Analyzing the Linear Keystream Biases in AEGIS

Maria Eichlseder **Marcel Nageler** Robert Primas

FSE 2020 - Athens

Outline

Motivation

The CAESAR Competition

Competition for Authenticated Encryption: Security, Applicability, Robustness

2013	2014	2015	2016	2018	2019
Call	Round 1	Round 2	Round 3	Finalists	Portfolio

- Use-case 1: Lightweight applications
- Use-case 2: High performance applications
  - `AEGIS-128`
  - `OCB`
- Use-case 3: Defense in depth

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 initialization and finalization
- 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

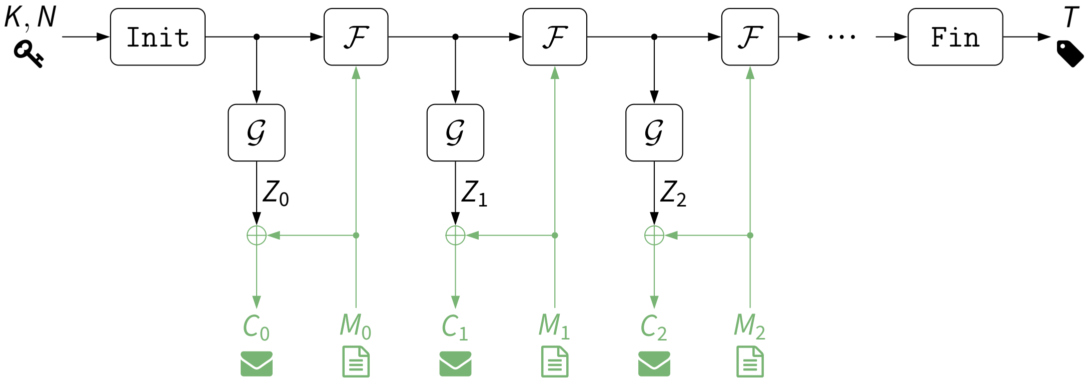

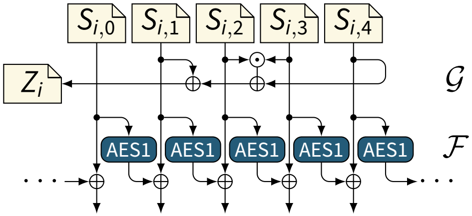

(a) AEGIS-**128**

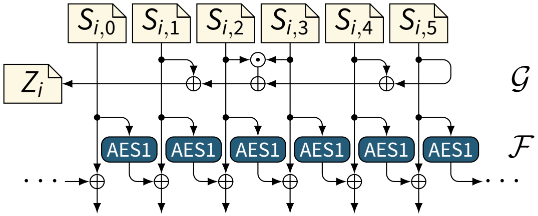

(b) AEGIS-**256**

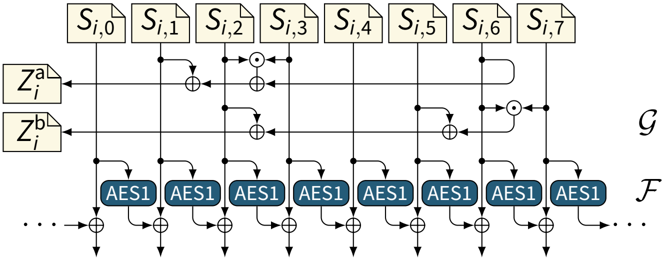

(c) AEGIS-**128L**

Linear keystream distinguisher [Min14]

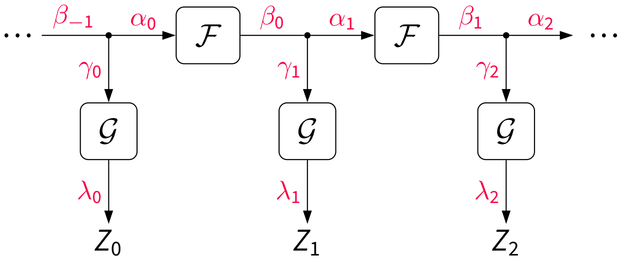

- Exploit keystream bias of $\lambda_0 \cdot Z_0 \oplus \lambda_1 \cdot Z_1 \oplus \lambda_2 \cdot Z_2$
- Correlation contribution $c = \prod_i (2p_i - 1) \to$ data complexity about c^{-2} KP

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

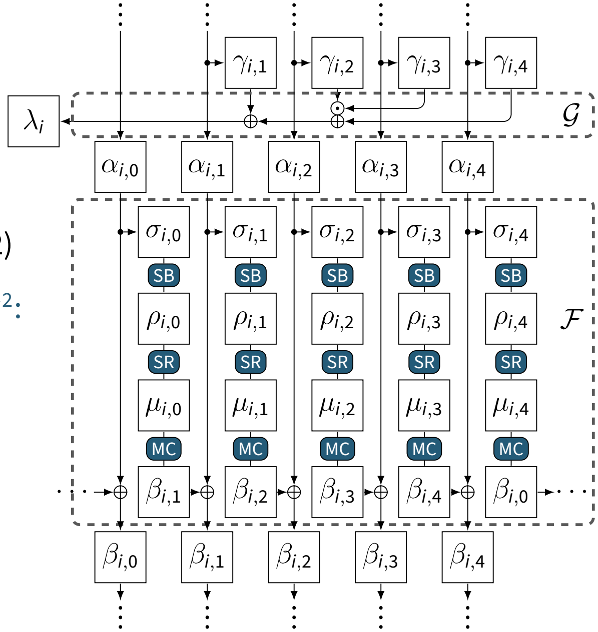

Simple Truncated Model – Results

- 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)

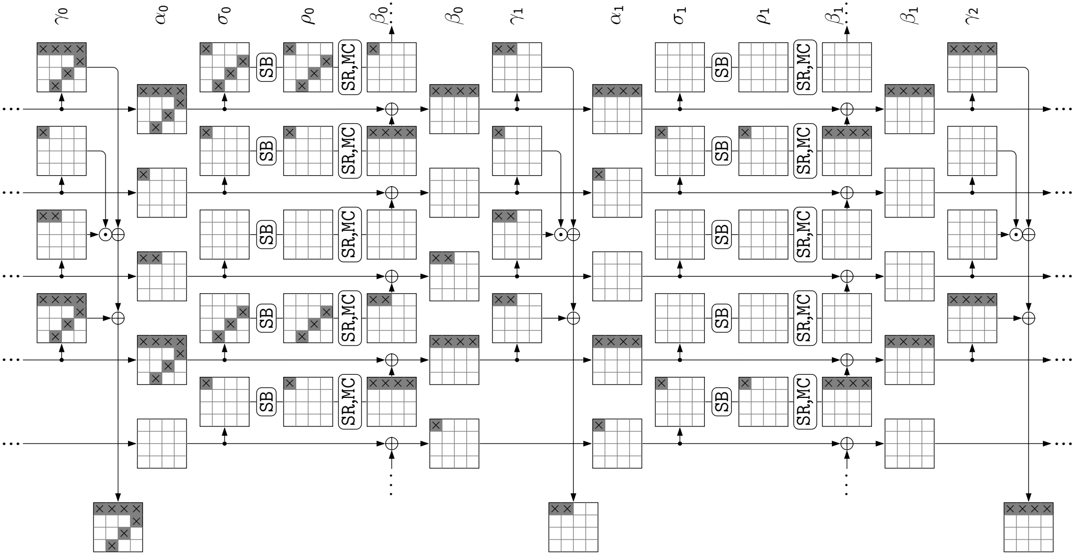

Inconsistent Trunc. Char. for AEGIS-128 (11 S-boxes, 13 ANDs)

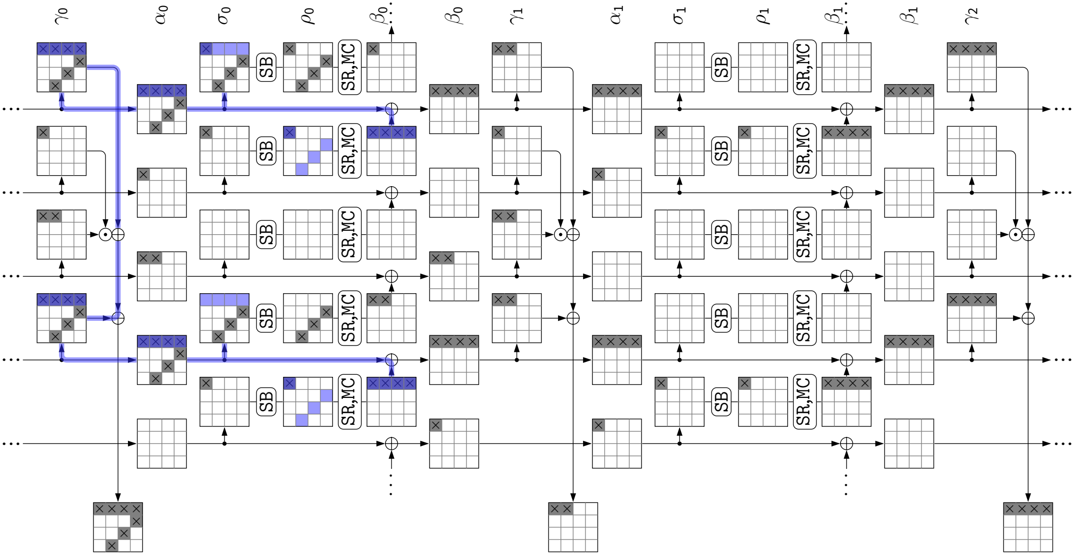

Inconsistent Trunc. Char. for AEGIS-128 (11 S-boxes, 13 ANDs)

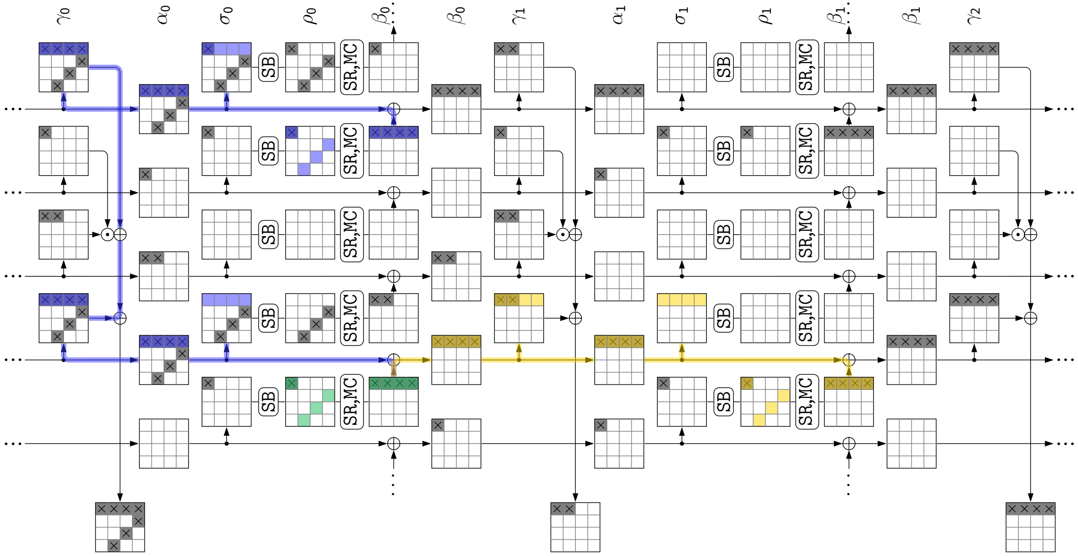

Improved Truncated Model

- Take any pair $(\mu_0, \beta_0), (\mu_1, \beta_1)$ of input/output masks for `MixColumns`. Consider their difference $(\mu_0 \oplus \mu_1, \beta_0 \oplus \beta_1)$.
- 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 $(\mu_0, \beta_0), (\mu_1, \beta_1)$ of input/output masks for `MixColumns`. Consider their difference $(\mu_0 \oplus \mu_1, \beta_0 \oplus \beta_1)$.
- 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 $(\mu_0, \beta_0), (\mu_1, \beta_1)$ of input/output masks for `MixColumns`. Consider their difference $(\mu_0 \oplus \mu_1, \beta_0 \oplus \beta_1)$.
- 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

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

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

Alternative Branch Model

First model

- $a + b + c = 2 \cdot d^{\vdash}, \qquad d^{\vdash} \in \{0, 1\}$

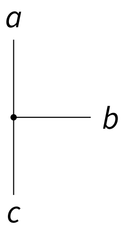

Alternative model

- $a + b \geq c, \quad a + c \geq b, \quad b + c \geq a,$
- $a + b + c \leq 2$

- $100\times$ speed-up through alternative branch model
  - 80 minutes instead of 5 days

Alternative Branch Model

First model

- $a + b + c = 2 \cdot d^{\vdash}, \qquad d^{\vdash} \in \{0, 1\}$

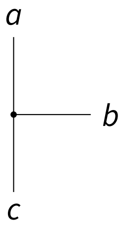

Alternative model

- $a + b \geq c, \quad a + c \geq b, \quad b + c \geq a,$
- $a + b + c \leq 2$

- $100\times$ speed-up through alternative branch model
  - 80 minutes instead of 5 days

Alternative Branch Model

First model

- $a + b + c = 2 \cdot d^{\vdash}, \qquad d^{\vdash} \in \{0, 1\}$

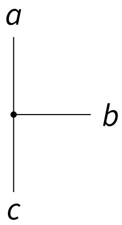

Alternative model

- $a + b \geq c, \quad a + c \geq b, \quad b + c \geq a,$
- $a + b + c \leq 2$

- $100\times$ speed-up through alternative branch model
  - 80 minutes instead of 5 days

Constraint Programming (CP) Model for Finding Characteristics

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

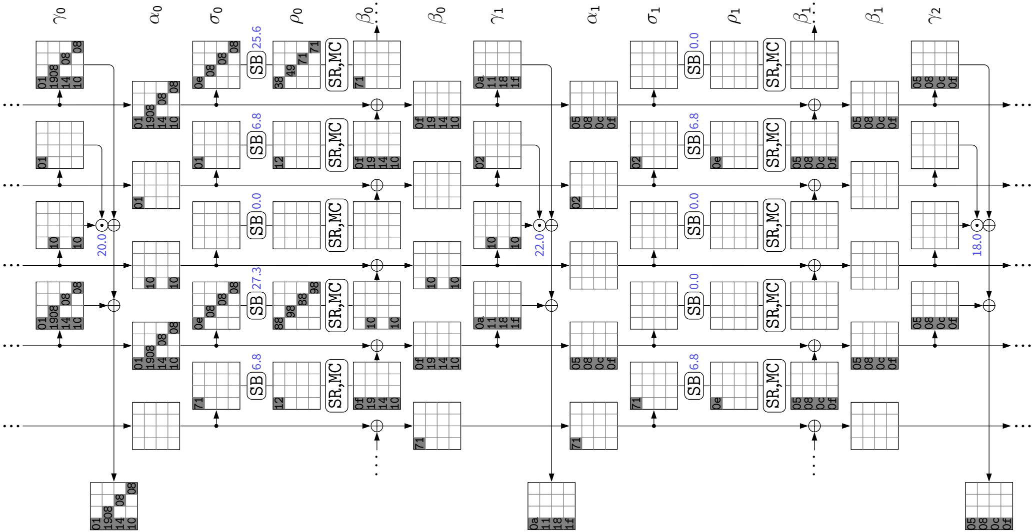

Linear approximation for AEGIS-128 with squared correlation contribution 2^{-140}

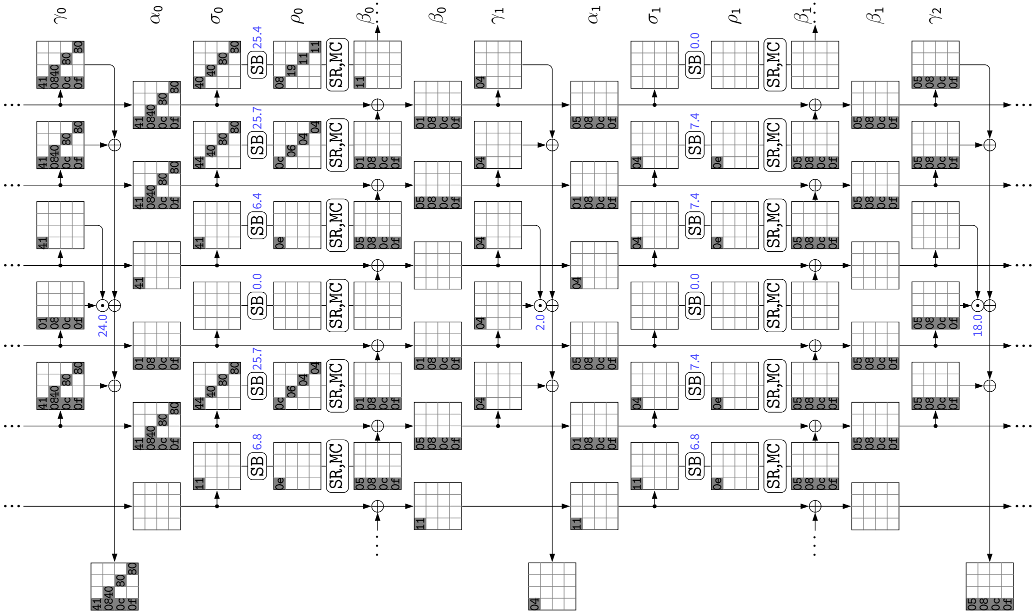

Linear approximation for AEGIS-256 with squared correlation contribution 2^{-162}

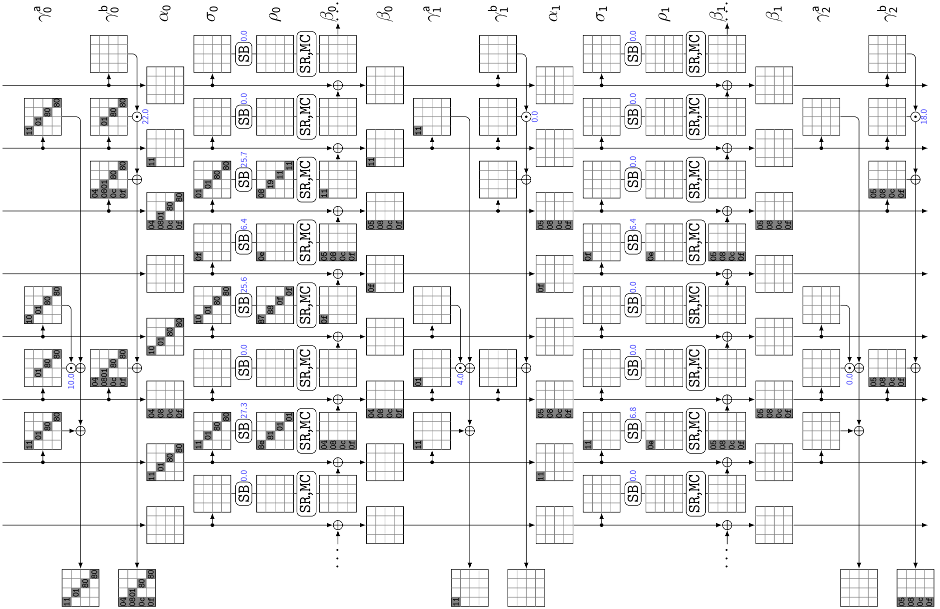

Linear approximation for AEGIS-128L with squared correlation contrib. 2^{-152}

Final Remarks

Experimental Verification

Note the dependencies between consecutive AND & SB:

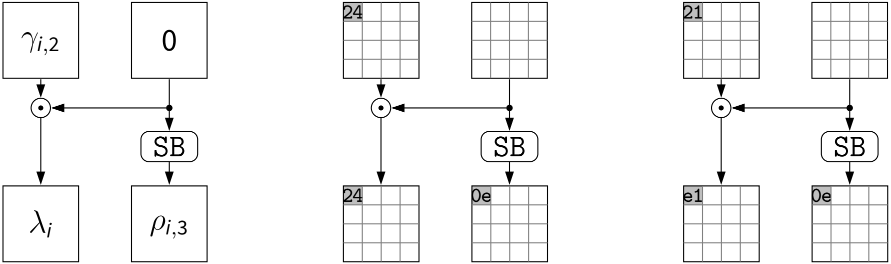

(a) $\beta_{i-1,3} \oplus \beta_{i,3} = 0$ (b) $2^{-7.66}$ instead of $2^{-4-6.39}$ (c) 0 instead of $2^{-8-6.39}$

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

Experimental Verification

Note the dependencies between consecutive AND & SB:

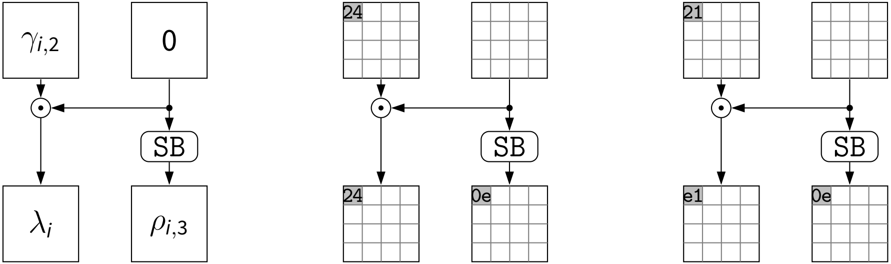

(a) $\beta_{i-1,3} \oplus \beta_{i,3} = 0$ (b) $2^{-7.66}$ instead of $2^{-4-6.39}$ (c) 0 instead of $2^{-8-6.39}$

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

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

SCIENCE
PASSION
TECHNOLOGY

Analyzing the Linear Keystream Biases in AEGIS

Maria Eichlseder **Marcel Nageler** Robert Primas

FSE 2020 - Athens

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.

[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.

[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.

[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.

Bibliography II

[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.

[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.