Improved Related-Key Attacks on AES-192

Florent Mazelet, María Naya-Plasencia

Inria, France

SAC, 2026-08-27

Outline

This section

The AES [DR02]

- Widely used symmetric cryptography NIST standard.
- Many cryptanalysis papers, but only broken in the Related-Key setting for the 192-bit and 256-bit versions.
- We are still interrested in improving the existing attacks.

The AES

- 10,12,14 rounds of a fonction :

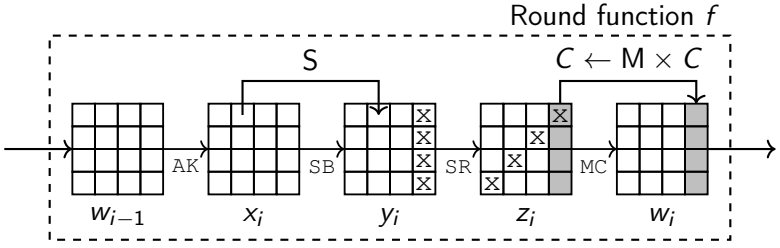

An AES round ([Jea16])

AES key schedule

Three key lengths : 128, 192, 256 bits.

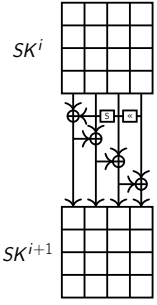

(a) AES-128

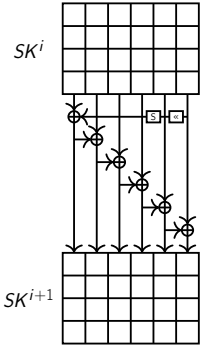

(b) AES-192

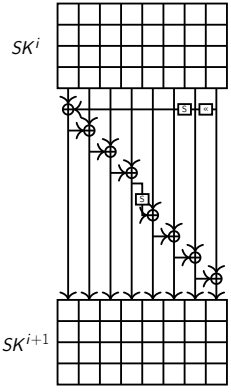

(c) AES-256

The key schedules of the AES

Differential cryptanalysis [BS91]

- Instead of looking at single messages, we follow the propagation of the difference $x \oplus x'$ of a *pair* of messages through the cipher.

Differential cryptanalysis [BS91]

- Instead of looking at single messages, we follow the propagation of the difference $x \oplus x'$ of a *pair* of messages through the cipher.
- A *differential characteristic* $\alpha \xrightarrow{p} \beta$ predicts that an input difference α becomes an output difference β with probability p.

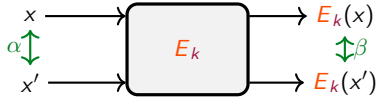

$$\Pr[\alpha \to \beta] = p \gg 2^{-n}$$

Differential cryptanalysis [BS91]

- Instead of looking at single messages, we follow the propagation of the difference $x \oplus x'$ of a *pair* of messages through the cipher.
- A *differential characteristic* $\alpha \xrightarrow{p} \beta$ predicts that an input difference α becomes an output difference β with probability p.

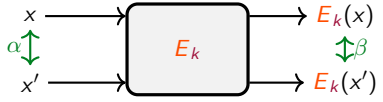

$$\Pr[\alpha \to \beta] = p \gg 2^{-n}$$

- If p is high enough, we can distinguish the cipher from a random permutation.

Differential cryptanalysis [BS91]

- Instead of looking at single messages, we follow the propagation of the difference $x \oplus x'$ of a *pair* of messages through the cipher.
- A *differential characteristic* $\alpha \xrightarrow{p} \beta$ predicts that an input difference α becomes an output difference β with probability p.

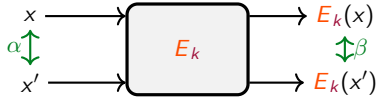

$$\Pr[\alpha \to \beta] = p \gg 2^{-n}$$

- If p is high enough, we can distinguish the cipher from a random permutation.
- A good non linear layer should prevent from high probability differentials.

Differential properties of the AES S-box

Let Δ_{in} be an 8-bit input difference and Δ_{out} an 8-bit output difference. Three cases are possible:

Differential properties of the AES S-box

Let Δ_{in} be an 8-bit input difference and Δ_{out} an 8-bit output difference. Three cases are possible:

- The differential transition is impossible for 129/256 pairs of differences $(\Delta_{in}, \Delta_{out})$

Differential properties of the AES S-box

Let Δ_{in} be an 8-bit input difference and Δ_{out} an 8-bit output difference. Three cases are possible:

- The differential transition is impossible for 129/256 pairs of differences $(\Delta_{in}, \Delta_{out})$
- There exist two ordered pairs of values: $(x, x \oplus \Delta_{in}), (x \oplus \Delta_{in}, x)$ that satisfy the transition for 126/256 pairs of differences

Differential properties of the AES S-box

Let Δ_{in} be an 8-bit input difference and Δ_{out} an 8-bit output difference. Three cases are possible:

- The differential transition is impossible for 129/256 pairs of differences $(\Delta_{in}, \Delta_{out})$
- There exist two ordered pairs of values: $(x, x \oplus \Delta_{in}), (x \oplus \Delta_{in}, x)$ that satisfy the transition for 126/256 pairs of differences
- There exist four ordered pairs of values that satisfy the transition for 1/256 pair of differences.

$\Longrightarrow$ For a given input/output difference, we either get 2, 4 or 0 input/output possible values for a total average of 1 solution.

Differential properties of the AES S-box

Let Δ_{in} be an 8-bit input difference and Δ_{out} an 8-bit output difference. Three cases are possible:

- The differential transition is impossible for 129/256 pairs of differences $(\Delta_{in}, \Delta_{out})$
- There exist two ordered pairs of values: $(x, x \oplus \Delta_{in}), (x \oplus \Delta_{in}, x)$ that satisfy the transition for 126/256 pairs of differences
- There exist four ordered pairs of values that satisfy the transition for 1/256 pair of differences.

$\Longrightarrow$ For a given input/output difference, we either get 2, 4 or 0 input/output possible values for a total average of 1 solution.

$\Longrightarrow$ The maximum transition probability is 2^{-6}.

The boomerang attack

It is unlikely to have a characteristic covering the whole cipher.

The boomerang attack

It is unlikely to have a characteristic covering the whole cipher.

- 1 differential characteristic ($\alpha \xrightarrow{p} \beta$) for the upper part of the cipher, and one for the bottom part ($\gamma \xrightarrow{q} \phi$).

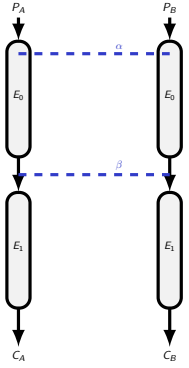

The boomerang attack

It is unlikely to have a characteristic covering the whole cipher.

- 1 differential characteristic ($\alpha \xrightarrow{p} \beta$) for the upper part of the cipher, and one for the bottom part ($\gamma \xrightarrow{q} \phi$).
- We encrypt one pair (P_A, P_B) with $P_A \oplus P_B = \alpha$ to (C_A, C_B), and we construct (C_C, C_D) such that $C_A \oplus C_C = C_B \oplus C_D = \gamma$.

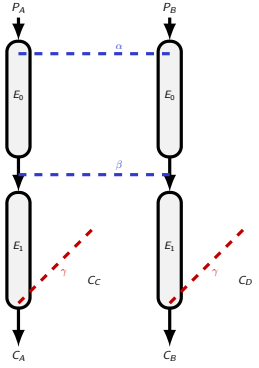

The boomerang attack

It is unlikely to have a characteristic covering the whole cipher.

- 1 differential characteristic ($\alpha \xrightarrow{p} \beta$) for the upper part of the cipher, and one for the bottom part ($\gamma \xrightarrow{q} \phi$).
- We encrypt one pair (P_A, P_B) with $P_A \oplus P_B = \alpha$ to (C_A, C_B), and we construct (C_C, C_D) such that $C_A \oplus C_C = C_B \oplus C_D = \gamma$.

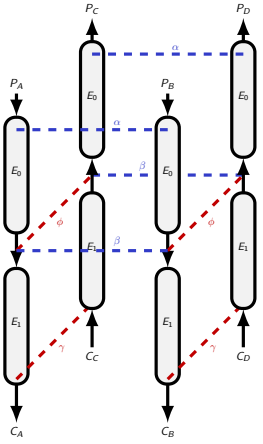

The boomerang attack

It is unlikely to have a characteristic covering the whole cipher.

- 1 differential characteristic ($\alpha \xrightarrow{p} \beta$) for the upper part of the cipher, and one for the bottom part ($\gamma \xrightarrow{q} \phi$).
- We encrypt one pair (P_A, P_B) with $P_A \oplus P_B = \alpha$ to (C_A, C_B), and we construct (C_C, C_D) such that $C_A \oplus C_C = C_B \oplus C_D = \gamma$.
- We decrypt C_C and C_D : then $(P_C \oplus P_D) = \alpha$ with probability p^2q^2, which is unlikely to happen for a random permutation.

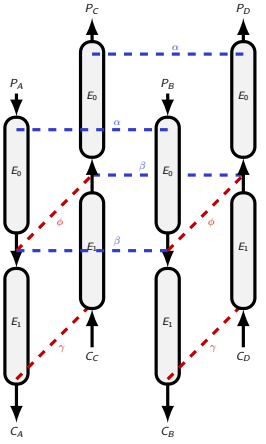

The boomerang attack

It is unlikely to have a characteristic covering the whole cipher.

- 1 differential characteristic ($\alpha \xrightarrow{p} \beta$) for the upper part of the cipher, and one for the bottom part ($\gamma \xrightarrow{q} \phi$).
- We encrypt one pair (P_A, P_B) with $P_A \oplus P_B = \alpha$ to (C_A, C_B), and we construct (C_C, C_D) such that $C_A \oplus C_C = C_B \oplus C_D = \gamma$.
- We decrypt C_C and C_D : then $(P_C \oplus P_D) = \alpha$ with probability p^2q^2, which is unlikely to happen for a random permutation.

It can be adapted to an attack where we don't use the decryption oracle, and the probability is then $2^{-n}p^2q^2$.

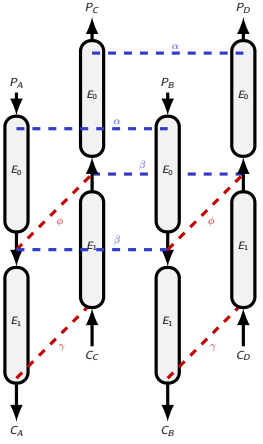

The related-key framework

- We allow the attacker to have more than one key, which are related to each other by mathematical relations.

The related-key framework

- We allow the attacker to have more than one key, which are related to each other by mathematical relations.
- This is used to evaluate the resistance of the cipher when it is used for hashing, where the message plays the role of the key.

The related-key framework

- We allow the attacker to have more than one key, which are related to each other by mathematical relations.
- This is used to evaluate the resistance of the cipher when it is used for hashing, where the message plays the role of the key.
- Versions 192 and 256 are already broken in this setting :
  - $2^{135}, 2^{120.5}, 2^{127.5}$ Time, data and memory for AES-192 [YSZ+24].
  - $2^{99.5}, 2^{99.5}, 2^{77}$ Time, data, and memory for AES-256 [BK09].

A related-key trail for a Boomerang attack on AES

$K_A \oplus K_B$, $K_C \oplus K_D$:

A related-key trail for a Boomerang attack on AES

$K_A \oplus K_B$, $K_C \oplus K_D$:

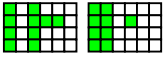

A related-key trail for a Boomerang attack on AES

$K_A \oplus K_B$, $K_C \oplus K_D$:

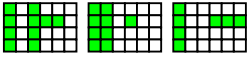

A related-key trail for a Boomerang attack on AES

$K_A \oplus K_B$, $K_C \oplus K_D$:

A related-key trail for a Boomerang attack on AES

$K_A \oplus K_B$, $K_C \oplus K_D$:

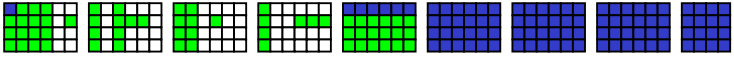

A related-key trail for a Boomerang attack on AES

$K_A \oplus K_B$, $K_C \oplus K_D$:

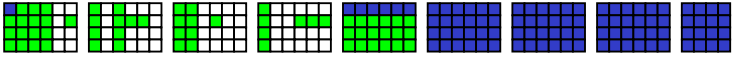

$K_A \oplus K_C$, $K_B \oplus K_D$:

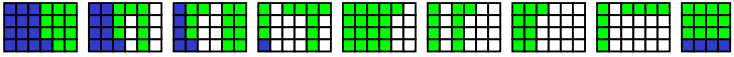

Disturbances and corrections

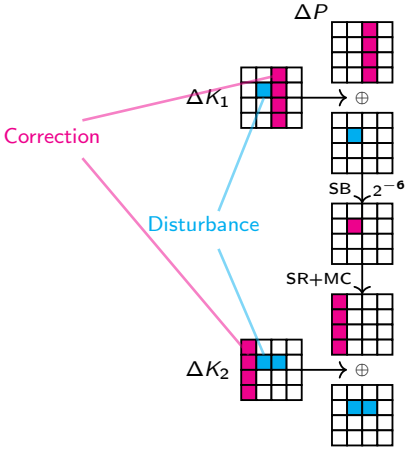

Disturbances and corrections in 1 round of AES

This section

Previous best attack [DEFN22]

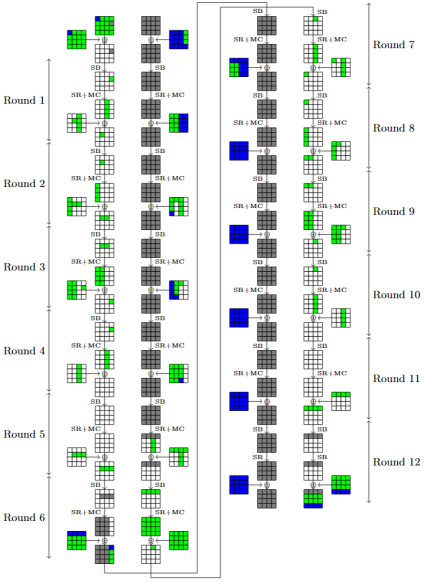

Previous best attack [DEFN22]

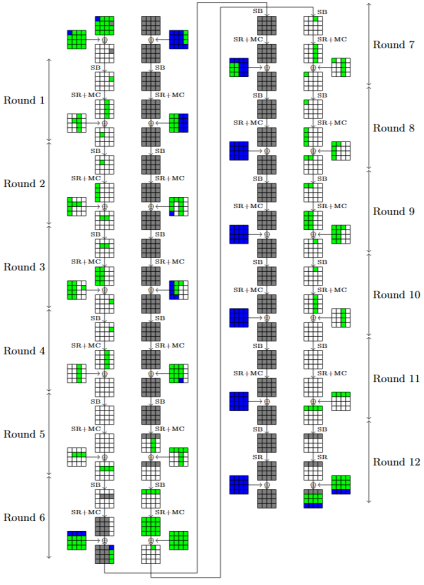

- Boomerang distinguisher of probability $2^{-108} = 2^{-2(4\times6+5\times6)}$

Previous best attack [DEFN22]

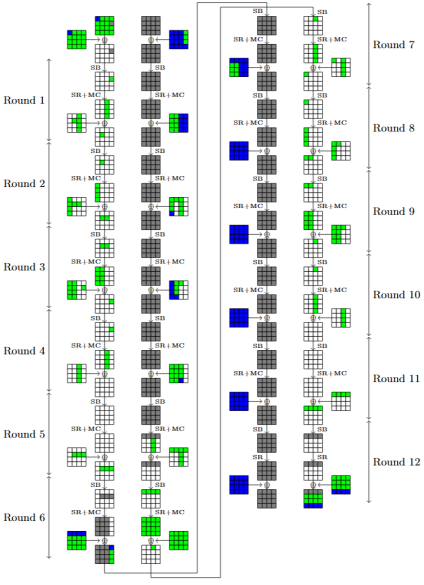

- Boomerang distinguisher of probability $2^{-108} = 2^{-2(4\times6+5\times6)}$
- Once we have a good quartet, some key bytes are recovered using the property of the s-box of the AES.

Key recovery and Filtering

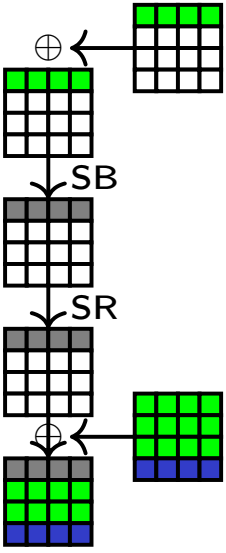

Key recovery from the last round

- A quartet gives us the value of the gray difference $S(a) + k, S(a + \delta) + k$.

Key recovery and Filtering

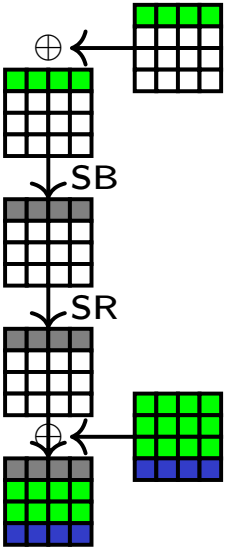

Key recovery from the last round

- A quartet gives us the value of the gray difference $S(a) + k, S(a + \delta) + k$.
- Using the property of the S-box, we can deduce the possible values $a, a + \delta, S(a), S(a + \delta)$, and have candidates k.

Key recovery and Filtering

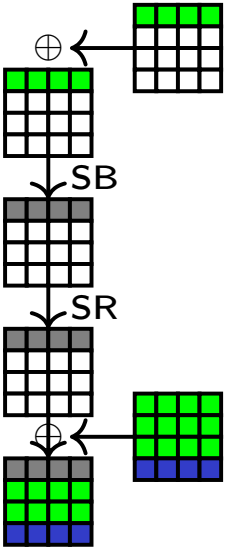

Key recovery from the last round

- A quartet gives us the value of the gray difference $S(a) + k, S(a + \delta) + k$.
- Using the property of the S-box, we can deduce the possible values $a, a + \delta, S(a), S(a + \delta)$, and have candidates k.
- If the difference cannot hold, then we filter the quartet, and if it holds, we get key candidates. In average we get 1 key candidate per quartet.

This section

"New" distinguisher

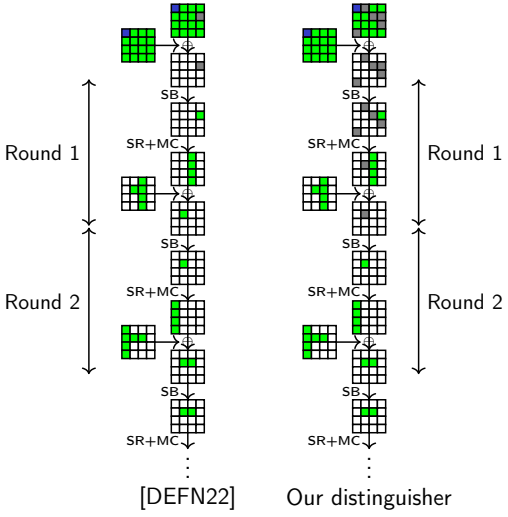

"New" distinguisher

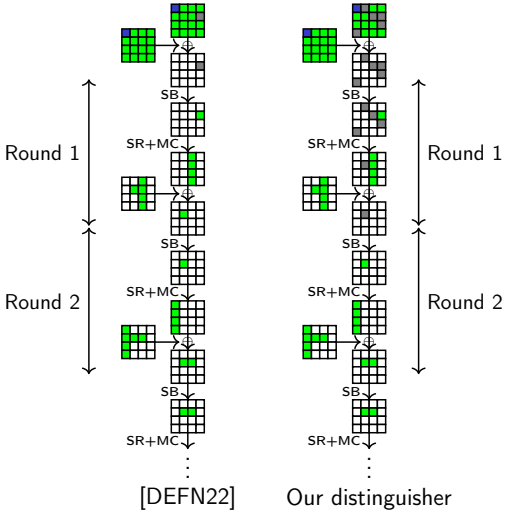

- Almost the same trail as in Asiacrypt 2022

"New" distinguisher

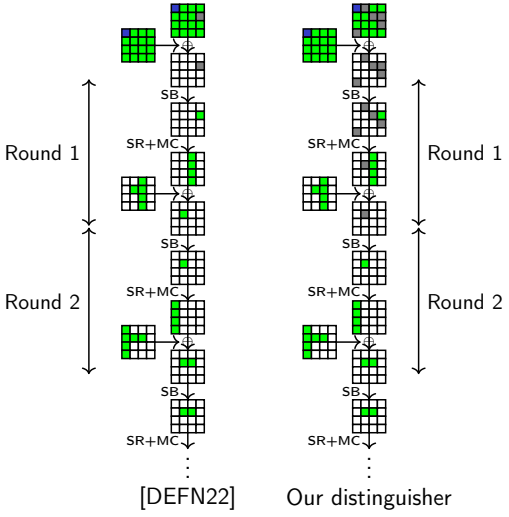

- Almost the same trail as in Asiacrypt 2022
- Some conditions in the first round are relaxed, increasing the probability of the boomerang distinguisher from 2^{-108} to 2^{-96}

"New" distinguisher

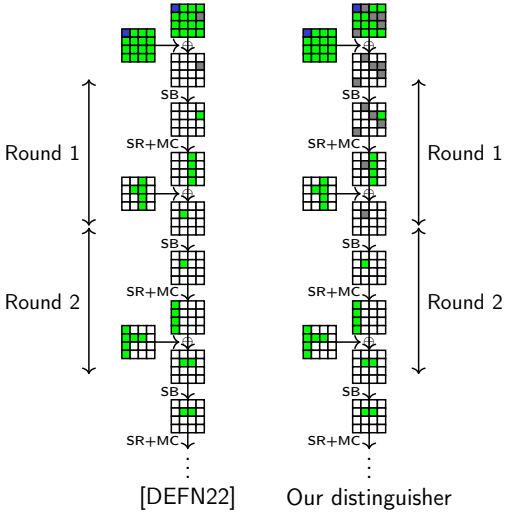

- Almost the same trail as in Asiacrypt 2022
- Some conditions in the first round are relaxed, increasing the probability of the boomerang distinguisher from 2^{-108} to 2^{-96}
- Different framework : Amplified boomerang vs boomerang

The attack

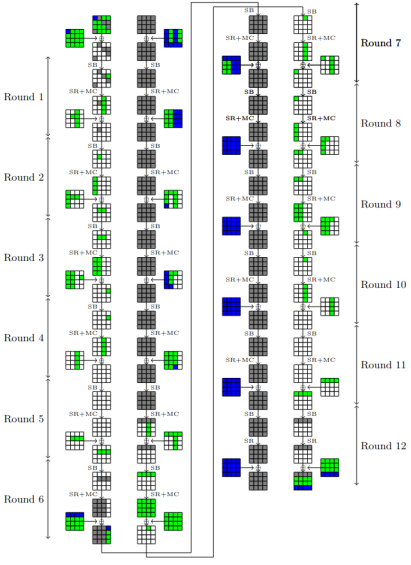

- Encrypt 2^{64} structures of 2^{48} plaintexts under the four related keys K_A, K_B, K_C, K_D and store the ciphertexts in four hashtables S_A, S_B, S_C, S_D.

The attack

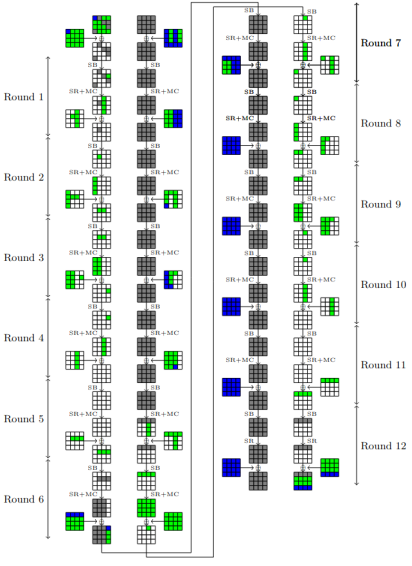

- Encrypt 2^{64} structures of 2^{48} plaintexts under the four related keys K_A, K_B, K_C, K_D and store the ciphertexts in four hashtables S_A, S_B, S_C, S_D.
- Look for collisions between S_A, S_C and between S_B, S_D. We will get 2^{262} possible quartets, that we filter according to the values of the input differences.

The attack

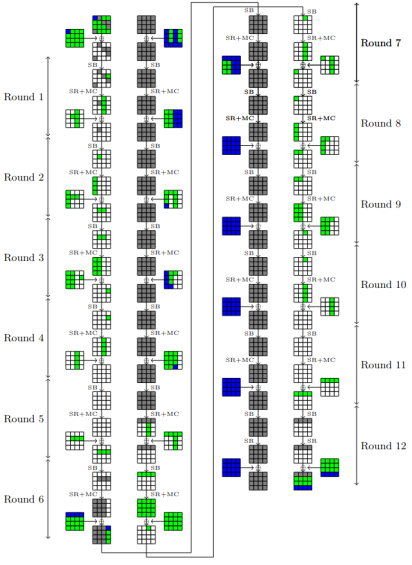

- Encrypt 2^{64} structures of 2^{48} plaintexts under the four related keys K_A, K_B, K_C, K_D and store the ciphertexts in four hashtables S_A, S_B, S_C, S_D.
- Look for collisions between S_A, S_C and between S_B, S_D. We will get 2^{262} possible quartets, that we filter according to the values of the input differences.
- Before the key recovery step, we will have 2^{130} quartets remaining.

Comparison between the Key trails

$K_A \oplus K_B$, $K_C \oplus K_D$:

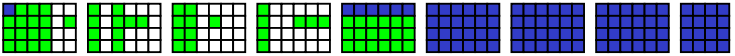

$K_A \oplus K_C$, $K_B \oplus K_D$:

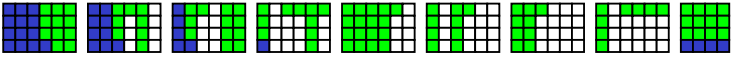

Related-keys trails from [DEFN22].

Comparison between the Key trails

$K_A \oplus K_B$, $K_C \oplus K_D$:

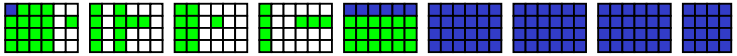

$K_A \oplus K_C$, $K_B \oplus K_D$:

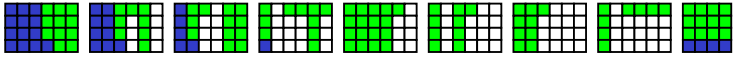

Related-keys trails from [DEFN22].

$K_A \oplus K_B$, $K_C \oplus K_D$:

$K_A \oplus K_C$, $K_B \oplus K_D$:

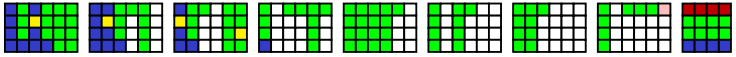

The same trail, but we deduce differences and values from known differences

Filtering

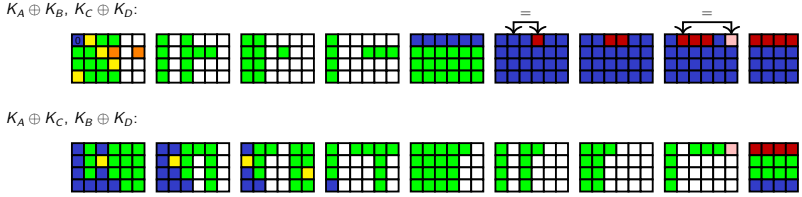

The same trail, but we deduce differences and values from known differences

Filtering

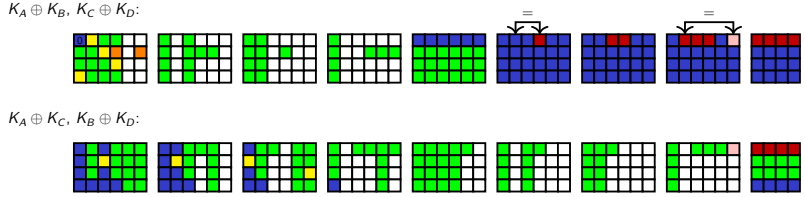

The same trail, but we deduce differences and values from known differences

- The AES key schedule is non linear, but have some linear properties that we can use. Overall we can reduce the quartets from 2^{130} to 2^{88}.

Filtering

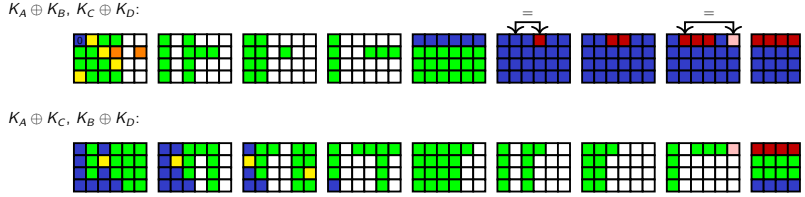

The same trail, but we deduce differences and values from known differences

- The AES key schedule is non linear, but have some linear properties that we can use. Overall we can reduce the quartets from 2^{130} to 2^{88}.
- The filtering also gives us key bytes, and we recover 18 key bytes per quartet. We determine the right key bytes by repeating the attack.

Filtering

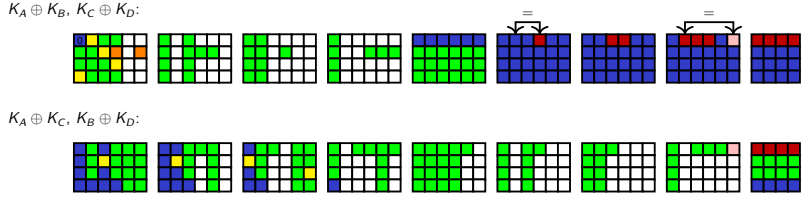

The same trail, but we deduce differences and values from known differences

- The AES key schedule is non linear, but have some linear properties that we can use. Overall we can reduce the quartets from 2^{130} to 2^{88}.
- The filtering also gives us key bytes, and we recover 18 key bytes per quartet. We determine the right key bytes by repeating the attack.

$\Longrightarrow$ Attack in $2^{133.5}, 2^{115.5}, 2^{131}$ time, data, memory.

This section

Previous attack on AES-256 [GSW22]

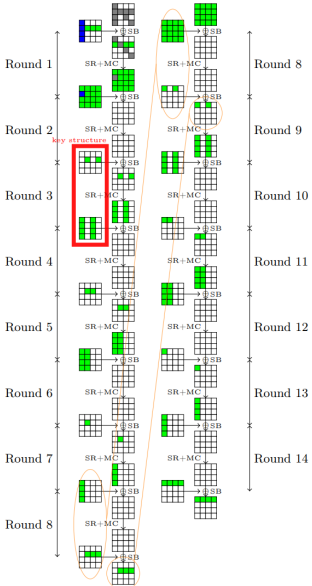

Previous attack on AES-256 [GSW22]

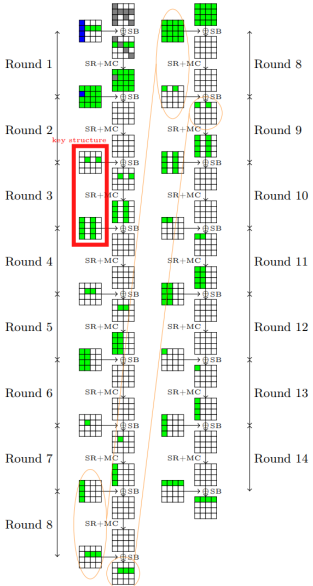

- Guo et al. build a key structure $S = \{K_i | K \oplus K_i = \Delta_i\}$ for the upper trail where K is the master key.

Previous attack on AES-256 [GSW22]

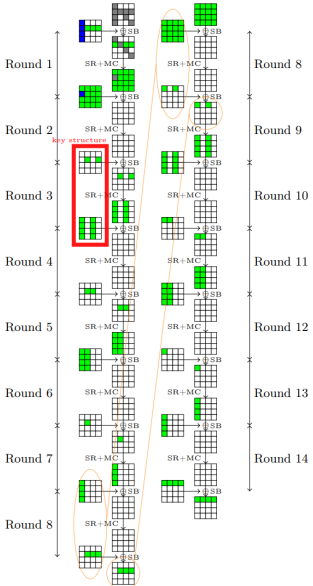

- Guo et al. build a key structure $S = \{K_i | K \oplus K_i = \Delta_i\}$ for the upper trail where K is the master key.
- For $1/2^8$ of the keys the transition probabilities will be 2^{-6} per S-box, but for the $\frac{2^7-2}{2^8}$, it will be 2^{-7}.

Previous attack on AES-256 [GSW22]

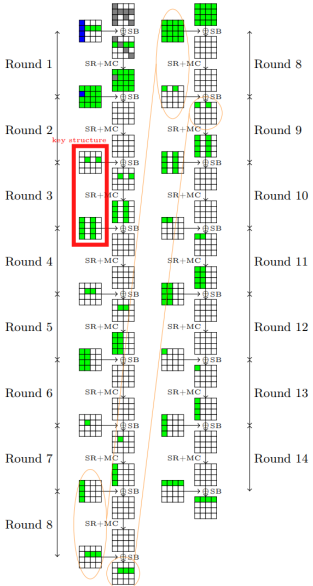

- Guo et al. build a key structure $S = \{K_i | K \oplus K_i = \Delta_i\}$ for the upper trail where K is the master key.
- For $1/2^8$ of the keys the transition probabilities will be 2^{-6} per S-box, but for the $\frac{2^7-2}{2^8}$, it will be 2^{-7}.
- the probability of the whole distinguisher can either be 2^{-106} or $2^{2(-6\times5-6\times3)} = 2^{-96}$.

Advantage of key structures

- 1 plaintext structure of size m with 1 key structure of size n, we would get $2^{2m+2n-1}$ pairs instead of having 2^{2n+m-1} pairs for the same data cost obtained by encrypting n plaintext structures of size m.

Advantage of key structures

- 1 plaintext structure of size m with 1 key structure of size n, we would get $2^{2m+2n-1}$ pairs instead of having 2^{2n+m-1} pairs for the same data cost obtained by encrypting n plaintext structures of size m.
- We need to know the differences between the related key bytes recovered and the master key, otherwise we learn the value of a key in the structure but not of the master key.

Key structures on AES-192

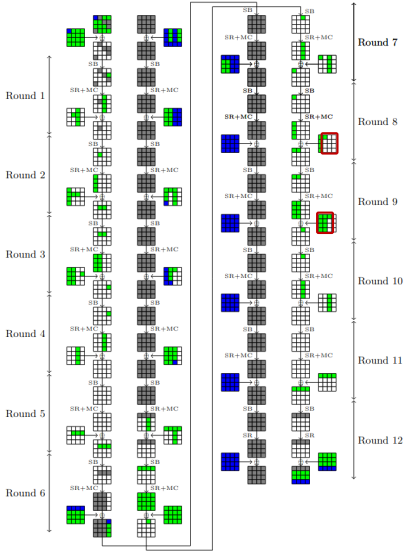

- We use a key structure of size 2^{16} from the last three columns of K^5 and the first three columns of K^6.

Key structures on AES-192

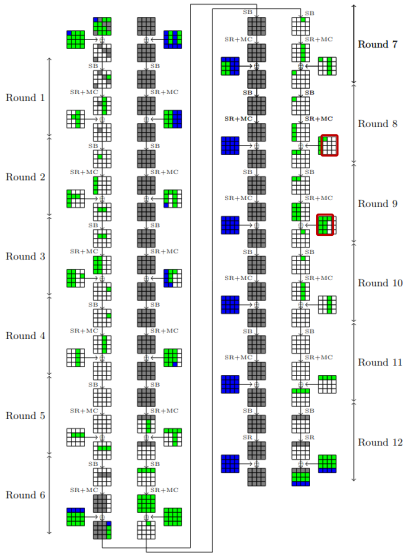

- We use a key structure of size 2^{16} from the last three columns of K^5 and the first three columns of K^6.
- The probability of the distinguisher will be either $2^{-2(4\times6)+2(5\times7)} = 2^{-118}$ or 2^{-108}.

Key structures on AES-192

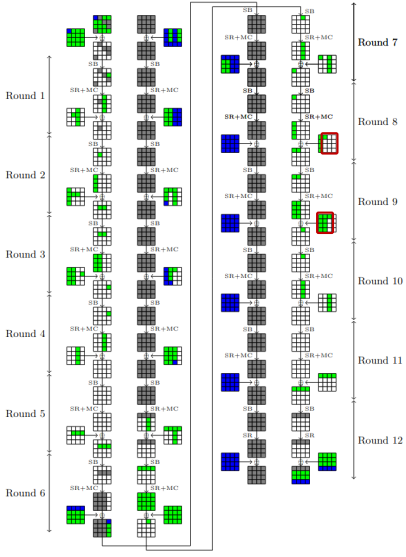

- We use a key structure of size 2^{16} from the last three columns of K^5 and the first three columns of K^6.
- The probability of the distinguisher will be either $2^{-2(4\times6)+2(5\times7)} = 2^{-118}$ or 2^{-108}.
- Starting from the decryption side, we decrypt 2^{40+16} ciphertexts, and we encrypt 2^{56+14} plaintexts.

Key structures on AES-192

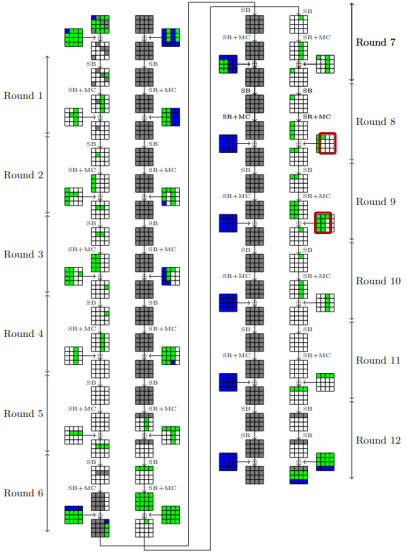

- We use a key structure of size 2^{16} from the last three columns of K^5 and the first three columns of K^6.
- The probability of the distinguisher will be either $2^{-2(4\times6)+2(5\times7)} = 2^{-118}$ or 2^{-108}.
- Starting from the decryption side, we decrypt 2^{40+16} ciphertexts, and we encrypt 2^{56+14} plaintexts.
- The average number of quartets we should get is $2^{71} \times \frac{2^7-2}{2^8} \times 2^{-118} = 2^{-48}$ or $2^{71} \times \frac{1}{2^8} \times 2^{-108} = 2^{-45}$, depending on the key pair used, And we will need 2^{47} structures using the distinguisher with probability 2^{-108}.

Results

Key size	#keys	Time	Data	Memory	Reference
192	4	2^{176}	2^{123}	2^{152}	[BK09]
	4	2^{124}	2^{124}	$2^{79.8}$	[DEFN22]
	4	$2^{135.5}$	$2^{120.5}$	$2^{127.5}$	[YSZ+24]
	4	$2^{133.5}$	$2^{115.5}$	2^{131}	Our work
	2^{17}	2^{117}	2^{117}	2^{78}	Our work
256	4	$2^{99.5}$	$2^{99.5}$	2^{77}	[BK09]
	2^{19-s}	2^{92+s}	2^{91+s}	2^{89-s}	[GSW22]

Best related-key attacks on AES

This section

Conclusion

- We found useful relations in the key schedule of AES.
- Using them, changing the framework from Boomerang to Amplified Boomerang, and relaxing the first round conditions allowed us to reduce the data of the attack on AES-192.
- We found new tradeoffs by using related key structures.
- Even if automated tools help a lot, we often have to try to improve the attacks by hand.

Conclusion

- We found useful relations in the key schedule of AES.
- Using them, changing the framework from Boomerang to Amplified Boomerang, and relaxing the first round conditions allowed us to reduce the data of the attack on AES-192.
- We found new tradeoffs by using related key structures.
- Even if automated tools help a lot, we often have to try to improve the attacks by hand.

Thank you !

References I

Alex Biryukov and Dmitry Khovratovich.

Related-key cryptanalysis of the full AES-192 and AES-256.

In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2009.

Eli Biham and Adi Shamir.

Differential cryptanalysis of DES-like cryptosystems.

Journal of Cryptology, 4(1):3–72, 1991.

Patrick Derbez, Marie Euler, Pierre-Alain Fouque, and Phuong Hoa Nguyen.

Revisiting related-key boomerang attacks on AES using computer-aided tool.

In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology - ASIACRYPT 2022 - 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5-9, 2022, Proceedings, Part III*, volume 13793 of *Lecture Notes in Computer Science*, pages 68–88. Springer, 2022.

References II

Joan Daemen and Vincent Rijmen.

The Design of Rijndael: AES - The Advanced Encryption Standard (Information Security and Cryptography).

Springer, 1 edition, 2002.

Jian Guo, Ling Song, and Haoyang Wang.

Key structures: Improved related-key boomerang attack against the full aes-256.

In Khoa Nguyen, Guomin Yang, Fuchun Guo, and Willy Susilo, editors, *Information Security and Privacy*, pages 3–23, Cham, 2022. Springer International Publishing.

Jérémy Jean.

TikZ for Cryptographers.

https://www.iacr.org/authors/tikz/, 2016.

Qianqian Yang, Ling Song, Nana Zhang, Danping Shi, Libo Wang, Jiahao Zhao, Lei Hu, and Jian Weng.

Optimizing rectangle and boomerang attacks: A unified and generic framework for key recovery.

J. Cryptol., 37(2):19, 2024.