detail, we highlight that the same strategy can be used to extend the 7-round known-key distinguisher of App. F.1 both at the beginning and at the end in order to set up the best 9-round known-key distinguisher from the computational point of view - its time complexity is approximately of 2^{23} computations. All the details are given in App. F.3.

The Shortcut Player. First of all, we study the computational cost of the player who knows the key. For this player, the two scenarios (with/without key schedule) are completely identical. Indeed, using the 8-round distinguisher described in [11] and in Sect. 4.3, he is able to generate n tuples that satisfy all the conditions (included the key schedule without any additional cost). The computational cost of this player is well approximated by $n \cdot 2^{44}$ computations.

7.1 Independent Subkeys: No Key Schedule

The strategies used by the verifier and by the generic player depend on which scenario one considers, that is depend on the fact that the two keys k^0 and k^{10} are independent or that a key schedule holds. Following the same strategy adopted by Gilbert in [8], as first case we assume that these two keys are independent.

The Generic Player. For the 9-round distinguisher, the best strategy that the generic player could adopt was to choose plaintexts in the same coset of \mathcal{D}_j , in order to fulfill the requested property on the plaintexts. The idea is simply to adapt this strategy for this case, that is the idea is to choose plaintexts such that the condition on the plaintexts is fulfilled with probability 1.

To do this, the generic player must fix a random key \hat{k} , and computes for a certain $j \in \{0, \dots, 3\}$ and for a random $a \in \mathcal{D}_j^\perp$ the following set:

$$D_a := R_{\hat{k}}^{-1}(\mathcal{D}_j \oplus a). \quad (7)$$

The idea is choose/use plaintexts in this set D_a just defined. In other words, the player works in the same way described for the 9-round distinguisher but using D_a defined above instead of a coset of \mathcal{D}_j . The corresponding ciphertxts are simply got by oracle-queries. Since the cardinality of a coset of \mathcal{D}_j is 2^{32} , the computation of a set D_a requires $2^{32+4} = 2^{36}$ S-Box look-ups for each coset $\mathcal{D}_j \oplus a$. Note that if the player needs more than 2^{32} (plaintext, ciphertxt) pairs, he simply chooses another $a' \in \mathcal{D}_j^\perp$ (or/and another j) and, using the *same* key \hat{k} , he computes the corresponding set $D_{a'}$ defined as before. We emphasize that the player must use always the same key \hat{k} to compute these sets, in order to fulfill the property on the plaintexts.

Given the set D_a , the idea is to use the same strategy presented for the 9-round distinguisher in the previous section in order to find the n tuples with

the required properties. Since the procedure to choose tuples such that the requirement on the ciphertexts is fulfilled is identical to the one presented for the 9-round distinguisher, we refer to that section for more details. We stress that given plaintexts in the same set D_a , the requirement on the plaintexts is always fulfilled since by construction there exists a key (which is \hat{k}) such that the plaintexts of each tuple belong to the same coset of \mathcal{D}_j after one round.

As a result, the strategy and the computational cost used to find these n tuples are (approximately) identical to the one presented in the previous section - note that the cost to compute the set D_a is negligible compared to the total cost. It follows that $n \geq 8$ tuples are sufficient for the case in which the cost of the generic player is approximated by the number of oracle-queries, while $n \geq 3$ tuples are sufficient for the case in which all the costs (oracle-queries + cost of the player) are considered. As before, we choose an (arbitrary) value of $n = 64$ in order to make the advantage of the shortcut player more significant.

The Verifier. Given n tuples, the verifier has to check the existence of keys k^0 and k^{10} as defined previously. Since no key schedule is considered, the idea is simply to work independently on the plaintexts (in order to find k^0) and on the ciphertexts (in order to find k^{10}). Since the verifier performs two independent attacks (as described in Sect. 5.1) on the plaintexts and on the ciphertexts, the cost doubles with respect to the 9-round case. As for the previous case, note that the verification cost is much lower than the players costs.

7.2 The Key Schedule Case

The scenario in which a key schedule holds is more complicated to analyze. Before we present our strategy, we recall the one adopted by Gilbert to set up his 10-round distinguisher. First he considers the case of AES with independent subkeys - denoted by AES_{10}^* , and he presents a 10-round known-key distinguisher for AES^* . Then, he simply observes that this known-key distinguisher on AES_{10}^* “*is obviously applicable without any modification to AES_{10} , i.e. the full AES-128*” (see Sect. 4.2 - page 221 of [8]). Using the same argumentation, we can easily conclude that also our distinguisher can be applied to real AES, i.e. to the case in which the key schedule holds. Indeed, as we have already pointed out, note that nothing changes for the shortcut player, while this scenario is more complicated for the generic player who doesn’t know the key, since a further condition on k^0 and k^{10} (the key schedule) is imposed. Even if it is possible to refer to previous results, here we show that a less number of tuples can be sufficient to set up this distinguisher in the case in which the key schedule holds.

The Verifier. Given n tuples, the verifier has to check the existence of k^0 and k^{10} that satisfy the AES key schedule and for which the properties on the plaintexts and on the ciphertexts are fulfilled. Working as before, the verifier can use several (equivalent) strategies, and here we focus on two of them.

In the first case, the idea is to work again independently on the plaintexts and on the ciphertexts, and find independently the two keys. Only as final step,

she checks if there exist keys k^0 and k^{10} (among the ones found previously) that satisfy the key schedule. In the second case, the idea is to work only on the plaintexts and to find k^0 such that the property on the plaintexts is satisfied. When a candidate for k^0 is found using the n tuples, the verifier finds k^{10} using the key schedule and checks if the requirement on the ciphertexts is satisfied.

For both these two cases, since on average only one key k^0 and one key k^{10} is found if the number of tuples n is greater or equal than 2, the computational cost for the verifier is comparable and well approximated by the cost of the (previous) case in which the subkeys are independent.

The Generic Player. When the key schedule holds, the strategy presented before for the generic player must be modified since it is no more the best one. Indeed, suppose this player fixes a key $k^0 = \hat{k}$ as before. It follows that the probability that \hat{k} (fixed) and a suitable k^{10} satisfy the key schedule is only 2^{-128} , which implies that the probability of success is very low.

For this reason, we present a modified strategy that he can use in this scenario. The idea is to look for plaintexts that maximize the number of keys k^0 and k^{10} for which the requirements are satisfied (included the key-schedule). If we consider two random pairs of texts (p^1, c^1) and (p^2, c^2) , there are on average 2^{34} keys k^0 such that $R_{k^0}(p^1) \oplus R_{k^0}(p^2) \in \mathcal{D}_j$ and 2^{34} keys k^{10} such that $R_{k^{10}}^{-1}(c^1) \oplus R_{k^{10}}^{-1}(c^2) \in \mathcal{M}_l$ for certain j and l . Thus, an initial key and a final one that satisfy the key schedule exist only with probability $(2^{34})^2 \cdot 2^{-128} = 2^{-60}$. Consider instead two plaintexts that belong to the same coset of \mathcal{D}_j . Since a coset of \mathcal{D}_j is mapped into a coset of \mathcal{C}_j (see Sect. 3.1), after one round the two texts belong to the same coset of \mathcal{C}_j for all the possible keys with probability 1. At the same time, it is possible to prove that there exist 2^{106} keys for which the two plaintexts belong to the same coset of $\mathcal{C}_j \cap \mathcal{D}_l \subseteq \mathcal{D}_l$ after one round.

Proposition 1. *Let p^1 and p^2 two plaintexts that belong to the same coset of \mathcal{D}_j for a certain j , that is $p^1 \oplus p^2 \in \mathcal{D}_j$. Moreover, assume that $p^1 \oplus p^2 \notin \mathcal{D}_j \cap \mathcal{C}_L$ for each $L \subseteq \{0, 1, 2, 3\}$ with $|L| \leq 3$. Then there exist on average 2^{106} different keys k such that $R_k(p^1) \oplus R_k(p^2) \in \mathcal{D}_l$ for a certain $l \in \{0, 1, 2, 3\}$.*

The proof is given in App. G. Thus, if one considers two couples (p^1, c^1) and (p^2, c^2) that satisfy the hypothesis of the previous proposition (in particular, $p^1 \oplus p^2 \in \mathcal{D}_l$ for a certain l), then there are on average 2^{106} keys k^0 such that $R_{k^0}(p^1) \oplus R_{k^0}(p^2) \in \mathcal{D}_i$ and 2^{34} keys k^{10} such that $R_{k^{10}}^{-1}(c^1) \oplus R_{k^{10}}^{-1}(c^2) \in \mathcal{M}_j$. It follows that there exist on average $2^{106} \cdot 2^{34} \cdot 2^{-128} = 2^{12}$ combinations of initial and final subkeys k^0 and k^{10} that satisfy the key schedule. Even if we don't exclude better strategies, we conjecture that this is one of the best strategy that this player can use in order to maximize the number of keys (k^0, k^{10}) that satisfy the key schedule and the other required properties.

Number n of Tuples: Oracle-Queries. Starting by these considerations, we show that $n = 4$ tuples are sufficient to set up the distinguisher when a key-schedule holds and when only the number of oracle-queries is considered (remember that for independent subkeys n must be equal or greater than 8). First of all,

working as in Sect. 6, note that given n tuples (where the plaintexts are chosen as described previously), the probability that there exist keys (k^0, k^{10}) that satisfy the key schedule and for which the properties on the plaintexts/ciphertexts are satisfied is $2^{12 \cdot n} \cdot 2^{-128 \cdot (n-1)} = 2^{-116 \cdot n + 128}$ instead of $2^{-94 \cdot n + 128}$ (see Eq. (6)), since for each couple there are only 2^{12} possible combinations¹⁴ of keys (k^0, k^{10}) instead of 2^{34} . Thus, using similar argumentation as before, in order to win the game the generic player must consider s different couples, where s is given by

$$s \simeq 2^{116 - \frac{128}{n}} \cdot (n!)^{\frac{1}{n}}. \quad (8)$$

In particular, he has to consider at least $2^{85.14}$ different couples in order to find $n = 4$ tuples that satisfy the requirements. Since each coset of \mathcal{D}_j contains 2^{32} different plaintexts (or approximately 2^{63} different couples), he must consider approximately $2^{22.14}$ different cosets of \mathcal{D}_j defined as in (7), for a total of $2^{54.14}$ (plaintexts, ciphertexts) pairs. Thus, in the case in which the cost of the generic player is approximated by the number of oracle-queries, his cost is of approximately $2^{54.2}$ oracle-queries. On the other hand, the cost for the shortcut player to generate the same number of different tuples with the required properties is approximately of $4 \cdot 2^{44} = 2^{46}$ computations, which is lower.

Number n of Tuples: Oracle-Queries and Cost of Generic Player. As for the 9-round case, if one considers all the costs (that is the number of oracle-queries and the computational cost of the generic player), it turns that a lower number of tuples (precisely $n = 2$) is sufficient. We refer to App. E.2 for details.

8 12-Round Distinguisher of AES

As one of the major contributions of this paper, in this section we present the first known-key distinguisher for 12 rounds of AES. This distinguisher is obtained by extending the previous 10-round distinguisher both at the end and at the beginning, or equivalently by extending two times at the end and at the beginning the 8-round known-key distinguisher presented in [11] and in Sect. 4.3. We highlight that this is the *first known-key distinguisher for full AES-192* (and on 12 rounds of AES-128, i.e. full AES-128 with two more rounds) and it also provides a counterexample to the claims made in [8].

In the know-key distinguisher scenario, the players have to send to the verifier n different tuples of (plaintext, ciphertext) pairs, that is $\{(p_i^1, c_i^1), (p_i^2, c_i^2)\}$ for $i = 0, \dots, n - 1$, with the following properties:

1. there exist keys k^0, k^1 s.t. for each tuple there exists j for which the two plaintexts belong to the same coset of \mathcal{D}_j after two rounds, that is

$$\exists k^0, k^1 \text{ s.t. } \forall i = 0, \dots, n-1 \quad \exists j \in \{0, \dots, 3\} \text{ s.t. } R_{k^0, k^1}^2(p_i^1) \oplus R_{k^0, k^1}^2(p_i^2) \in \mathcal{D}_j;$$

¹⁴ If two combinations $(\tilde{k}^0, \tilde{k}^{10})$ and $(\hat{k}^0, \hat{k}^{10})$ satisfy the key schedule, then they are equal with prob. 2^{-128} (e.g. if $\tilde{k}^0 = \hat{k}^0$ then $\tilde{k}^{10} = \hat{k}^{10}$ due to the key schedule).

2. there exist keys k^{11}, k^{12} s.t. for each tuple there exists l for which the two ciphertexts belong to the same coset of \mathcal{M}_l two rounds before, that is

$$\exists k^{11}, k^{12} \text{ s.t. } \forall i = 0, \dots, n-1 \exists l \in \{0, \dots, 3\} \text{ s.t. } R_{k^{11}, k^{12}}^{-2}(c_i^1) \oplus R_{k^{11}, k^{12}}^{-2}(c_i^2) \in \mathcal{M}_l;$$

where $R_{k^0, k^1}^2(\cdot) \equiv R_{k^1}(R_{k^0}(\cdot))$ and $R_{k^{11}, k^{12}}^{-2}(\cdot) \equiv R_{k^{11}}^{-1}(R_{k^{12}}^{-1}(\cdot))$.

As for the known-key distinguisher for 10-round AES, two scenarios can be considered, that is the case of independent subkeys and the case in which the key schedule holds. For the following, we consider only the first scenario, that is we limit ourselves to present a known-key distinguisher for 12-round of AES with independent subkeys. However, using similar argumentation as before, we claim that the same distinguisher can be applied to the case in which the key schedule holds. In particular, we remember that nothing change for the shortcut player (who knows the key) in this second case, while the challenge becomes much harder for the other player.

The strategy used by the players and by the verifier is very similar to the one presented for the 10-round distinguisher in the case of no key-schedule. Thus, we refer to the previous section for all the details, and we limit here to highlight the idea and the major differences.

The Two Players. Exactly as before, the shortcut player can generate n tuples with the required properties for a cost of $n \cdot 2^{44}$ computations.

The generic player exploits the same strategy proposed for the 10-round distinguisher with no key-schedule. First he fixes random keys \hat{k}^0, \hat{k}^1 and \hat{k}^{12} , and using the keys \hat{k}^0 and \hat{k}^1 , he computes the set $D_a = R_{\hat{k}^0}^{-1}(R_{\hat{k}^1}^{-1}(\mathcal{D}_j \oplus a))$. Similar to the previous case, the idea is to work with plaintexts in the same set D_a . He then gets the corresponding ciphertexts by oracle-queries, and the idea is simply to decrypt them using the key \hat{k}^{12} . As a result, using the same strategy proposed for the 9- and 10-round distinguisher, he can construct n tuples that satisfy the relation \mathcal{R} , that is he is able to find n tuples for which a common key k^{11} exists such that the requirement on the ciphertexts is satisfied.

By analogous calculation as before, at least $n \geq 8$ tuples are sufficient to set up the distinguisher when only the number of the oracle-queries is considered.

The Verifier. When the verifier receives the n tuples, she can use the following strategy to check if the required properties are satisfied or not. First of all, since there is no key schedule, the verifier can work independently on k^0, k^1 (that is on the plaintexts) and on k^{11}, k^{12} (that is on the ciphertexts). Similarly to the previous cases where the verifier uses the key-recovery attack of Sect. 5.1 to find the keys, for this 12-round distinguisher the idea is to exploit the key-recovery attack presented in Sect. 5.2 to find (if exist) the four keys k^0, k^1 and k^{11}, k^{12} .

We present in details the verification procedure for the ciphertexts case (analogous for the plaintexts). Given the first tuple and using the strategy described in Sect. 5.2, the verifier guesses eight bytes of the final subkey k^{12} (two diagonals), decrypts partially, and finds 2^{34} values for eight bytes of k^{11} working as in Sect. 5.1, for a total of $2^{34} \cdot 2^{64} = 2^{98}$ candidates. Then, she eliminates wrong

candidates by testing them using the other tuples - to reduce the computational cost, she can work independently on each column of k^{11} . Note that the probability that found subkeys k^{11} and k^{12} satisfy the required property for another tuple is $4 \cdot 2^{-32} = 2^{-30}$. Thus, using other four tuples, with high probability the verifier finds approximately only one pair of subkeys k^{11} and k^{12} for which the property on the ciphertexts is satisfied (note $2^{98} \cdot (2^{-30})^4 = 2^{-22}$). The cost of this step is of 2^{76} table look-ups (using the pre-computation phase), as shown in Sect. 5.2 or in [10] in more details. The remaining eight bytes of k^{11} and of k^{12} and the subkeys k^0 and k^1 can be found in a similar way.

As a result, given 5 different tuples, the total cost for this attack is approximately of $4 \cdot 2^{76} = 2^{78}$ table look-ups (using the pre-computation phase). When the verifier has found possible candidates for the four keys, she checks that also the other $n - 5$ tuples satisfy the relation \mathcal{R} for the found keys. In conclusion, given $n \geq 5$ tuples, the total cost for the verifier can be approximated at $2 \cdot (2^{78} + 2^6 \cdot (n - 5))$ table look-ups. If $n \ll 2^{72}$, then the computational cost of the verifier is approximately $2^{71.1}$ twelve-round encryptions.

Number n of Tuples. As we have just seen, it is possible to set up the distinguisher for n equal or greater than 8. However, if $n = 8$ then the cost of the shortcut player (2^{47} computations) is much lower than the cost of the verifier ($2^{71.1}$ computations), which is not consistent with the given definition of known-key distinguisher (see Sect. 4.1). Indeed, by definition the verification cost must be less than the cost of the shortcut players (and so the cost of the generic player), that is the entire cost of the distinguisher (computational cost of the shortcut player + verification cost) must be well approximated by the cost of the shortcut player. In order to fulfill this condition, it is sufficient to choose a number of tuple n that satisfy the condition $n \cdot 2^{44} \gg 2^{71.1}$ (and $n \ll 2^{72}$). It follows that a good (arbitrary) choice for this distinguisher¹⁵ could be $n \geq 2^{38}$.

In conclusion, to win the game, the two players have to send 2^{38} tuples of (plaintext, ciphertext) pairs with the required properties. The cost for the shortcut player is of 2^{82} computations, while the verification cost is of $2^{71.1}$ computations. Note that even if this result is obtained considering only the number of the oracle-queries and the case of independent subkeys, it holds also for the cases in which all the costs are considered and/or the key schedule holds. Indeed, it is simple to observe that also in these cases (1) the choice of a suitable number n is more influenced by the request that the verification cost is lower than the cost to generate the n tuples and (2) the game becomes harder for the generic player, while nothing changes for the shortcut one.

Finally, in App. H we show that a similar strategy can be (theoretically) used to extend both at the end and at the beginning the Gilbert's 10-round known-key distinguisher, obtaining in a similar way a 12-round known-key distinguisher based on the balance property.

¹⁵ By previous analysis, we remember that the cost of the shortcut player is always lower than the cost of the generic player for each value of n that satisfies $n \geq 8$.

9 Infeasibility of a 14-round Known-Key Distinguisher

In this paper, we have shown that Gilbert’s known-key distinguisher model can lead to results on more rounds than previous expected. Even though the core distinguisher remains at 8 rounds, 12 instead of 10 rounds are achieved. This may raise the question: *How meaningful is this distinguisher model?* We claim that it appears meaningful in the sense that it does not seem to allow results on an arbitrary number of rounds.

We analyze this claim in more details, assuming by contradiction the existence of a meaningful known-key distinguisher on 14 rounds of AES (to be meaningful, we assume that the probability of the shortcut player to win the game is higher than the one of the generic player). The main criticism in order to extend a known-key distinguisher both at the end and at the beginning as in the Gilbert model regards the computational cost to verify the existence of keys such that the n tuples of (plaintexts, ciphertexts) pairs satisfy the relation \mathcal{R} . We stress that the verification cost must be lower than the players costs. Thus, consider the known-key distinguishers that exploit the balance property or a truncated differential trail. In order to extend 1 round at the beginning and at the end, a classical key recovery attack - as the integral attack [4] and the truncated differential attack [10] - is sufficient for this task. In order to extend 2 rounds as for the distinguishers presented in this paper, the idea is to use a key recovery attack with an extension at the end, e.g. the integral attack with an extension at the end [4] and the truncated differential attack of Sect. 5.2. In a similar way, in order to extend for $r \geq 3$ rounds, one needs a key-recovery attack with two extensions at the end, that is more than a single one¹⁶. Since balance and/or truncated differential attacks with this property don’t exist in literature for AES-128¹⁷ and since it seems very unlikely to set up them without guessing an entire subkey (which leads to a brute force attack), we claim that it is not possible to extend the 8-round distinguishers currently present in the literature for more than 4 rounds, that is 2 rounds at the end and 2 at the beginning. We leave the open problem to confute our claims for future investigations.

10 Conclusion, Discussion, and Open Problems

In this paper, we improve all the known-key distinguishers currently present in the literature for AES up to 10 rounds of AES and we set up the first known-key distinguisher on 12 rounds of AES, by extending distinguishers based on truncated differential trails using the technique proposed by Gilbert in [8]. In particular, we set up a 9-round known-key distinguisher for AES with a time complexity of 2^{23} , a 10-round known-key distinguisher with a time complexity of 2^{50} , and a 12-round known-key distinguisher with a time complexity of 2^{82} .

¹⁶ In App. I we show why attacks with both an extension at the end and at the beginning are completely useless for this scope, taking as example the partial-sum attack on 7 rounds of AES-128 presented in [6] - also known as herds attack.

¹⁷ Note that for AES-256 it is possible to set up such attacks by simply guessing an entire subkey. However, since the complexity of such attacks is higher than 2^{128} (an entire subkey is guessed), the verification cost is higher than the costs of the players.

In order to set up our known-key distinguishers presented in this paper, for each case we consider the *intuitively* best strategy that the generic player can exploit to win the game, giving strong argumentations to justify it. In this sense, we adopt the same strategy also used by authors of [15] and [9] to present their 7- and 8-rounds known-key distinguishers. Indeed, also for these works, authors exploit the intuitively best strategy that the generic player can use, while no formal proof of this fact is given. In other words, both in these papers and in our work, we conjecture that the (intuitively) best strategy of the generic player is (1) to consider plaintexts that satisfy the required properties (e.g. that belong to the same coset of a diagonal space \mathcal{D}_i), and (2) to look for the ones for which the required properties on the corresponding ciphertexts are also satisfied. We leave the problem to formally prove the statements that our distinguisher works (equivalently, that our conjectures are correct) for future investigation. We emphasize that for the cases of [15] and [9], a formal proof has been provided only after approximately two years (see [17] for details).

Drawing conclusions about AES because of our concrete results seems premature. Perhaps more interestingly our example of AES emphasizes a gap in our understanding of distinguishers in open-key models (which includes chosen-key and known-key models). Our results exploit subtle properties in the known-key model which seem currently not possible in the chosen-key model. Indeed, whereas in the chosen-key model the best result (excluding related-key variants as in [2]) is on 9 rounds [7], we now have results on 12 rounds in the known-key model. This seems to hint towards more possibilities in the chosen-key model.

Acknowledgements. The work in this paper has been partially supported by the Austrian Science Fund (project P26494-N15).

References

1. E. Andreeva, A. Bogdanov, and B. Mennink, “Towards Understanding the Known-Key Security of Block Ciphers,” in *FSE 2013*, ser. LNCS, vol. 8424, 2014, pp. 348–366.
2. A. Biryukov, D. Khovratovich, and I. Nikolić, “Distinguisher and Related-Key Attack on the Full AES-256,” in *CRYPTO 2009*, ser. LNCS, vol. 5677, 2009, pp. 231–249.
3. R. Canetti, O. Goldreich, and S. Halevi, “The Random Oracle Methodology, Revisited,” *Journal ACM*, vol. 51, no. 4, pp. 557–594, 2004.
4. J. Daemen, L. R. Knudsen, and V. Rijmen, “The Block Cipher Square,” in *FSE 1997*, ser. LNCS, vol. 1267, 1997, pp. 149–165.
5. J. Daemen and V. Rijmen, *The Design of Rijndael: AES - The Advanced Encryption Standard*, ser. Information Security and Cryptography. Springer, 2002.
6. N. Ferguson, J. Kelsey, S. Lucks, B. Schneier, M. Stay, D. Wagner, and D. Whiting, “Improved Cryptanalysis of Rijndael,” in *FSE 2000*, ser. LNCS, vol. 1978, 2001, pp. 213–230.
7. P. Fouque, J. Jean, and T. Peyrin, “Structural Evaluation of AES and Chosen-Key Distinguisher of 9-Round AES-128,” in *CRYPTO 2013*, ser. LNCS, vol. 8042, 2013, pp. 183–203.

8. H. Gilbert, “A Simplified Representation of AES,” in *ASIACRYPT 2014*, ser. LNCS, vol. 8873, 2014, pp. 200–222.
9. H. Gilbert and T. Peyrin, “Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations,” in *FSE 2010*, ser. LNCS, vol. 6147, 2010, pp. 365–383.
10. L. Grassi, C. Rechberger, and S. Rønjom, “Subspace Trail Cryptanalysis and its Applications to AES,” *IACR Transactions on Symmetric Cryptology*, vol. 2016, no. 2, pp. 192–225, 2017.
11. J. Jean, M. Naya-Plasencia, and T. Peyrin, “Multiple limited-birthday distinguishers and applications,” in *SAC 2013*, ser. LNCS, vol. 8282, 2014, pp. 533–550.
12. L. Knudsen and D. Wagner, “Integral Cryptanalysis,” in *FSE 2002*, ser. LNCS, vol. 2365, 2002, pp. 112–127.
13. L. R. Knudsen and V. Rijmen, “Known-Key Distinguishers for Some Block Ciphers,” in *ASIACRYPT 2007*, ser. LNCS, vol. 4833, 2007, pp. 315–324.
14. M. Lamberger, F. Mendel, M. Schl  ffer, C. Rechberger, and V. Rijmen, “The Rebound Attack and Subspace Distinguishers: Application to Whirlpool,” *J. Cryptology*, vol. 28, no. 2, pp. 257–296, 2015.
15. F. Mendel, T. Peyrin, C. Rechberger, and M. Schl  ffer, “Improved Cryptanalysis of the Reduced Gr  stl Compression Function, ECHO Permutation and AES Block Cipher,” in *SAC 2009*, ser. LNCS, vol. 5867, 2009, pp. 16–35.
16. F. Mendel, C. Rechberger, M. Schl  ffer, and S. S. Thomsen, “The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Gr  stl,” in *FSE 2009*, ser. LNCS, vol. 5665, 2009, pp. 260–276.
17. I. Nikoli  c, J. Pieprzyk, P. Sokolowski, and R. Steinfeld, “Known and Chosen Key Differential Distinguishers for Block Ciphers,” in *ICISC 2010*, ser. LNCS, vol. 6829, 2011, pp. 29–48.
18. L. Wei, T. Peyrin, P. Sokolowski, S. Ling, J. Pieprzyk, and H. Wang, “On the (In)Security of IDEA in Various Hashing Modes,” in *FSE 2012*, ser. LNCS, vol. 7549, 2012, pp. 163–179.

A Considerations on Gilbert’s 10-round Distinguisher

In this section, we show that the verification cost for the Gilbert’s 10-round distinguisher is not lower than the cost of the shortcut player, i.e. the relationship \mathcal{R} defined in [8] is not efficiently checkable, despite what it is claimed.

We refer to Sect. 4.4 for a detailed presentation of the Gilbert’s known-key distinguisher. We limit to recall that the two players have to find 2^{64} (plaintext, ciphertext) pairs, i.e. (p^i, c^i) for $i = 0, \dots, 2^{64} - 1$, s.t. there exist keys k^0 and k^{10} for which the following sums $\bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p^i) = 0$ and $\bigoplus_{i=0}^{2^{64}-1} R_{k^{10}}^{-1}(c^i) = 0$ are satisfied. When the verifier receives the set of 2^{64} (plaintext, ciphertext) pairs from the players, she checks if the required properties are satisfied or not by finding the two keys. Since no key-schedule holds, the verifier can work independently on k^0 and k^{10} . Both for k^0 and k^{10} , Gilbert proposes to work on 4 bytes of the key at the same time, that is to work on entire column in the case of k^{10} (the first operation of $R_{k^{10}}^{-1}(\cdot)$ to compute is $MC^{-1}(\cdot)$) and to work on entire diagonal in the case of k^0 (the first operations of $R_{k^0}(\cdot)$ to compute are $MC \circ SR(\cdot)$). Since there are 4 columns/diagonals and $(2^8)^4 = 2^{32}$ possible values for each of them, the total cost is of approximately $2 \cdot 4 \cdot 2^{32} \cdot 2^{64} = 2^{99}$ S-Box look-ups.

As we are going to show, it is not necessary to work on 4 bytes of the subkeys k^0 and k^{10} simultaneously, but it is possible to find k^0 and k^{10} working on single bytes (independently of the others). The idea for k^{10} (very common in the literature) is simply to change the positions of the final MixColumns operation and of the final AddRoundKey operation, using the fact that the MixColumns is linear. In this way, the verifier can work on single byte of $\hat{k}^{10} \equiv MC^{-1}(k^{10})$ - the existence of the key \hat{k}^{10} obviously implies the existence of k^{10} .

For the case of k^0 , remember that a set of balanced texts is mapped into a set of balanced texts by the MixColumns operation (since it is linear), that is given a set of texts $\{t^i\}_{i=0,\dots,n}$ then $\bigoplus_{i=0}^n t^i = 0$ if and only if $\bigoplus_{i=0}^n MC(t^i) = 0$. Thus, in order to verify that $\bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p^i) = 0$, the verifier can simply check if the condition $\bigoplus_{i=0}^{2^{64}-1} \text{S-Box}(p_{j,l}^i \oplus k_{j,l}^0) = 0$ holds for each byte, i.e. $\forall j, l = 0, \dots, 3$. It follows that the verifier can work on single byte of k^0 .

Using this strategy, the verification cost can be approximated by $2 \cdot 16 \cdot 2^8 \cdot 2^{64} = 2^{77}$ S-Box look-ups, or equivalently $2^{69.36}$ ten-round AES encryptions - better than the strategy proposed in [8], which is still (much) higher than the cost of the shortcut player even in this case (approximately of 2^{64} ten-round AES encryptions). Even if we don't exclude that better strategies exist, it follows that using a (classical) integral attack the claim made in [8] "*the overall complexity of checking \mathcal{R} is strictly smaller than $N = 2^{64}$ AES $_{10}^*$ operations*" (see Sect. 4.2 - page 218) is false, i.e. the computational cost of the verifier is not negligible w.r.t. the cost of the players. On the other hands, in App. H we show how to modify the Gilbert's 10 rounds distinguisher in order to overcome this problem.

B A possible Variant of Gilbert's Distinguisher - Details

In Sect. 4.4, we proposed a possible variant of the Gilbert's distinguisher - that also applies to all our proposed distinguishers present in the paper - which allows to better understand it. Consider the case in which the two players have to send to the verifier the N -tuple that verify the required relation \mathcal{R} together with the subkeys for which such relation is satisfied.

In more details, assume that the relationship \mathcal{R} depends on the existence of subkey(s) such that the required property is not directly verified on the plaintexts or/and on the ciphertexts but one (or more) round(s) before/after. As an example, consider the 10-round known-key distinguisher proposed by Gilbert and based on the balance propoerty. In this case, the two players have to send 2^{64} (plaintexts, ciphertexts) pairs, i.e. (p^i, c^i) for $i = 0, \dots, 2^{64}-1$ and the two subkeys k^0 and k^{10} such that the plaintexts are uniformly distributed after one round in the cosets of \mathcal{D}_I and the ciphertexts are uniformly distributed one round before in the cosets of \mathcal{M}_J , or equivalently that $\bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p^i) = \bigoplus_{i=0}^{2^{64}-1} R_{k^{10}}^{-1}(c^i) = 0$.

In this case, the task of the verifier is to check if the relation \mathcal{R} is satisfied (or not) *only* for the subkeys she received by the players. It follows that her computational cost is negligible, in the sense that it is comparable to the computational cost of the 8-round integral distinguisher presented in [8] where

the required property \mathcal{R} can be directly verified on the plaintexts/ciphertexts (or equivalently comparable to the computational costs of the other known-key distinguishers present in literature up to 8 rounds). Here we show in details why such distinguisher is meaningless.

The main problem of such a distinguisher regards the fact that it can be set up for any number of rounds. To explain this problem, consider our known-key distinguisher on $r = 8 + 2 \cdot r'$ rounds of AES, for $r' \geq 1$ (the same considerations apply e.g. to the Gilbert integral distinguisher). The players have to send to the verifier n different tuples of (plaintext, ciphertext) pairs, that is $\{(p_i^1, c_i^1), (p_i^2, c_i^2)\}$ for $i = 0, \dots, n-1$, and $2 \cdot r'$ subkeys $k^0, \dots, k^{r'-1}$ and $k^r, \dots, k^{r-r'+1}$ such that

1. for each tuple there exists j for which the two plaintexts belong to the same coset of \mathcal{D}_j after r' rounds, that is

$$\forall i = 0, \dots, n-1 \exists j \in \{0, \dots, 3\} \quad \text{s.t.} \quad R_{k^0, \dots, k^{r'-1}}(p_i^1) \oplus R_{k^0, \dots, k^{r'-1}}(p_i^2) \in \mathcal{D}_j;$$

2. for each tuple there exists l for which the two ciphertexts belong to the same coset of \mathcal{M}_l r' rounds before, that is

$$\forall i = 0, \dots, n-1 \exists l \in \{0, \dots, 3\} \quad \text{s.t.} \quad R_{k^r, \dots, k^{r-r'+1}}^{-1}(c_i^1) \oplus R_{k^r, \dots, k^{r-r'+1}}^{-1}(c_i^2) \in \mathcal{M}_l,$$

where $R_{k^0, \dots, k^{r'-1}}(\cdot) \equiv R_{k^{r'-1}} \circ \dots \circ R_{k^0}(\cdot)$ and $R_{k^r, \dots, k^{r-r'+1}}^{-1}(\cdot) \equiv R_{k^{r-r'+1}}^{-1} \circ \dots \circ R_{k^r}^{-1}(\cdot)$.

Consider now the costs of the verifier and of the two players. As we have already said, the cost of the verifier is negligible, since she has to check if the relation \mathcal{R} is satisfied only for the received subkeys. The cost of the shortcut player is approximately of $n \cdot 2^{44}$ computations for n tuples, since he can use the rebound attack (see Sect. 4.2 and App. C for details) to find them. The generic player instead can use the strategy proposed in details Sect. for the 10 rounds case and in Sect. for the 12 rounds one in order to win the game. Such strategy allows the player to find plaintexts (or ciphertexts) that satisfy the required condition with negligible computational cost. However, the only way to satisfy both the conditions (i.e. the relation \mathcal{R}) is to test the texts found in the first step by brute force. It follows that increasing the number n of required tuples (and the number of rounds r'), the computational cost of the generic player grows faster than the cost of the shortcut player. In other words, even if we don't exclude that a better strategy exists, it seems hard that the cost of the generic player can be lower than the cost of the shortcut one. By definition of known-key distinguisher given in Sect. 4.1, it follows that such a distinguisher can be set up for any number of rounds (of AES), which is meaningless according to our definition given in Sect. 9.

C The Rebound Attack - Details

In the 7- and 8-round known-key distinguishers proposed in [15] and [9], the goal of the two players is to find two pairs of (plaintexts, ciphertexts) - i.e. (p^1, c^1)

and (p^2, c^2) - with the following property: the two plaintexts belong to the same coset of \mathcal{D}_i - i.e. $p^1 \oplus p^2 \in \mathcal{D}_i$ - and the two ciphertexts belong to the same coset of \mathcal{M}_i - i.e. $c^1 \oplus c^2 \in \mathcal{M}_i$ - for a fixed index i .

Consider the known-key distinguisher setting of the two players proposed in Sect. 4.1. In order to win the proposed game, the technique that the shortcut player (i.e. the player that knows the key) should use is the *Rebound Attack*.

The rebound attack was proposed in [16] for the cryptanalysis of AES-based hash functions. The rebound attack consists of two phases, called inbound and outbound phase. According to these phases, the internal permutation of the hash function is split into three sub-parts. Let f be the permutation, then we get $f = f_{fw} \circ f_{in} \circ f_{bw}$. The part of the inbound phase is placed in the middle of the permutation and the two parts of the outbound phase are placed next to the inbound part. In the outbound phase, two high-probability (truncated) differential trails are constructed, which are then connected in the inbound phase.

Since the rebound attack is a differential attack, as first thing an attacker needs to construct a “good” (*truncated*) *differential trail*. A good trail used for a rebound attack should have a high probability in the outbound phases and can have a rather low probability in the inbound phase. In particular, two properties are important: first, the system of equations that determine whether a pair follows the differential trail in the inbound phase should be under-determined. This contributes to the fact that many solutions (starting points for the outbound phase) can be found efficiently by using guess-and-determine strategies. Second, the outbound phases need to have high probability in the outward direction.

When searching for solutions of the *inbound part*, the attacker first guesses some variables such that the remaining system is easier to solve. Hence, the inbound phase of the attack is similar to message modification in an attack on hash functions. The available freedom in terms of the actual values of the internal variables is used to find a solution deterministically or with a very high probability.

In the *outbound phase*, the attacker verifies whether the solutions of the inbound phase also follow the differential trail in the outbound parts. Note that in the outbound phase, there are usually only a few or no free variables left. Hence, a solution of the inbound phase will lead to a solution of the outbound phase with a low probability. Therefore, the attacker aims for narrow (truncated) differential trails in the outbound parts, which can be fulfilled with a probability as high as possible (in the outward directions). The advantage of using an inbound phase in the middle and two outbound phases at the beginning and end is that one can construct differential trails with a higher probability in the outbound phase and at the same time cover a relatively high number of rounds.

The AES Case. Here we consider in details the strategy of the shortcut player for 7- and 8-round of AES. The truncated differential trails used for 7- and 8-round AES are depicted in Fig. 4. Referring to the 8-round trail (the 7-round case is analogous), the inbound phase is composed of the states from S_2 to S_5 , which are highlight in Fig. 5. The player chooses differences in 8 bytes, that is 4 bytes in S'_{start} (i.e. S_2 after the S-Box) and the 4 output bytes in S_{end} (i.e.




Fig. 4. 7- and 8-round differential paths for AES-128.




Fig. 5. A detail of the inbound phase (rounds 2 - 4) of the 8-round differential.

S_5 before the S-Box). Since ShiftRows and MixColumns are linear operations, the player can propagate these difference through these operations in order to compute S_{in} and S_{out} . We define the operations between these two states as *Super-SB*:

$$Super-SB(\cdot) := S-Box \circ ARK \circ MC \circ S-Box(\cdot), \quad (9)$$

where note that the key is known. The player has to look for two states S_{in} and S_{out} such that the differential trail is satisfied though this *Super-SB* operation. When the player finds these two states, he can easily compute the corresponding states S_2 and S_5 .

In the outbound phase, the player simply propagates the results found in the previous step in the backward and in the forward directions, and checks if they satisfy the entire differential trail.

As proved in [9], in the case of a perfect random permutation on average 2^{64} operations are required to find two (plaintexts, ciphertexts) pairs that satisfy the 8-round differential trail. Instead, in the AES case and when the initial and the final subspaces are fixed, it requires 2^{48} computations and 2^{32} memory.

D Known-Key Distinguishers for 7- and 8-Round of AES based on the Square Property

The 7- and the 8-round known-key distinguisher based on the balance property are a direct application of the 3- and 4-round secret-key distinguishers based on the square property and used in an inside-out fashion.

First of all, we re-call some definitions. Given a set of texts, we say that a byte X could be:

- Active (*A*): Every value in \mathbb{F}_{2^8} appears the same number of times in X ;
- Balance (*B*): The XOR of all values in X is 0;
- Constant (*C*): The value is fixed to a constant for all texts in X .

First, we formally define the 7- and the 8-round known-key distinguisher based on the balance property. Assume there are two players - one knows the key and the other not, and the verifier. To win the game, the players have to send to the verifier 2^n (plaintext, ciphertext) pairs, that is (p^i, c^i) for $i = 0, \dots, 2^n - 1$, such that the balance property holds both on the plaintexts and on the ciphertexts:

$$\bigoplus_{i=0}^{2^n-1} p^i = \bigoplus_{i=0}^{2^n-1} c^i = 0.$$

A suitable value of n is 56 for 7 rounds of AES and 64 for 8 rounds case.

What is the best strategy that the shortcut player can use to win the game? Consider 2^{32} plaintexts with one active diagonal (i.e. 4 bytes), and all the others 12 bytes constants. It is a well-known fact that the sum of 2^{32} corresponding ciphertexts after four rounds is equal to zero. A similar property holds in the decryption direction, that is the following integral properties hold:

$$\begin{bmatrix} B & B & B & B \\ B & B & B & B \\ B & B & B & B \\ B & B & B & B \end{bmatrix} \xleftarrow{R^{(-3)}} \begin{bmatrix} A & C & C & C \\ A & C & C & C \\ A & C & C & C \\ A & C & C & C \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} A & C & C & C \\ C & A & C & C \\ C & C & A & C \\ C & C & C & A \end{bmatrix} \xrightarrow{R^{(4)}} \begin{bmatrix} B & B & B & B \\ B & B & B & B \\ B & B & B & B \\ B & B & B & B \end{bmatrix}$$

where $R^{(4)}$ denotes 4 consecutive AES encryption rounds and $R^{(-3)}$ denotes 3 full AES decryption rounds.

Equivalent, this means that if one takes a coset of \mathcal{D}_i for a certain i , then the sum of the corresponding ciphertexts after 4 rounds is equal to zero. Again, if one takes a coset of \mathcal{C}_j for a certain j as the set of ciphertexts, the sum of the corresponding plaintexts 3 rounds before is equal to 0. Thus, starting in the middle with a coset of $\mathcal{D}_i \oplus \mathcal{C}_j$ for a certain i and j , then the sum of the corresponding plaintexts 3 rounds before and the ciphertexts after 4 rounds is equal to 0:

$$\begin{bmatrix} B & B & B & B \\ B & B & B & B \\ B & B & B & B \\ B & B & B & B \end{bmatrix} \xleftarrow{R^{-3}} \begin{bmatrix} A & C & C & C \\ A & A & C & C \\ A & C & A & C \\ A & C & C & A \end{bmatrix} \xrightarrow{R^4} \begin{bmatrix} B & B & B & B \\ B & B & B & B \\ B & B & B & B \\ B & B & B & B \end{bmatrix}.$$

This distinguisher on 7 rounds AES was proposed for the first time by Knudsen and Rijmen in [13], and it has a complexity of 2^{56} . In particular, in this case it is possible to prove that the probability of success of the player who doesn't know the key is strictly less than 1. In other words, the other player needs more computations to generate a set of (plaintexts, ciphertexts) pairs with the required properties.

Since a coset of \mathcal{C}_j is mapped into a coset of \mathcal{M}_j after one round with prob. 1, then given a coset of \mathcal{M}_j for a certain j as the set of ciphertexts, the sum of the corresponding plaintexts 4 rounds before is equal to 0. Equivalently, starting in the middle with a coset of $\mathcal{D}_i \oplus \mathcal{M}_j$ for a certain i and j , then the sum of the

corresponding plaintexts 4 rounds before and of the ciphertexts after 4 rounds is equal to 0:

$$\begin{bmatrix} B & B & B & B \\ B & B & B & B \\ B & B & B & B \\ B & B & B & B \end{bmatrix} \xleftarrow{R^{-4}} \mathcal{D}_i \oplus \mathcal{M}_j \oplus a \xrightarrow{R^4} \begin{bmatrix} B & B & B & B \\ B & B & B & B \\ B & B & B & B \\ B & B & B & B \end{bmatrix}$$

for a constant a . A similar distinguisher was proposed for the first time by Gilbert in [8], and it has a complexity of 2^{64} .

D.1 The 8-round Known-Key Distinguisher based on the Balance Property - A Formal Description.

To set up a known-key distinguisher on 8 rounds, the idea is simply to connect two 4-round trails in the middle and to choose a middle space $\mathcal{D}_i \oplus \mathcal{M}_j$ for i and j fixed (with $|i| = |j| = 1$). In the middle, the set $\mathcal{D}_i \oplus \mathcal{M}_j$ can be re-written as

$$\bigcup_{b \in \mathcal{D}_i} \mathcal{M}_j \oplus b = \bigcup_{a \in \mathcal{M}_j} \mathcal{D}_i \oplus a,$$

that is as union of cosets of the space \mathcal{D}_i or as union of cosets of the space \mathcal{M}_j .

Forward Direction. If one encrypts $\mathcal{D}_i \oplus a$ for four rounds ($a \in \mathcal{M}_j$), then the set $R^{(4)}(\mathcal{D}_i \oplus a)$ is a set of $(2^8)^4 = 2^{32}$ ciphertexts where each ciphertext belongs to a different coset of a mixed space \mathcal{M}_I of dimension 12. Thus if one encrypts all 2^{32} cosets of \mathcal{D}_i , we get all the 2^{32} cosets of \mathcal{M}_I , where each coset contains exactly 2^{32} ciphertexts. Only for completeness, if the final MixColumns operation is omitted, then the encryption of all 2^{32} cosets of \mathcal{D}_i results in all the 2^{32} cosets of $\mathcal{I}\mathcal{D}_I$, where each coset contains exactly 2^{32} ciphertexts.

Indeed, note that by Theorem 2 two elements that belong to the same coset of \mathcal{D}_I can not belong to the same coset of \mathcal{M}_J for $|I| + |J| \leq 4$. Thus, given a coset of \mathcal{D}_i with $|i| = 1$, after 4 rounds each element is distributed in a different coset of \mathcal{M}_J for $|J| = 3$. Note that $\mathcal{D}_i \oplus \mathcal{M}_j = \bigcup_{a \in \mathcal{M}_j} \mathcal{D}_i \oplus a$. Thus, since the coset of \mathcal{D}_i contains 2^{32} elements and since there are exactly 2^{32} cosets of \mathcal{M}_J , the elements of $\mathcal{D}_i \oplus \mathcal{M}_j$ are uniformly distributed in each coset of \mathcal{M}_I .

Backward Direction. If one decrypts $\mathcal{M}_j \oplus b$ for four rounds ($b \in \mathcal{D}_i$), then - due to Theorem 2 - the set $R^{(-4)}(\mathcal{M}_j \oplus b)$ is a set of 2^{32} plaintexts where each plaintext belongs to a different coset of a diagonal space \mathcal{D}_J of dimension 12. If one decrypts all 2^{32} cosets of \mathcal{M}_j , one gets all the 2^{32} cosets of \mathcal{D}_J , where each coset contains exactly 2^{32} plaintexts.

A Distinguisher with complexity 2^{64} : Uniform Distribution. Suppose to start in the middle with 2^{64} texts in the same coset of $\mathcal{D}_i \oplus \mathcal{M}_j$, and let

J and I fixed such that $|I| = |J| = 3$. As we have seen, the ciphertexts are uniform distributed in all the cosets of \mathcal{M}_I , that is each coset contains exactly 2^{32} ciphertexts. In the same way, the plaintexts are uniform distributed in all the cosets of \mathcal{D}_J , that is each coset contains exactly 2^{32} plaintexts. Thus, one needs only to count the number of elements in the ciphertexts and plaintexts that belongs to each coset to distinguish an 8-round AES permutation from a random one.

Description as Zero-Sum Distinguisher. An even simpler approach is possible: the simplest method is to XOR the 2^{64} plaintexts *and* ciphertexts and verify that the result is zero. The complexity is 2^{64} . This is the distinguisher which exploits the integral property proposed by Gilbert in [8].

Even if we've already presented it, we recall it using the subspace trail notation, which allows an easier explanation than using the *Super-SB* operation (9) introduced by Gilbert. To do this, we recall the 7-round AES distinguisher proposed by Knudsen and Rijmen in [13], which has a complexity of 2^{56} and which exploits the following integral property:

$$\text{Zero-Sum} \xleftarrow{R^{-3}} \mathcal{D}_0 \oplus \mathcal{C}_0 \oplus a \xrightarrow{R^4} \text{Zero-Sum},$$

where $a \in (\mathcal{D}_0 \oplus \mathcal{C}_0)^\perp$. Equivalently, this means that if one starts from a (collection of) coset(s) of \mathcal{D}_0 then after four rounds (without the final MixColumns operation) the integral property holds. In a similar way, if one starts from a (collection of) coset(s) of \mathcal{C}_0 , then the integral property holds three rounds before.

As shown in detail in [10] and in Sect. 3, for each $a \in \mathcal{C}_0^\perp$ there exists unique $b \in \mathcal{M}_0^\perp$ such that $R^{-1}(\mathcal{M}_0 \oplus b) = \mathcal{C}_0 \oplus a$. This means that if one starts from a coset of \mathcal{M}_0 , then the integral property holds four rounds before. Indeed, this coset of \mathcal{M}_0 is mapped into a coset of \mathcal{C}_0 , and then the integral property holds. Thus if one takes a collection of cosets of \mathcal{M}_0 , then the integral property holds four rounds before. In conclusion, if one starts in the middle with a coset of $\mathcal{D}_0 \oplus \mathcal{M}_0$ instead of a coset of $\mathcal{D}_0 \oplus \mathcal{C}_0$, then the integral property holds both after four rounds and four rounds before:

$$\text{Zero-Sum} \xleftarrow{R^{-4}} \mathcal{D}_0 \oplus \mathcal{M}_0 \oplus a' \xrightarrow{R^4} \text{Zero-Sum}$$

where $a' \in (\mathcal{D}_0 \oplus \mathcal{M}_0)^\perp$. The complexity is 2^{64} since $|\mathcal{D}_0 \oplus \mathcal{M}_0 \oplus a'| = 2^{64}$.

Finally, we show that the probability of success of a generic player is negligible, and we give an approximation of the computational cost of the generic player. First, note that given 2^{64} random plaintexts which are balanced and uniformly distributed, the probability that also the corresponding ciphertexts have these properties is upper bounded by 2^{-128} .

What is the cost of such a player? The authors of [13] *conjecture* that the complexity to find an N (plaintexts, ciphertexts) pairs for which the sum in k bits is equal to zero is $\mathcal{O}(N \cdot 2^{k/1+\log_2 N})$. Even this is an inaccurate estimation of the complexity we are looking for (due to the fact that the complexity estimate above is in the big \mathcal{O} notation), it gives an idea of the cost of the generic player

- approximately of $\mathcal{O}(2^{64} \cdot 2^{128/65}) \approx \mathcal{O}(2^{65.97})$ computations - with respect to the shortcut one - approximately 2^{64} computations.

E Details of Known-Key Distinguisher when the Computational Cost of the Generic Player is Considered

Referring to the known-key distinguisher scenario described in Sect. 4.1, the generic player depends by the oracle to generate the N -tuple (i.e. he cannot work alone to generate it). As a consequence, note that two possible settings can be analyzed. In the first one, only the number of oracle queries is considered to determine the computational cost of this player, that is the number of encryptions/decryptions required by the generic player to the oracle. In the second one, both the number of oracle queries and any further computational cost of the generic player (which is in general not negligible) are considered. As we have already said, we expect that a known-key distinguisher in the first setting works also in the second one but not viceversa. If the total cost of the generic player is well approximated by the number of queries (assuming 1 oracle-query \approx 1 computation/encryption), these two settings are completely equivalent.

In the main text, we have focused only on the first case, that is we have approximated the cost of the generic player by the number of oracle-queries necessary to generate a sufficient number of (plaintexts, ciphertexts) pairs such that n tuples with the required properties exist with a good probability. However, note that the player has also to identify the n tuples with the required properties before sending them to the verifier. As we are going to show, this computational cost is not negligible. In this section, we present the details of this case both for the 9-round distinguisher presented in Sect. 6 and for the 10-round one with key schedule presented in Sect. 7.2.

E.1 Known-Key Distinguisher on 9-Round AES

For a complete description of the 9-round known-key distinguisher for AES, we refer to Sect. 6. Here we limit to consider the cost of the generic player to find the n tuples with the required properties. In particular, we are going to show that if this cost is taken into account, then $n = 3$ tuples are sufficient for our distinguisher on 9 rounds of AES.

By formula (6), if $n = 3$ then $2^{52.2}$ different couples, or approximately $2^{26.6}$ plaintexts/ciphertexts pairs are sufficient to find the 3-tuples with the required properties (where the plaintexts belong to the same coset of \mathcal{D}_i). Indeed, note that with $2^{52.2}$ different couples it is possible to construct approximately 2^{154} different sets of 3 tuples. Since the probability that a set satisfies the required properties is 2^{-154} , there is at least one set that satisfies the property with non-negligible property. Thus, the cost to generate them is of $2^{26.6}$ oracle-queries.

Given these $2^{26.6}$ (plaintexts, ciphertexts) pairs, the generic player must work on the ciphertexts (note that the property on the plaintexts is already

satisfied) in order to find the 3-tuples with the required properties. For each couple $\{(p^1, c^1), (p^2, c^2)\}$, a possible strategy is to find the key k such that $R_k^{-1}(c^1) \oplus R_k^{-1}(c^2) \in \mathcal{M}_i$, using the attack of Sect. 5.1, and then to find 3 couples with a common key k . In the following, we present a modified strategy that allows to reduce the computational cost.

A possible way to reduce the total computational cost is to work first on only two couples (instead of three), that is to find two couples with the same key for which the required property is satisfied. Since there are $2^{52.2}$ couples, the player can construct approximately $2^{103.4}$ 2-tuples (i.e. different sets of two different couples). Approximately, there are $2^{103.4} \cdot (4 \cdot 2^{32})^2 \cdot 2^{-128} = 2^{43.4}$ different sets with on average one key in common for the two couples. For this step and using the attack of Sect. 5.1, the cost can be approximated at $2^{103.4} \cdot 4 \cdot 2^{11.6} = 2^{117}$ table look-ups. Then, given two couples with a common key k , the attacker looks for a third couple for which the required property is satisfied by the found key k . Note that for a given key, the probability that a pair of ciphertexts belong to the same coset of \mathcal{M}_l one round before for that key is only $2^{34-128} = 2^{-94}$. It follows that the player has to consider all the $2^{43.4}$ possible sets of two couples just found and all the possible $2^{52.4}$ couples (for a total of approximately $2^{43.4} \cdot 2^{52.2} = 2^{95.6}$ possibilities) in order to find the three tuples. Thus, given two couples with a common key, the idea is simply to test this found key on all the other couples, until one couple that satisfies the required property is found. To do this, the player computes $2^{43.4} \cdot 2^{52.2} \cdot 2 \cdot 2^4 = 2^{100.6}$ S-Box look-ups tables. It follows that the attacker is able to find the three desired couples, with a cost of approximately 2^{117} table look-ups or $2^{109.5}$ nine-round encryptions, besides the (non-negligible) memory cost to store the couples found at the first step with the corresponding key.

The cost of the shortcut player can instead be approximated by $3 \cdot 2^{44} = 2^{45.6}$ nine-round encryptions. Thus, $n = 3$ tuples are sufficient to set up the 9-round known-key distinguisher when all the costs (oracle queries + computational cost of generic player). Even if we don't exclude that the generic player can use better strategies to find these three couples, it seems improbable that the generic player is able to find the 3-tuples faster than the shortcut player when all the costs are considered. It follows that if the two players have to send 3 different tuples with the desired properties, then the game is win (with very high probability) by the player who knows the key. We leave as an open problem the research of a better strategy that the generic player can use to find these n tuples. For completeness, using the above strategy it is possible to prove that $n = 2$ tuples are not sufficient to set up this distinguisher¹⁸.

¹⁸ In this case, $2^{30.5}$ different couples instead of $2^{52.2}$ are sufficient, that is $2^{16.75}$ different (plaintexts, ciphertexts) pairs, for a cost of $2^{16.75}$ oracle-queries. In order to find the key, the total cost can be approximated at $2^{42.1}$ table look-ups, or $2^{34.6}$ nine-round encryptions. The cost for the player who knows the key is of 2^{45} nine-round encryptions. It follows that $n = 2$ tuples are not sufficient.

E.2 Known-Key Distinguisher on 10-Round AES with Key Schedule

For a complete description of the 10-round known-key distinguisher for AES with key schedule, we refer to Sect. 7.2. Here we limit to consider the cost of the generic player to find the n tuples with the required properties. In this case, it turns that $n = 2$ tuples are sufficient.

Indeed, in this case the cost for the shortcut player is of 2^{45} computations. Instead, if $n = 2$ and using formula (8), the other player must consider at least $s = 2^{52.5}$ different couples, that is approximately $2^{26.75}$ different (plaintext, ciphertext) pairs (with a cost of $2^{26.75}$ oracle-queries), where all the plaintexts belong to the same coset of \mathcal{D}_j . The player can construct approximately 2^{104} 2-tuples. First of all, for each one the player look for a final key k^{10} (if exists). Since the probability that such key exists for a given 2-tuple is only 2^{-60} , only 2^{44} 2-tuples survive this step. The cost of this step is well approximated by $2^{104} \cdot 4 \cdot 2^{11.6} = 2^{117.6}$ table look-ups, using the attack described in Sect. 5.1.

Given these 2^{44} 2-tuples just found with the corresponding key k^{10} , for each key k^{10} the player can simply find the j -th column of the first key k^0 (note that three columns of k^0 can take any possible values), and checks if the property on the plaintexts is satisfied. Since this happens with probability¹⁹ 2^{-44} , such key usually exists. The cost of this step is well approximated by $2^{44} \cdot 2^{34} \cdot 40 = 2^{83.1}$ S-Box look-ups to check the key schedule, that is approximately $2^{75.5}$ ten-round encryptions, besides the (not-negligible) memory cost. It follows that the computational cost for this player is much higher than the one of the shortcut player. However, since we don't exclude that the generic player can use a better strategy to win the game, we leave the open problem to improve the strategy that we have just presented here. On the other hand, even if a better strategy is found, it seems improbable that the generic player is able to find the 2-tuples faster than the shortcut player when all the costs are considered.

F New 7-, 8- and 9-round Known-Key Distinguishers for AES

In this section, we propose new 8- and 9-round known-key distinguisher for AES, which are obtained extending at the end or/and at the beginning a 7-round known-key distinguisher for AES. The strategy to set up them is the same used in Sect. 6 and 7. For this reason, we refer to those sections for all the details. We highlight that the *9-round known-key distinguisher proposed in this section is the best one* both for the computational and data cost among those currently present in the literature.

¹⁹ Remember that we're working with plaintexts in the same coset of \mathcal{D}_j . After one round, they are mapped into the same coset of \mathcal{C}_j . Thus, two texts belong to the same coset of $\mathcal{C}_j \cap \mathcal{D}_i$ for a certain i , if three bytes of the j -th column are equal to zero. This happens with probability $4 \cdot 2^{-24} = 2^{-22}$.




Fig. 6. 7-round differential characteristic for known-key distinguisher of AES-128.

F.1 7-Round Known-Key Distinguisher

For the following, we briefly recall the currently best known distinguisher on 8 rounds of AES (proposed in [11] and already presented in Sect. 4.3). This distinguisher is obtained starting from the 8-round distinguisher presented in [9], and depicted in Fig. 4. Using the subspace trail notation and the known-key distinguisher scenario, the goal of the two players in this distinguisher is to find a pair of (plaintexts, ciphertexts) - i.e. (p^1, c^1) and (p^2, c^2) - with the following properties: the two plaintexts belong to the same coset of \mathcal{D}_i - i.e. $p^1 \oplus p^2 \in \mathcal{D}_i$ - and the two ciphertexts belong to the same coset of \mathcal{M}_i - i.e. $c^1 \oplus c^2 \in \mathcal{M}_i$, where the index i is fixed. The idea proposed in [11] to improve this distinguisher is simply to not fix the initial subspace \mathcal{D}_i and the final one \mathcal{M}_j , that is to leave i and j completely arbitrary (i.e. they can take any possible values). It follows that the probability that a solution of the inbound phase of the rebound attack satisfies the outbound phase is higher, which implies that a complexity of 2^{44} is sufficient (instead of 2^{48}) for the shortcut player.

The same strategy can be applied to the 7 rounds distinguisher presented in [15] and recalled in Sect. 4.2. In particular, using the same argumentation of [11], the computational cost of the distinguisher illustrated in Fig. 6 is 2^{20} instead of 2^{24} . Indeed, note that for free \mathcal{D}_i and \mathcal{M}_j , the probability that a solution of the inbound phase satisfies the outbound phase increases of a factor $4^2 = 2^4$.

F.2 8-Round Known-Key Distinguisher

A possible 8-round known-key distinguisher can be set up starting from the 7-round distinguisher just presented and extending it at the end (or at the beginning) using a similar technique presented in Sect. 6 for the 9-round distinguisher. We refer to Sect. 6 for a complete discussion of this technique and we limit here to give a formal definition of the distinguisher and to do some considerations about the data and the computational cost.

In the known-key distinguisher scenario, the two players have to send to the verifier n different tuples of (plaintext, ciphertext) pairs, that is $\{(p_i^1, c_i^1), (p_i^2, c_i^2)\}$ for $i = 0, \dots, n - 1$, with the following properties:

1. for each tuple, there exists j s.t. the two plaintexts belong to the same coset of \mathcal{D}_j , that is

$$\forall i = 0, \dots, n-1, \quad \exists j \in \{0, \dots, 3\} \quad \text{s.t.} \quad p_i^1 \oplus p_i^2 \in \mathcal{D}_j;$$

2. there exists a key k s.t. for each tuple there exists l for which the two ciphertexts belong to the same coset of \mathcal{M}_l one round before, that is

$$\exists! k \quad \text{s.t.} \quad \forall i = 0, \dots, n-1, \quad \exists l \in \{0, \dots, 3\} \quad \text{s.t.} \quad R_k^{-1}(c_i^1) \oplus R_k^{-1}(c_i^2) \in \mathcal{M}_l.$$

If only the number of oracle-queries is considered, it is possible to prove that $n \geq 3$ tuples are sufficient to set up this distinguisher. Indeed, using the same argumentation of Sect. 6, the generic player has to consider approximately $2^{52.18}$ different couples (see (6)), that is approximately $2^{26.59}$ different (plaintexts, ciphertexts) pairs, for a cost of $2^{26.6}$ oracle-queries, in order to have good probability to construct 3 tuples with the required properties. On the other hand, the cost for the shortcut player is only of $3 \cdot 2^{20} = 2^{21.6}$ computations. In order to make the advantage of the shortcut player more significant, we choose an (arbitrary) value of $n = 8$, which implies a cost for the shortcut player of 2^{23} computations and of $2^{48.9}$ computations for the generic player.

In a similar way, it is possible to prove that $n = 2$ tuples are sufficient to set up this 8-round distinguisher when all the costs (number of oracle-queries + cost of generic player) are taken into account. Indeed, in this case and in order to construct 2 tuples that satisfy the required property for the same key, the second player has to consider approximately $2^{30.4}$ different tuples, that is $2^{15.74}$ different (plaintexts, ciphertexts) pairs. Using the same analysis proposed in App. E.1, the cost to find the 2-tuples that satisfy the relation \mathcal{R} can be approximated at $2^{30.4} \cdot 2^{11.6} = 2^{42}$ table look-ups, that is $2^{34.68}$ eight-round AES encryptions. On the other hand, the cost of the player that knows the key is of $2 \cdot 2^{20} = 2^{21}$ computations, which is (much) lower than $2^{35.68}$.

For both cases, the verifier uses the same strategy presented in Sect. 6, and her cost is well approximated by $2^{11.6}$ eight-round encryptions.

F.3 9-Round Known-Key Distinguisher

An efficient 9-round known-key distinguisher can be set up by extending the previous 8-round distinguisher at the beginning, or equivalent by extending the 7-round known-key distinguisher both at the beginning and at the end. Such a distinguisher is the best one both for the computational and data cost among those presented in literature.

In order to set up this distinguisher on 9 rounds, we exploit the same strategy proposed for 10-round one. For this reason, we refer to Sect. 7 for a complete discussion. In the known-key distinguisher scenario, the two players have to send to the verifier n different tuples of (plaintext, ciphertext) pairs, that is $\{(p_i^1, c_i^1), (p_i^2, c_i^2)\}$ for $i = 0, \dots, n-1$, with the following properties:

1. there exists a key k^0 s.t. for each tuple there exists j for which the two plaintexts belong to the same coset of \mathcal{D}_j after one round, that is

$$\exists! k^0 \quad \text{s.t.} \quad \forall i = 0, \dots, n-1, \quad \exists j \in \{0, \dots, 3\} \quad \text{s.t.} \quad R_{k^0}(p_i^1) \oplus R_{k^0}(p_i^2) \in \mathcal{D}_j;$$

2. there exists a key k^9 s.t. for each tuple there exists l for which the two ciphertexts belong to the same coset of \mathcal{M}_l one round before, that is

$$\exists! k^9 \quad \text{s.t.} \quad \forall i = 0, \dots, n-1, \quad \exists l \in \{0, \dots, 3\} \quad \text{s.t.} \quad R_{k^9}^{-1}(c_i^1) \oplus R_{k^9}^{-1}(c_i^2) \in \mathcal{M}_l.$$

We discuss here the two scenarios in which (1) the subkeys are independent and in which (2) a key schedule holds. Since the strategies used by the players are equivalent of the ones proposed in Sect. 7, we refer to that section for all the details and we limit here to do some considerations about the computational and data cost.

Independent Subkeys: No Key Schedule. As for the 10-round known-key distinguisher, we first consider the case in which there is no key schedule. The idea for the generic player is to choose an initial key \hat{k} and to choose the plaintexts in the set $D_a = R_{\hat{k}}^{-1}(\mathcal{D}_i \oplus a)$. If this player needs more plaintexts, the idea is to compute other $D_{a'}$ sets for another $a' \in \mathcal{D}_i^\perp$ using the same key \hat{k} , as for the 10-round distinguisher case.

In a similar way as before, $n \geq 3$ tuples are sufficient for the case in which only the number of oracle-queries is considered, while $n \geq 2$ tuples are sufficient to set up the distinguisher for the case in which all the cost are considered. As before, we choose an (arbitrary) value of $n = 8$ to make the advantage of the shortcut player more significant. The costs of the two players are well approximated by the costs given for the previous 8-round distinguisher (note that the cost to compute $R_{\hat{k}}^{-1}(\mathcal{D}_i \oplus a)$ is negligible compared to the total cost). Finally, the cost of the verifier is double with respect the previous case (since she has to check the existence of two keys).

The Key Schedule Case. Similar to what done in Sect. 7, the idea for the generic player is to choose the plaintexts in the same coset of \mathcal{D}_i in order to maximize the possible number of key k^0 and k^{10} for which the required properties are satisfied - we refer to Sect. 7 - Prop. 1 for all the details. By analogous calculation of Sect. 7, if $n = 2$ then this player needs approximately $2^{52.48}$ different couples in order to have a good probability of success, that is he must do approximately $2^{26.74}$ oracle-queries. On the other hand, the computational cost for the shortcut player is of 2^{21} . Thus, $n = 2$ tuple is sufficient for this setting. Since n must be at least equal to 2, note that the same result holds also for the case in which we consider the total computational cost (oracle cost + player cost). The verification cost is (approximately) equivalent to the one given for the case of independent subkeys, due to the (same) argumentations given in Sect. 7.

F.4 Consideration and Comparison with Gilbert’s Distinguisher

Using the technique described in Sect. 8, one can theoretically extend again the previous 9-round known-key distinguisher at the end or/and at the beginning, obtaining a 10- or/and 11-round known-key distinguisher. Even if this is possible, we show that this distinguisher is (much) less competitive than the one described in Sect. 7 - obtained by a single extension at the end and at the beginning of a 8-round distinguisher - and the one proposed by Gilbert in [8]. The main problem of a 10-round distinguisher obtained extending a 7-round distinguisher two times at the end and one at the beginning (or viceversa) regards the computational cost of the verifier. Indeed, it is possible to show that this cost is much higher than 2^{64} , using similar argumentations proposed in Sect. 8 and due to the complexity of the key-recovery attack described in Sect. 5.2. Thus, since by definition the computational cost of verifier must be smaller than the costs of the two players, it follows that the overall computational cost of such a distinguisher is much higher than 2^{64} , that is it is much higher than the computational cost of our 10-round known-key distinguisher proposed in Sect. 7 and the one proposed by Gilbert in [8].

A final observation regards the possibility to set up a 8- and 9-round known-key distinguisher by extending at the beginning or/and at the end the 7-round known-key distinguisher proposed by Knudsen and Rijmen in [13] and based on the balance property (in a similar way of what Gilbert did to set up his distinguisher on 10-round). Note that the 7-round distinguisher based on the balance property [13] has a complexity of 2^{56} . Thus, the 8-round known-key distinguisher obtained by extending the 7-round distinguisher at the end (or at the beginning) has at least a complexity of 2^{56} , which is higher than the 8-round known-key distinguisher proposed in [11] (complexity of 2^{44}) and our one proposed in App. F.2 (complexity of 2^{23}). Similar argumentation holds for the 9-round distinguisher (obtained by extending at the beginning and at the end the cited 7-round distinguisher). It follows that such distinguishers are not competitive with respect to the others currently present in literature.

G Proof of Proposition 1 - Sect. 7.2

Proposition 2. *Let p^1 and p^2 two plaintexts that belong to the same coset of \mathcal{D}_j for a certain j , that is $p^1 \oplus p^2 \in \mathcal{D}_j$. Moreover, assume that $p^1 \oplus p^2 \notin \mathcal{D}_j \cap \mathcal{C}_L$ for each $L \subseteq \{0, 1, 2, 3\}$ with $|L| \leq 3$. Then there exist on average 2^{106} different keys k such that $R_k(p^1) \oplus R_k(p^2) \in \mathcal{D}_l$ for a certain $l \in \{0, 1, 2, 3\}$.*

Proof. First of all, suppose by contradiction that $p^1 \oplus p^2 \in \mathcal{D}_j \cap \mathcal{C}_L$ for a certain $L \subseteq \{0, 1, 2, 3\}$ with $|L| \leq 3$. Since $\mathcal{D}_j \cap \mathcal{C}_L \subseteq \mathcal{C}_L$, it follows that $R(p^1) \oplus R(p^2) \in \mathcal{C}_j \cap \mathcal{M}_L \subseteq \mathcal{M}_L$ for $|L| \leq 3$. By (3), it follows that if $R(p^1) \oplus R(p^2) \in \mathcal{M}_L$ for $|L| \leq 3$, then $R(p^1) \oplus R(p^2) \notin \mathcal{D}_j$ for $|j| = 1$.

Now, without loss of generality assume $j = 0$ (the proof can be easily generalized for each j). The idea is to look for the number of keys such that $R_k(p^1) \oplus R_k(p^2) \in \mathcal{C}_j \cap \mathcal{D}_l \subseteq \mathcal{D}_l$ for a certain $l \in \{0, 1, 2, 3\}$.

By definition, $p_{i,l}^1 = p_{i,l}^2$ for each $i \neq l$. Thus, it is easy to note that for each value of $k_{i,l}$ for $i \neq l$ (that is 12 bytes) then $R(p^1)_{i,l} = R(p^2)_{i,l}$ for each i and for each $l = 1, 2, 3$ (i.e. the second, the third and the fourth columns of $R(p^1)$ and $R(p^2)$ are equal), for a total of 2^{96} possibilities.

Consider now the bytes on the first diagonal, that is in positions $i = l$. In this case, one has to guarantee that after one round three bytes of the two texts are equal, in order to have $R_k(p^1) \oplus R_k(p^2) \in \mathcal{C}_j \cap \mathcal{D}_l$. As shown in the attack on 3 rounds of Sect. 5.1, for each l on average there are 2^8 possible combinations of these four bytes such that this condition is satisfied. Since there are four different possible values of l , the number of possible keys for this second point are on average $4 \cdot 2^8 = 2^{10}$.

In conclusion, the number of keys such that $R_k(p^1) \oplus R_k(p^2) \in \mathcal{C}_j \cap \mathcal{D}_l$ for a certain l are $(2^8)^{12} \cdot 2^{10} = 2^{106}$.

□

Let p^1 and p^2 in the same coset of \mathcal{D}_j (that is $p^1 \oplus p^2 \in \mathcal{D}_j$), and without loss of generality assume $j = \{0\}$. We do a consideration about the hypothesis that $p^1 \oplus p^2 \notin \mathcal{D}_0 \cap \mathcal{C}_L$ for each $L \subseteq \{0, 1, 2, 3\}$ with $|L| \leq 3$. If $p^1 \oplus p^2 \in \mathcal{D}_0$, then $p^1 \oplus p^2 \notin \mathcal{D}_0 \cap \mathcal{C}_L$ for each $L \subseteq \{0, 1, 2, 3\}$ with $|L| \leq 3$ if and only if $p_{i,i}^1 \neq p_{i,i}^2$ for each $i = 0, \dots, 3$. By simple calculation, this happens with probability $(255/256)^4 \simeq 2^{-0.0225}$.

Thus, given a coset of \mathcal{D}_0 , it is possible to construct approximately $2^{31} \cdot (2^{32} - 1) = 2^{62.9999\dots} \simeq 2^{63}$ different couples. If one eliminates all the pairs (p^1, p^2) for which there exists at least one i such that $p_{i,i}^1 = p_{i,i}^2$, then the number of survived pairs is $2^{62.9999\dots} \cdot (255/256)^4 = 2^{62.9775\dots}$, which is still well approximated by 2^{63} for our scope. The cases for the other subspaces \mathcal{D}_j are similar.

H Gilbert's Known-Key Distinguisher on 12 Rounds

In this section, using a technique similar to the one presented in Sect. 8, we show how to extend the Gilbert's 10-round known-key distinguisher both at the end and at the beginning in order to set up (theoretically) a 12-round known-key distinguisher based on the balance property. Before we do this, we first do some considerations about the computational cost of the verification step for the Gilbert's 10-round known-key distinguisher model.

H.1 Considerations about the Computational Cost of the Verifier - Gilbert's 10-round Distinguisher

First of all, we recall a formal definition of the Gilbert's distinguisher on 10 rounds of AES, already given in Sect. 4.4. In the known-key distinguisher scenario and in order to win the game, the players have to send to the verifier 2^{64} (plaintext, ciphertext) pairs, that is (p^i, c^i) for $i = 0, \dots, 2^{64} - 1$, with the following properties:

1. there exists a key k^0 such that the texts $\{R_{k^0}(p^i)\}_i$ are uniform distributed among the cosets of \mathcal{D}_I with $|I| = 3$ fixed, or equivalently s.t. the sum of the plaintexts after one round is equal to zero, that is $\bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p^i) = 0$;
2. there exists a key k^{10} such that the texts $\{R_{k^{10}}^{-1}(c^i)\}_i$ are uniform distributed among the cosets of \mathcal{M}_J with $|J| = 3$ fixed, or equivalently s.t. the sum of the ciphertexts one round before is equal to zero, that is $\bigoplus_{i=0}^{2^{64}-1} R_{k^{10}}^{-1}(c^i) = 0$.

As for our distinguisher, we assume that all the subkeys are independent, that is no key schedule holds. Note that the same distinguisher works also in the case in which the subkeys are not independent.

When the verifier receives the set of 2^{64} (plaintext, ciphertext) pairs from one of the two players, she check if the properties are satisfied or not by finding the two keys. Note that since no key schedule holds, the verifier can work independently on k^0 and k^{10} . As we have already showed in App. A, the verifier can work on single bytes of k^0 and k^{10} (i.e. independently of the others) in order to check the existence of these keys. By simple computation, the verification cost can be approximated by $2 \cdot 16 \cdot 2^8 \cdot 2^{64} = 2^{77}$ S-Box look-ups, or equivalently $2^{69.36}$ ten-rounds encryptions of AES. Note that even using an improved strategy with respect to the one proposed in [8], the computational cost of the verifier is still higher than the cost of the shortcut player, which is approximately of 2^{64} ten-rounds encryptions of AES. Thus, the claim “*the overall complexity of checking \mathcal{R} is strictly smaller than $N = 2^{64}$ AES $_{10}^*$ operations*” made in [8] (see Sect. 4.2 - page 218) is false, that is the computational cost of the verifier is not negligible with respect to the cost of the players. We refer to App. A for details.

Before we go on, we stress the importance of this fact. If the verifier guesses 4 bytes of the key at the same time, then there is no way to extend this distinguisher again at the beginning and at the end. This justifies the claim made in [8] that it is not possible to extend again Gilbert’s 10-round known-key distinguisher in order to set up a 12-round one. Indeed, a further extension requires to guess the entire initial and final subkey, which implies a brute force attack for the verifier. Instead, the possibility to work on single bytes of k^0 and k^{10} independently of the other allows to set up a 12-round known-key distinguisher, as shown in the following. To do this, the idea is simply to use the square attack on 5-round AES with the extension at the end for the verification process.

Before we present how to extend this 10-round distinguisher both at the end and at the beginning, we show a possible way to modify the Gilbert’s 10-round distinguisher in order to overcome the problem regarding the cost of the verification step. In the known-key distinguisher setting, the players have to send to the verifier a suitable number n of sets of 2^{64} (plaintext, ciphertext) pairs, i.e. (p_i^j, c_i^j) for $i = 0, \dots, 2^{64} - 1$ and $j = 0, \dots, n - 1$, where $n > 2^6$ - note that $2^{69.4} \cdot 2^{-64} \simeq 2^{5.4}$, where $2^{69.4}$ is the verification cost - such that:

1. there exists a keys k^0 such that for all $j = 0, \dots, n - 1$ there exists $I \subseteq \{0, 1, 2, 3\}$ with $|I| = 3$ fixed such that the texts $\{R_{k^0}(p_i^j)\}_i$ are uniform distributed among the cosets of \mathcal{D}_I , or equivalently such that the sum of the

plaintexts after one round is equal to zero:

$$\forall j = 0, \dots, n-1 : \quad \bigoplus_{i=0}^{2^{64}-1} R_{k^0}(p_i^j) = 0;$$

2. there exists a key k^{10} such that for all $j = 0, \dots, n-1$ there exists $J \subseteq \{0, 1, 2, 3\}$ with $|J| = 3$ fixed such that the texts $\{R_{k^{10}}^{-1}(c_i^j)\}_i$ are uniform distributed among the cosets of \mathcal{M}_J , or equivalently such that the sum of the ciphertexts one round before is equal to zero:

$$\forall j = 0, \dots, n-1 : \quad \bigoplus_{i=1}^{2^{64}} R_{k^{10}}^{-1}(c_i^j) = 0.$$

Note that as for our known-key distinguishers on 10- and 12-round of AES, it is very important that the keys k^0 and k^{10} are equals for all the sets of (plaintext, ciphertext) pairs, i.e. for each $j = 0, \dots, n-1$.

H.2 Known-Key Distinguisher on 12 Rounds Based on the Balance Property

A formal definition of the Gilbert's distinguisher on 12 rounds of AES can be the following. As before, assume there are two players - one knows the key and the other not, and the verifier. In the known-key distinguisher scenario, the players have to send to the verifier 2^{64} (plaintext, ciphertext) pairs, that is (p^i, c^i) for $i = 0, \dots, 2^{64} - 1$, with the following properties:

1. there exist keys k^0, k^1 for which there exists $I \subseteq \{0, 1, 2, 3\}$ with $|I| = 3$ fixed such that the texts $\{R_{k^1}(R_{k^0}(p^i))\}_i$ are uniform distributed among the cosets of \mathcal{D}_I , or equivalently such that the sum of the plaintexts after two rounds is equal to zero:

$$\bigoplus_{i=0}^{2^{64}-1} R_{k^1}(R_{k^0}(p^i)) = 0;$$

2. there exist keys k^{11}, k^{12} for which there exists $J \subseteq \{0, 1, 2, 3\}$ with $|J| = 3$ fixed such that the texts $\{R_{k^{11}}^{-1}(R_{k^{12}}^{-1}(c^i))\}_i$ are uniform distributed among the cosets of \mathcal{M}_J , or equivalently such that the sum of the ciphertexts two rounds before is equal to zero:

$$\bigoplus_{i=0}^{2^{64}-1} R_{k^{11}}^{-1}(R_{k^{12}}^{-1}(c^i)) = 0.$$

As for our distinguishers, we assume that all the subkeys are independent, that is no key schedule holds. However, such distinguisher works exactly in the same way also in the case in which a key schedule holds.

The two major problems that arise for this setting regard (1) the computational cost of the verifier and (2) the probability of success of a generic player (with respect to the shortcut player). We start to analyze the first one. For the 10 rounds distinguisher, the verifier can simply work on each byte of the keys independently of the others. In other words, the verifier simply performs a classical square attack on the ciphertexts and on the plaintexts in order to find (if exist) the initial and the final key, working independently on each byte. For the 12-round distinguisher, the idea is to perform a square attack with the extension at the end (see [4] for details). For example, working on the ciphertexts, the attacker guesses one diagonal of the final key k^{12} , decrypts that diagonal and then finds (independently) four bytes of k^{11} , using the classical integral attack. Repeating four times this attack, she is able to check the existence of k^{11} and k^{12} . The total cost of this step can be approximated by $2 \cdot 4 \cdot 2^{32} \cdot (16 \cdot 2^8) \cdot 2^{64} \cdot 32 = 2^{116}$ S-Box look-ups, since the attack must be performed on the plaintexts and on the ciphertexts (i.e. 2), and for each guess key (i.e. $4 \cdot 2^{32}$ (possibilities for k^{12}) $\cdot (4 \cdot 2^8)$ (possibilities for k^{11})) the attacker must compute $2^{64} \cdot 32$ S-Box (i.e. 32 for each text). Using the approximation of 20 S-Box look-ups \approx 1 round of AES, this corresponds to do approximately 2^{106} twelve rounds encryptions. Since the cost of the shortcut player is well approximated by 2^{64} computations (using the same procedure proposed for the 8-round known-key distinguisher based on the integral property), the requirement that the cost of the verifier is much smaller than the cost of the players is not satisfied.

The second and probably major problem is to prove (or to give a strong argumentation) that the probability of the generic player to win the game is (much) smaller than the one of the other player, or equivalently that the time to create the set of 2^{64} (plaintexts, ciphertexts) is higher for a generic player than for the shortcut player (for the same probability of success).

Open Problem. A possible way to fix both the problems is to send not a single set of 2^{64} (plaintexts, ciphertexts) but more sets, such that the zero-sum property is always satisfied for the same subkeys k^0, k^1, k^{11} and k^{12} . In particular, since the cost of the verifier much be lower than the costs of the players, the number of sets n should be greater than $n \geq 2^{106} \cdot 2^{-64} = 2^{42}$, for a total cost which is higher than our distinguisher. This implies that such a distinguisher can not be better than our one proposed in Sect. 8.

As first thing, we give a formal definition of this modified 12-round known-key distinguisher, obtained extending the Gilbert one. In the known-key distinguisher scenario, the players have to send to the verifier n sets of 2^{64} (plaintext, ciphertext) pairs, that is (p_i^j, c_i^j) for $i = 0, \dots, 2^{64} - 1$ and $j = 0, \dots, n - 1$, where $n \geq 1$, with the following properties:

1. there exist keys k^0, k^1 for which for all $j = 0, \dots, n - 1$ there exists $I \subseteq \{0, 1, 2, 3\}$ with $|I| = 3$ fixed such that the texts $\{R_{k^1}(R_{k^0}(p_i^j))\}_i$ are uniform distributed among the cosets of \mathcal{D}_I , or equivalently such that the sum of the

plaintexts after two rounds is equal to zero:

$$\forall j = 0, \dots, n - 1 : \bigoplus_{i=0}^{2^{64}-1} R_{k^1}(R_{k^0}(p_i^j)) = 0;$$

2. there exist keys k^{11}, k^{12} for which for all $j = 0, \dots, n - 1$ there exists $J \subseteq \{0, 1, 2, 3\}$ with $|J| = 3$ fixed such that the texts $\{R_{k^{11}}^{-1}(R_{k^{12}}^{-1}(c_i^j))\}_i$ are uniform distributed among the cosets of \mathcal{M}_J , or equivalently such that the sum of the ciphertexts two rounds before is equal to zero:

$$\forall j = 0, \dots, n - 1 : \bigoplus_{i=1}^{2^{64}} R_{k^{11}}^{-1}(R_{k^{12}}^{-1}(c_i^j)) = 0.$$

We leave as open problems (1) to research a suitable value of n for which this distinguisher works, and (2) to prove that such a distinguisher is meaningful, that is to prove that the probability of success of a generic player is lower than the one of the shortcut player for that value of n .

I The Herds Attack

N. Ferguson *et al.* [6] presented the first (and unique) integral attack on 7 rounds of AES-128. The attack is obtained by extending at the beginning the integral attack on 6 rounds of AES-128 [4]-[12]. This attack requires $2^{128} - 2^{119}$ chosen plaintexts, which are distributed in $2^{96} - 2^{87}$ different cosets of \mathcal{D}_i for a certain $i \subseteq \{0, 1, 2, 3\}$ with $|i| = 1$, it has a computational complexity of 2^{120} seven-round AES-encryptions and a memory cost of 2^{64} bits of memory.

Here we show why this attack can not be used to set up a 14-round known-key distinguisher for AES, based on the balanced property. Moreover, the same argumentation can also be used to justify why key-recovery attacks with both an extension at the end and at the beginning can not be used to set up known-key distinguisher in the Gilbert model.

We briefly recall the idea used in [6] to set up the attack on 7 rounds of AES. As we have already seen, given 2^{32} plaintexts in the same coset of a diagonal space \mathcal{D}_i , their sum after 4 rounds is equal to zero for each key. Thus, it is possible to set up an integral attack on 5 rounds of AES (working independently on each byte of the final subkey), and on 6 rounds. In this second case, assuming that the final MixColumns is omitted (otherwise the idea is simply to exchange the final MixColumns operation and the final AddRoundKey operation - they are linear), the idea is to guess one column of $SR(k)$ - where k is the final round, decrypt one round and repeat the attack on 5 rounds. We refer to [4] for more details.

In order to attack 7 rounds, a first possibility is to extend the previous attack at the end, by guessing the entire final subkey, decrypting one round and repeating the attack on 6 rounds. However, while for AES-192 and AES-256 this attack is better than a brute force one, this is not true for AES-128. The idea

of Ferguson *et al.* is the following. Consider the entire codebook, that is 2^{128} plaintexts (which can be seen as the union of 2^{96} different cosets of a diagonal space \mathcal{D}_i). Their sum after 5 rounds is equal to zero for each key. Similar to the previous attack on 6 rounds, the idea is to guess 4 byte of the final key (i.e. one column of $SR(k)$), decrypt one round and do a classical integral attack (as the one already described for 5-round). However, as the authors observe, the sum is zero also for wrong keys and not only for the right one, since the full codebook is used.

To solve the problem, the idea is not to consider the sum of all the ciphertexts, but only of part of them. In particular, given a byte $x \in \mathbb{F}_{2^8}$ fixed, one guesses four byte of the initial key and computes one-round encryption of each plaintext. The idea is to select the plaintexts such that after one round the byte in the j -th row and l -th column of the corresponding text is equal to x , for a total of 2^{120} plaintexts. Then, the idea is to consider only the sum of the ciphertexts of these 2^{120} plaintexts, which is equal to zero only for the right key. Thus, by guessing 4 bytes of the final key and working independently on each byte of the second to last key, one checks if the sum is equal to zero.

Finally the authors show how to improve this technique in order to use only $2^{128} - 2^{119}$ chosen plaintexts instead of the full codebook, besides other improvements on the computational cost. We refer to [6] for a complete description of the attack, and we limit ourselves to explain why it can not be used for a known-key distinguisher.

Given the details of the herds attack, we now focus on the 14-round known-key distinguisher. Suppose by contradiction that such distinguisher can be set up. Without being too formal, we first give a more precise idea of this distinguisher. In the known-key distinguisher scenario, the players have to send to the verifier 2^n (plaintext, ciphertext) pairs, that is (p^i, c^i) for $i = 0, \dots, 2^n - 1$ where $n \geq 64$, with the following properties:

1. there exist keys k^0, k^1, k^2 and $I \subseteq \{0, 1, 2, 3\}$ with $|I| = 3$ fixed such that the texts $\{R_{k^2}(R_{k^1}(R_{k^0}(p^i)))\}_i$ are uniform distributed among the cosets of \mathcal{D}_I , or equivalently such that the sum of the plaintexts after three rounds is equal to zero:

$$\bigoplus_{i=0}^{2^n-1} R_{k^2}(R_{k^1}(R_{k^0}(p^i))) = 0;$$

2. there exist keys k^{12}, k^{13}, k^{14} and $J \subseteq \{0, 1, 2, 3\}$ with $|J| = 3$ fixed such that the texts $\{R_{k^{12}}^{-1}(R_{k^{13}}^{-1}(R_{k^{14}}^{-1}(c^i)))\}_i$ are uniform distributed among the cosets of \mathcal{M}_J , or equivalently such that the sum of the ciphertexts three rounds before is equal to zero:

$$\bigoplus_{i=1}^{2^{64}} R_{k^{12}}^{-1}(R_{k^{13}}^{-1}(R_{k^{14}}^{-1}(c^i))) = 0.$$

First of all, note 2^n must be greater or equal than $2^{128} - 2^{119}$, that is $2^n \geq 2^{128} - 2^{119}$. Indeed, this is the minimum number of texts for which the

attacker is able to use the herds attack (i.e. the computational cost is lower than a brute force one). Assuming that this distinguisher is meaningful, we analyze the strategies of the players and of the verifier.

First consider the shortcut player, i.e. the player who knows the key. As for the 8-round known-key distinguisher based on the integral property (see App. D for details), in order to guarantee the balanced property holds both on the plaintexts and on the ciphertexts, the best strategy for this player is to consider the union of at least $2^{64} - 2^{55}$ different cosets of $\mathcal{D}_i \cap \mathcal{M}_j$, that is

$$\bigcup_{k=0}^{2^{64}-2^{55}-1} (\mathcal{D}_i \oplus \mathcal{M}_j \oplus a_k)$$

for $2^{64} - 2^{55}$ different $a_k \in (\mathcal{D}_i \oplus \mathcal{M}_j)^\perp$ (where $|(\mathcal{D}_i \oplus \mathcal{M}_j)^\perp| = 2^{64}$). Note that the previous union of cosets can be rewritten in the following way:

$$\bigcup_{k=0}^{2^{64}-2^{55}-1} (\mathcal{D}_i \oplus \mathcal{M}_j \oplus a_k) = \bigcup_{k=0}^{2^{64}-2^{55}-1} \left(\bigcup_{b \in \mathcal{M}_j} \mathcal{D}_i \oplus (b \oplus a_k) \right) = \bigcup_{k=0}^{2^{96}-2^{87}-1} \mathcal{D}_i \oplus \hat{a}_k,$$

that is the player is considering in the middle a union of cosets of \mathcal{D}_i (analogous for \mathcal{M}_j). In particular, the computational cost of this player is at least $2^{128} - 2^{119}$.

Consider now the verification strategy. The idea is to use the herds attack to find the subkeys and so to prove their existences. However, as we are going to show, a problem arises for the verifier, since when she receives the plaintexts and the ciphertexts by the player, the only way in which she can check the existence of the six subkeys is using a brute force attack. In other words, the verifier can not use in any way the herds attack presented before. Indeed, to do this, she has to know the “intermediate” texts, which corresponds to the 7-round encryption of the plaintexts or to the 7-round decryption of the ciphertexts. Since she doesn’t know them, she can not divide the texts in set of 2^{120} elements. In other words, the extension at the beginning of the herds attack creates a problem for the verifier, since she doesn’t have any access to the intermediate values of the plaintexts/ciphertexts. This justifies why this attack can not be used in order to set up a 14-round known-key distinguisher, and more generally why an attack with both the extension at the end and the beginning can not be used to set up a known-key distinguisher in the Gilbert model.