TU Graz

SCIENCE
PASSION
TECHNOLOGY

Analyzing the Linear Keystream Biases in AEGIS

Maria Eichlseder Marcel Nageler Robert Primas

Dagstuhl Seminar “Symmetric Cryptography” 2020

Motivation

The CAESAR Competition

Competition for **A**uthenticated **E**ncryption: **S**ecurity, **A**pplicability, **R**obustness

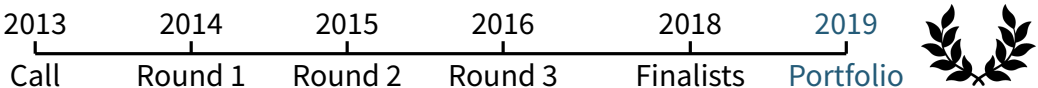

- Use-case 1: Lightweight cryptography
- Use-case 2: High SW performance
  - `AEGIS-128`
  - `OCB`
- Use-case 3: Robustness

AEGIS

- Design by Wu and Preneel [WP13; WP16]
- Family of authenticated ciphers:
  - AEGIS-128 (final portfolio)
  - AEGIS-128L (finalist)
  - AEGIS-256 (finalist)
- High SW performance thanks to AES-NI

AEGIS – Previous Analysis

Designers' analysis [WP16]

- Focus on weak states, generic and differential attacks
- Conservative bound based on diff. active S-boxes: $p < 2^{-150}$ for all variants

Analysis in misuse settings [KEM17; VV18]

Linear cryptanalysis by Minaud [Min14]

- Linear characteristics of the round function
- $c^2 = 2^{-154}$ for AEGIS-128, $c^2 = 2^{-178}$ for AEGIS-256
- Application: Linear approximation of the keystream (KP)
- Very high data requirements, but can be collected across different keys

In the Meantime…

Minaud's analysis inspired similar attacks on another CAESAR finalist, MORUS:

- Ashur et al. [AEL+18] proposed a linear distinguisher (found by hand) and discussed how such keystream correlations could be exploited in practice.
- Shi et al. [SSS+19] substantially improved the distinguishers using MILP. ("substantial" = from $c^{-2} \geq 2^{146}$ to $c^{-2} = 2^{76}$!)

Note: MORUS has a very different round function, but a somewhat similar mode.

Question: So what about AEGIS?

In the Meantime…

Minaud's analysis inspired similar attacks on another CAESAR finalist, `MORUS`:

- Ashur et al. [AEL+18] proposed a linear distinguisher (found by hand) and discussed how such keystream correlations could be exploited in practice.
- Shi et al. [SSS+19] substantially improved the distinguishers using MILP. ("substantial" = from $c^{-2} \geq 2^{146}$ to $c^{-2} = 2^{76}$!)

Note: `MORUS` has a very different round function, but a somewhat similar mode.

Question: So what about `AEGIS`?

AEGIS

AEGIS Authenticated Encryption

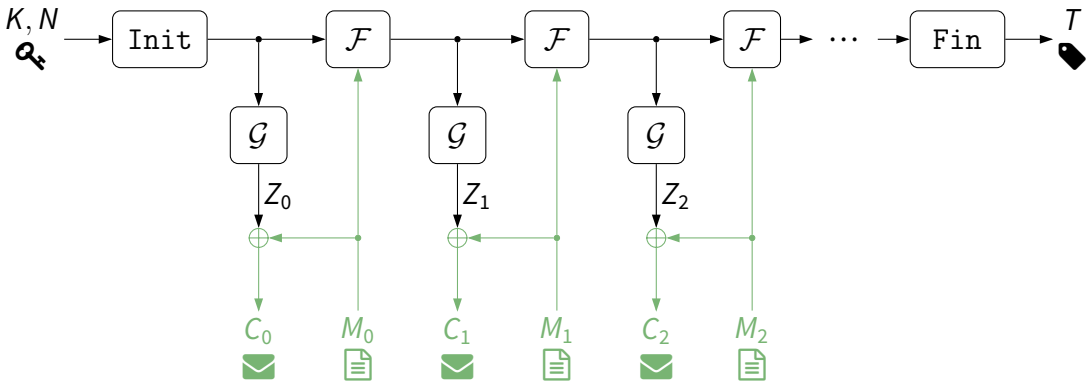

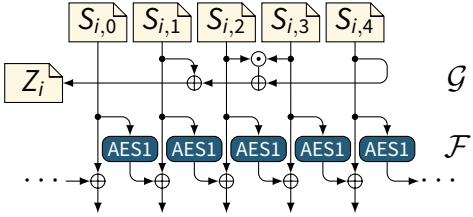

(a) AEGIS-**128**

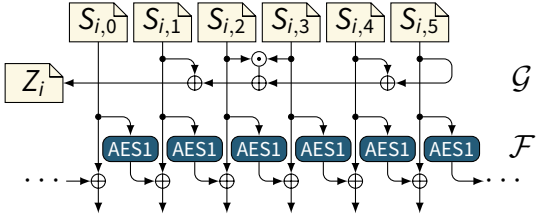

(b) AEGIS-**256**

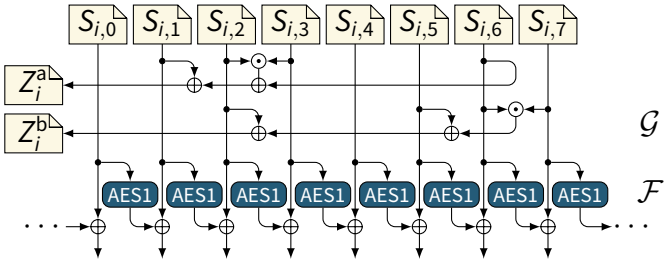

(c) AEGIS-**128L**

Linear keystream distinguisher [Min14]

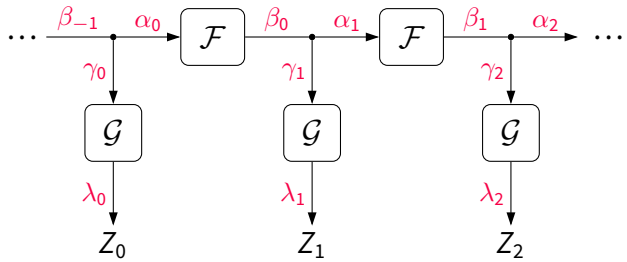

- Exploit keystream bias of $\lambda_0 \cdot Z_0 \oplus \lambda_1 \cdot Z_1 \oplus \lambda_2 \cdot Z_2$
- Correlation contribution $c = \prod_i (2p_i - 1) \rightarrow$ data complexity about c^{-2} KP

Our Results

Bounds for the inv. squared correlation contribution c^{-2} of best suitable linear charact.:

	AEGIS-128	AEGIS-256	AEGIS-128L
[Min14] (manual)	$c^{-2} \leq 2^{154}$	$c^{-2} \leq 2^{178}$	
Truncated model	$2^{92} \leq c^{-2}$	$2^{116} \leq c^{-2}$	$2^{114} \leq c^{-2} \leq 2^{172}$
Improved model	$2^{102} \leq c^{-2} \leq 2^{140}$	$2^{120} \leq c^{-2}$	
Bitwise model	$2^{132} \leq c^{-2} \leq 2^{140}$	$2^{152} \leq c^{-2} \leq 2^{162}$	$2^{140} \leq c^{-2} \leq 2^{152}$
	MILP	CP	

- Distinguishing complexity is likely a bit below c^{-2} blocks (using multiple approx.). [Min14] estimates 2^{140} for AEGIS-128.
- There are dependencies (no key, $\mathcal{F}$ and $\mathcal{G}$ evaluated separately), linear hull effect, …

Our Results

Bounds for the inv. squared correlation contribution c^{-2} of best suitable linear charact.:

	AEGIS-128	AEGIS-256	AEGIS-128L
[Min14] (manual)	$c^{-2} \leq 2^{154}$	$c^{-2} \leq 2^{178}$	
Truncated model	$2^{92} \leq c^{-2}$	$2^{116} \leq c^{-2}$	$2^{114} \leq c^{-2} \leq 2^{172}$
Improved model	$2^{102} \leq c^{-2} \leq 2^{140}$	$2^{120} \leq c^{-2}$	
Bitwise model	$2^{132} \leq c^{-2} \leq 2^{140}$	$2^{152} \leq c^{-2} \leq 2^{162}$	$2^{140} \leq c^{-2} \leq 2^{152}$

MILP CP

- Distinguishing complexity is likely a bit below c^{-2} blocks (using multiple approx.). [Min14] estimates 2^{140} for AEGIS-128.
- There are dependencies (no key, $\mathcal{F}$ and $\mathcal{G}$ evaluated separately), linear hull effect, …

Our Results

Bounds for the inv. squared correlation contribution c^{-2} of best suitable linear charact.:

	AEGIS-128	AEGIS-256	AEGIS-128L
[Min14] (manual)	$c^{-2} \leq 2^{154}$	$c^{-2} \leq 2^{178}$	
Truncated model	$2^{92} \leq c^{-2}$	$2^{116} \leq c^{-2}$	$2^{114} \leq c^{-2} \leq 2^{172}$
Improved model	$2^{102} \leq c^{-2} \leq 2^{140}$	$2^{120} \leq c^{-2}$	
Bitwise model	$2^{132} \leq c^{-2} \leq 2^{140}$	$2^{152} \leq c^{-2} \leq 2^{162}$	$2^{140} \leq c^{-2} \leq 2^{152}$

MILP CP

- Distinguishing complexity is likely a bit below c^{-2} blocks (using multiple approx.). [Min14] estimates 2^{140} for AEGIS-128.
- There are dependencies (no key, $\mathcal{F}$ and $\mathcal{G}$ evaluated separately), linear hull effect, …

Our Results

Bounds for the inv. squared correlation contribution c^{-2} of best suitable linear charact.:

	AEGIS-128	AEGIS-256	AEGIS-128L
[Min14] (manual)	$c^{-2} \leq 2^{154}$	$c^{-2} \leq 2^{178}$	
Truncated model	$2^{92} \leq c^{-2}$	$2^{116} \leq c^{-2}$	$2^{114} \leq c^{-2} \leq 2^{172}$
Improved model	$2^{102} \leq c^{-2} \leq 2^{140}$	$2^{120} \leq c^{-2}$	
Bitwise model	$2^{132} \leq c^{-2} \leq 2^{140}$	$2^{152} \leq c^{-2} \leq 2^{162}$	$2^{140} \leq c^{-2} \leq 2^{152}$

MILP CP

- Distinguishing complexity is likely a bit below c^{-2} blocks (using multiple approx.). [Min14] estimates 2^{140} for AEGIS-128.
- There are dependencies (no key, $\mathcal{F}$ and $\mathcal{G}$ evaluated separately), linear hull effect, …

Our Results

Bounds for the inv. squared correlation contribution c^{-2} of best suitable linear charact.:

	AEGIS-128	AEGIS-256	AEGIS-128L
[Min14] (manual)	$c^{-2} \leq 2^{154}$	$c^{-2} \leq 2^{178}$	
Truncated model	$2^{92} \leq c^{-2}$	$2^{116} \leq c^{-2}$	$2^{114} \leq c^{-2} \leq 2^{172}$
Improved model	$2^{102} \leq c^{-2} \leq 2^{140}$	$2^{120} \leq c^{-2}$	
Bitwise model	$2^{132} \leq c^{-2} \leq 2^{140}$	$2^{152} \leq c^{-2} \leq 2^{162}$	$2^{140} \leq c^{-2} \leq 2^{152}$
	MILP ↑ CP ↑		

- Distinguishing complexity is likely a bit below c^{-2} blocks (using multiple approx.). [Min14] estimates 2^{140} for AEGIS-128.
- There are dependencies (no key, $\mathcal{F}$ and $\mathcal{G}$ evaluated separately), linear hull effect, …

New Bounds and Attacks

Simple Truncated Model

- Variables: active bytes
- Constraints: MC ($\mathcal{B}=5$), $\vdash$ ($\mathcal{B}=2$)
- Minimize $\sum w = \sum \max \log_2 c^{-2}$: S-box ($w=6$), AND ($w=2$)
- Fix nr of blocks: $\beta_{-1}=\alpha_k=0$

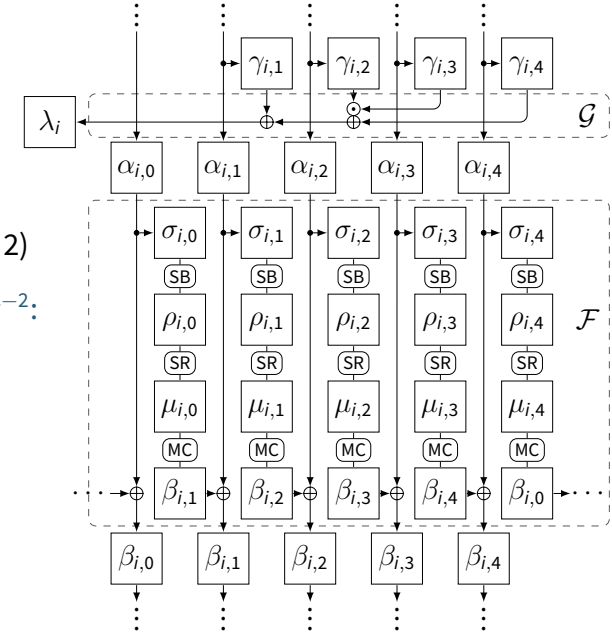

Simple Truncated Model – Results

- No solutions for ≤ 2 blocks
- Best results for 3 blocks
  - $c^{-2} \geq 2^{92}$ for AEGIS-128
  - $c^{-2} \geq 2^{116}$ for AEGIS-256
  - But impossible to find valid corresponding bitwise characteristics!
- Much higher cost ($>$ [Min14]) for ≥ 4 blocks

Inconsistent Trunc. Char. for AEGIS-128 (11 S-boxes, 13 ANDs)

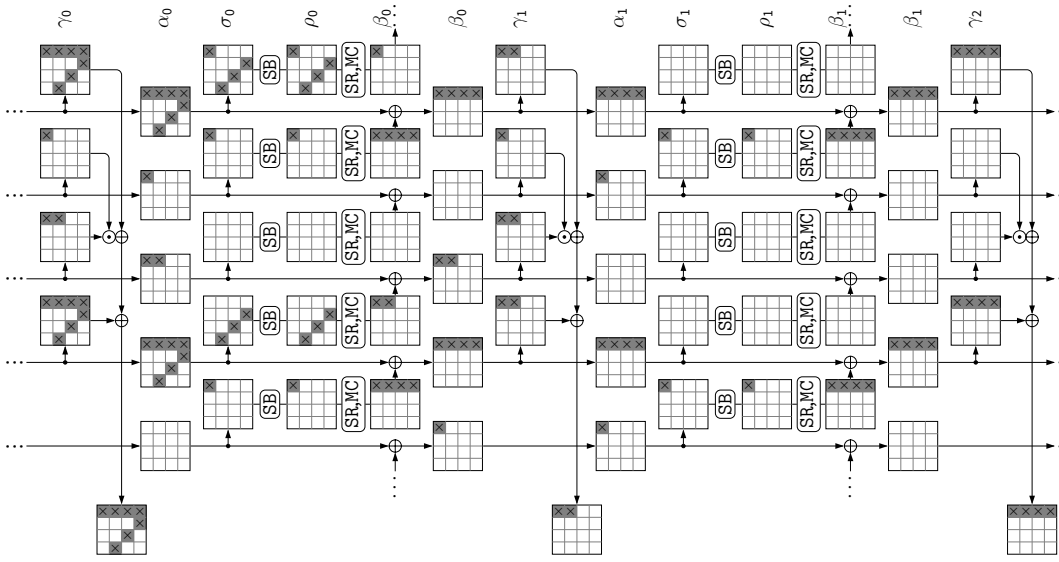

Inconsistent Trunc. Char. for AEGIS-128 (11 S-boxes, 13 ANDs)

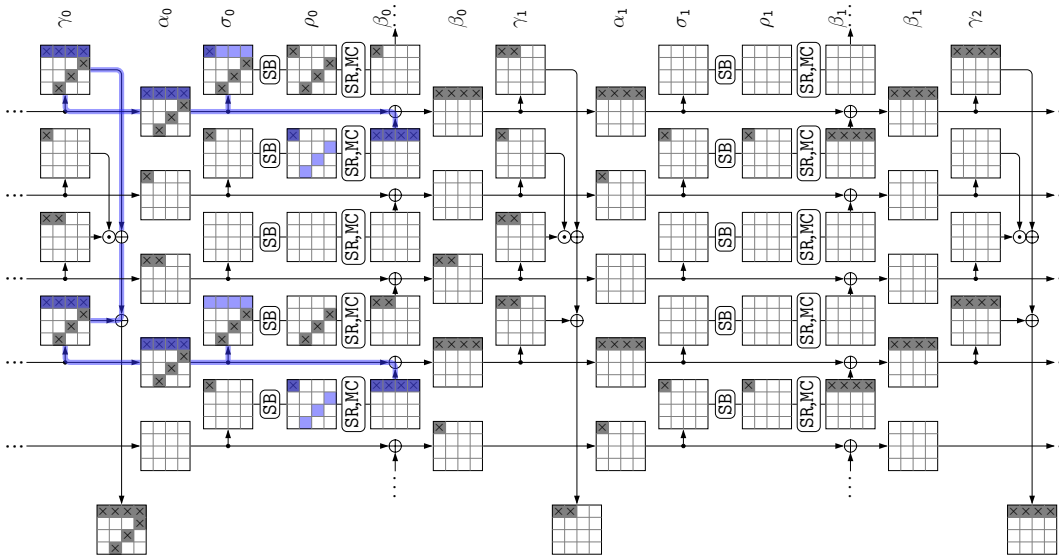

Inconsistent Trunc. Char. for AEGIS-128 (11 S-boxes, 13 ANDs)

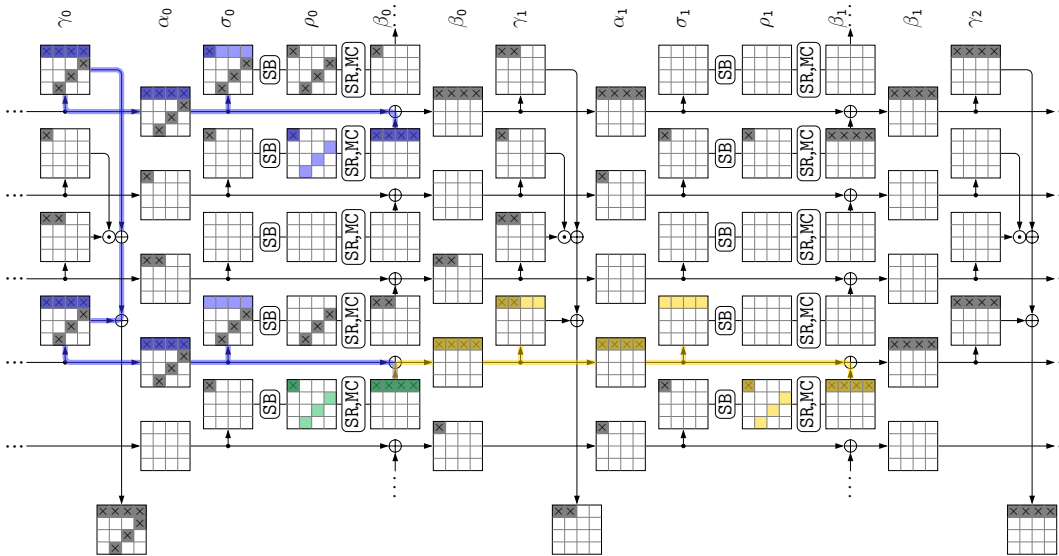

Improved Truncated Model

- Take any pair $(\alpha_0, \beta_0), (\alpha_1, \beta_1)$ of input/output masks for `MixColumns`. Consider their difference $(\alpha_0 \oplus \alpha_1, \beta_0 \oplus \beta_1)$.
- Then this difference must also satisfy the branch number $\mathcal{B} = 5$ of `MC`. The same is true for higher-order differences.

- Improved model: Identify paths of linear operations between pairs of `MixColumns` operations, which define such a (higher-order) difference.

  Add variables for their mask difference and constrain branch-number.

- Full 2-round model of `AEGIS-128`:
  ca. 2500 variables, 2500 constraints, consistent results $\rightarrow 2^{102} \leq c^{-2} \leq 2^{140}$

Improved Truncated Model

- Take any pair $(\alpha_0, \beta_0), (\alpha_1, \beta_1)$ of input/output masks for `MixColumns`. Consider their difference $(\alpha_0 \oplus \alpha_1, \beta_0 \oplus \beta_1)$.
- Then this difference must also satisfy the branch number $\mathcal{B} = 5$ of `MC`. The same is true for higher-order differences.

- Improved model: Identify paths of linear operations between pairs of `MixColumns` operations, which define such a (higher-order) difference.

  Add variables for their mask difference and constrain branch-number.

- Full 2-round model of `AEGIS-128`:
  ca. 2500 variables, 2500 constraints, consistent results $\rightarrow 2^{102} \leq c^{-2} \leq 2^{140}$

Improved Truncated Model

- Take any pair $(\alpha_0, \beta_0), (\alpha_1, \beta_1)$ of input/output masks for `MixColumns`. Consider their difference $(\alpha_0 \oplus \alpha_1, \beta_0 \oplus \beta_1)$.
- Then this difference must also satisfy the branch number $\mathcal{B} = 5$ of `MC`. The same is true for higher-order differences.

- Improved model: Identify paths of linear operations between pairs of `MixColumns` operations, which define such a (higher-order) difference.

  Add variables for their mask difference and constrain branch-number.

- Full 2-round model of `AEGIS-128`:
  ca. 2500 variables, 2500 constraints, consistent results $\rightarrow 2^{102} \leq c^{-2} \leq 2^{140}$

Bitwise MILP Model

Results for improved model:

- There is still a big gap between bound and best characteristic
- Main reason: Cost of AND-gates ($2^{-16} \leq c^2 \leq 2^{-2}$ per byte)

Partially bitwise model:

- AND, XOR, MC model: bitwise specification
- S-box model: any transition with $c^2 = 2^{-6}$ (reality: 93 % possible, $2^{-12} \leq c^2 \leq 2^{-6}$)

	AEGIS-128	AEGIS-256	AEGIS-128L
Bitwise model	$2^{132} \leq c^{-2} \leq 2^{140}$	$2^{152} \leq c^{-2} \leq 2^{162}$	$2^{140} \leq c^{-2} \leq 2^{152}$
	MILP ↑	CP ↑	

Constraint Programming (CP) Model for Finding Characteristics

- Fix truncated characteristic
  - because there are sooo many S-boxes
- Only allow a few (high-probability) S-box transitions
  - because the S-box LAT is sooo large and dense
- Soft constraints for heuristically minimizing the cost (Z3 solver)
  - “you must use the best transition; if you don’t, there’s a penalty”

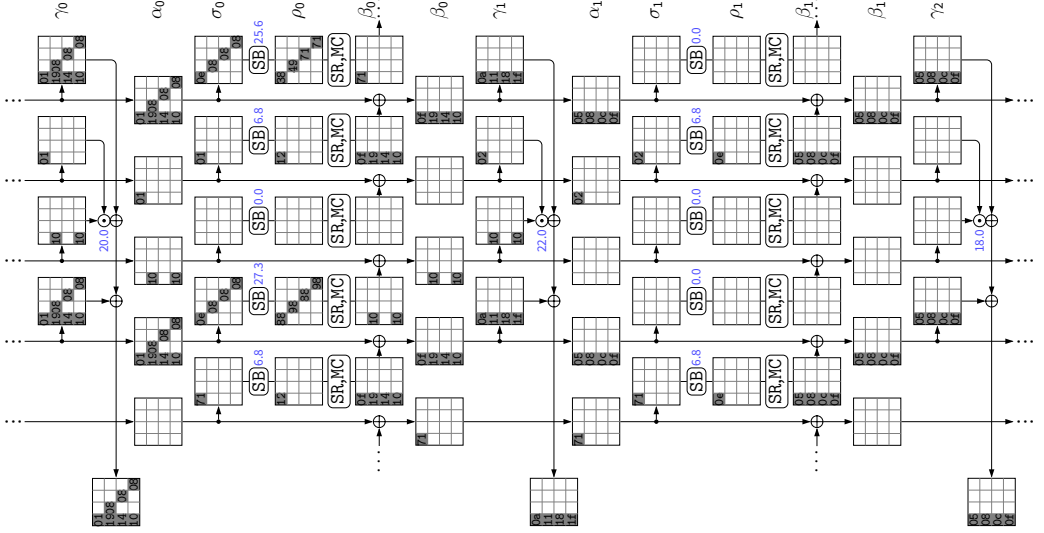

Linear approximation for AEGIS-128 with squared correlation contribution 2^{-140}

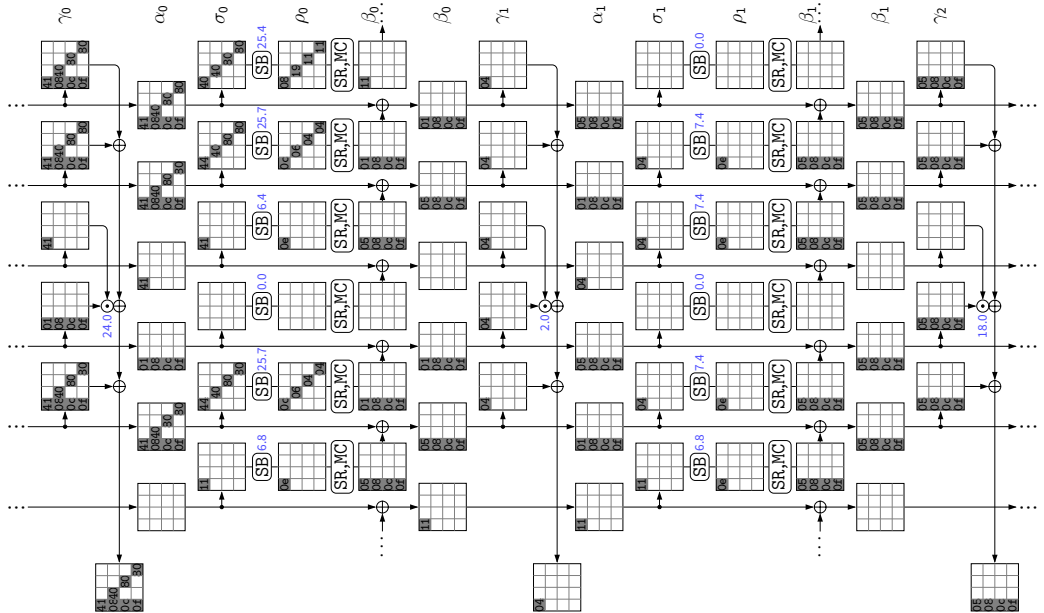

Linear approximation for AEGIS-256 with squared correlation contribution 2^{-162}

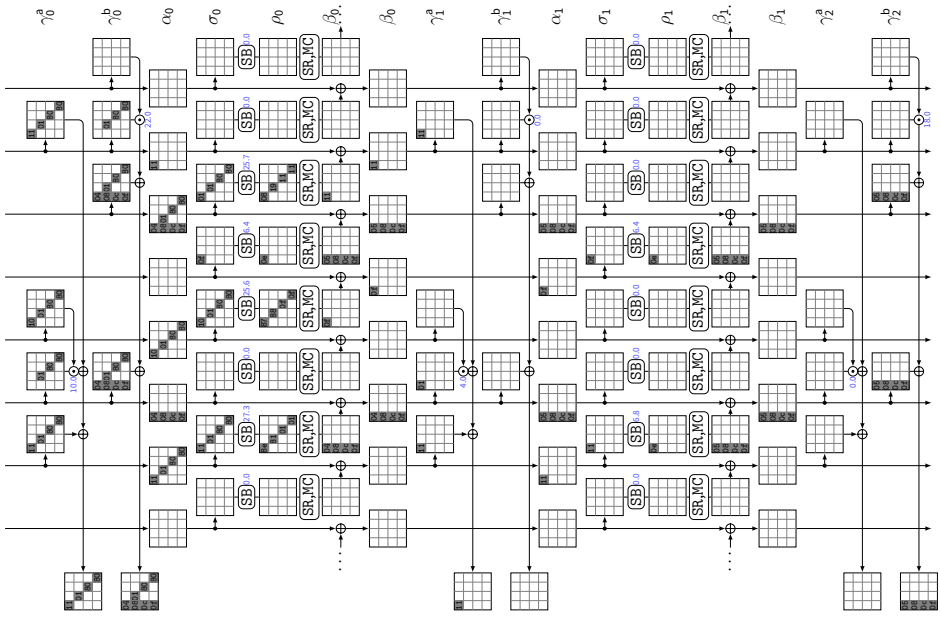

Linear approximation for AEGIS-128L with squared correlation contrib. 2^{-152}

Final Remarks

Experimental Verification

Note the dependencies between consecutive AND & SB:

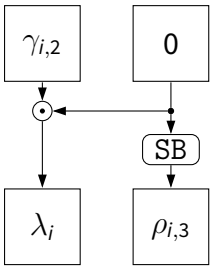

(a) $\beta_{i-1,3} \oplus \beta_{i,3} = 0$

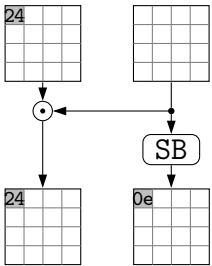

(b) $2^{-7.66}$ instead of $2^{-4-6.39}$

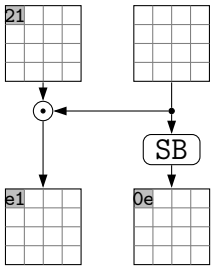

(c) 0 instead of $2^{-8-6.39}$

The best possible combined squared correlation is $2^{-7.36}$ (instead of 2^{-2-6}).

Experimental Verification

Note the dependencies between consecutive AND & SB:

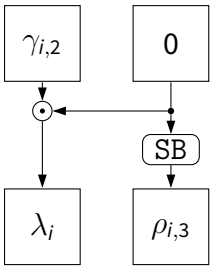

(a) $\beta_{i-1,3} \oplus \beta_{i,3} = 0$

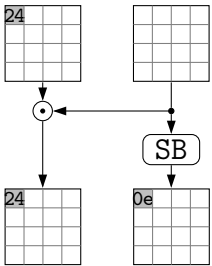

(b) $2^{-7.66}$ instead of $2^{-4-6.39}$

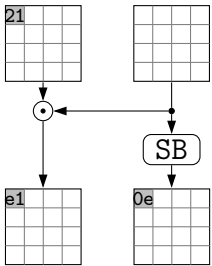

(c) 0 instead of $2^{-8-6.39}$

The best possible combined squared correlation is $2^{-7.36}$ (instead of 2^{-2-6}).

Cross-round Conditions in Characteristics

Some other examples of cross-round/-substate inconsistencies in characteristics:

In truncated characteristics:

- BCs: Related-key differential characteristics, e.g., `AES` [FJP13; GLMS18]
- TBCs: Related-tweakey differential characteristics, e.g., `Deoxys` [CHP+17]

In characteristics (also within substates):

- ARX hash functions, e.g., `MD4` [WLF+05], `SHA-2` [MNS11], `SHA-3` finalists [Leu12]

Conclusion

- We proposed improved keystream approximations for the AEGIS family, as well as upper bounds for c^2 below 2^{-128} (with some caveats).
- Straightforward models only produce very weak bounds and no solutions.
- https://extgit.iaik.tugraz.at/krypto/aegis_linear_trails

Questions

Bibliography I

[AEL+18] Tomer Ashur, Maria Eichlseder, Martin M. Lauridsen, Gaëtan Leurent, Brice Minaud, Yann Rotella, Yu Sasaki, and Benoît Viguier. **Cryptanalysis of MORUS**. Advances in Cryptology – ASIACRYPT 2018. Vol. 11273. LNCS. Springer, 2018, pp. 35–64. DOI: 10.1007/978-3-030-03329-3_2.

[CHP+17] Carlos Cid, Tao Huang, Thomas Peyrin, Yu Sasaki, and Ling Song. **A Security Analysis of Deoxys and its Internal Tweakable Block Ciphers**. IACR Transactions on Symmetric Cryptology 2017.3 (2017), pp. 73–107. DOI: 10.13154/tosc.v2017.i3.73-107.

[FJP13] Pierre-Alain Fouque, Jérémy Jean, and Thomas Peyrin. **Structural Evaluation of AES and Chosen-Key Distinguisher of 9-Round AES-128**. Advances in Cryptology – CRYPTO 2013. Vol. 8042. LNCS. Springer, 2013, pp. 183–203. DOI: 10.1007/978-3-642-40041-4_11.

Bibliography II

[GLMS18] David Gérault, Pascal Lafourcade, Marine Minier, and Christine Solnon. **Revisiting AES related-key differential attacks with constraint programming.** Information Processing Letters 139 (2018), pp. 24–29.

[KEM17] Daniel Kales, Maria Eichlseder, and Florian Mendel. **Note on the Robustness of CAESAR Candidates.** IACR Cryptology ePrint Archive, Report 2017/1137. 2017. URL: http://eprint.iacr.org/2017/1137.

[Leu12] Gaëtan Leurent. **Analysis of Differential Attacks in ARX Constructions.** Advances in Cryptology – ASIACRYPT 2012. Vol. 7658. LNCS. Springer, 2012, pp. 226–243. DOI: 10.1007/978-3-642-34961-4_15.

[Min14] Brice Minaud. **Linear Biases in AEGIS Keystream.** Selected Areas in Cryptography – SAC 2014. Vol. 8781. LNCS. Springer, 2014, pp. 290–305. DOI: 10.1007/978-3-319-13051-4_18.

Bibliography III

[MNS11] Florian Mendel, Tomislav Nad, and Martin Schläffer. **Finding SHA-2 Characteristics: Searching through a Minefield of Contradictions**. Advances in Cryptology – ASIACRYPT 2011. Vol. 7073. LNCS. Springer, 2011, pp. 288–307. DOI: 10.1007/978-3-642-25385-0_16.

[SSS+19] Danping Shi, Siwei Sun, Yu Sasaki, Chaoyun Li, and Lei Hu. **Correlation of Quadratic Boolean Functions: Cryptanalysis of All Versions of Full MORUS**. Advances in Cryptology – CRYPTO 2019. Vol. 11693. LNCS. Springer, 2019, pp. 180–209. DOI: 10.1007/978-3-030-26951-7_7.

[VV18] Serge Vaudenay and Damian Vizár. **Can Caesar Beat Galois? – Robustness of CAESAR Candidates Against Nonce Reusing and High Data Complexity Attacks**. Applied Cryptography and Network Security – ACNS 2018. Vol. 10892. LCNS. Springer, 2018, pp. 476–494. DOI: 10.1007/978-3-319-93387-0_25.

Bibliography IV

[WLF+05] Xiaoyun Wang, Xuejia Lai, Dengguo Feng, Hui Chen, and Xiuyuan Yu. **Cryptanalysis of the Hash Functions MD4 and RIPEMD**. Advances in Cryptology – EUROCRYPT 2005. Vol. 3494. LNCS. Springer, 2005, pp. 1–18. DOI: 10.1007/11426639_1.

[WP13] Hongjun Wu and Bart Preneel. **AEGIS: A Fast Authenticated Encryption Algorithm**. Selected Areas in Cryptography – SAC 2013. Vol. 8282. LNCS. Springer, 2013, pp. 185–201. DOI: 10.1007/978-3-662-43414-7_10.

[WP16] Hongjun Wu and Bart Preneel. **AEGIS: A Fast Authenticated Encryption Algorithm (v1.1)**. Submission to CAESAR: Competition for Authenticated Encryption. Security, Applicability, and Robustness (Round 3 and Final Portfolio). http://competitions.cr.yp.to/round3/aegisv11.pdf. Sept. 2016.