ESP IJAST
ESP International Journal of Advancements in Science & Technology
ISSN: 2583-9233/ Volume 4 Issue 4 October-December 2026
Paper Id: IJAST-V4I4P106/Page No: 40-51

Research Article

Truncated Authentication Tags for Ascon-AEAD128 over LoRa: A Risk Analysis of Forgery, Failure-Limit Denial of Service, Airtime and Packet Errors under NIST SP 800-232

Mohammed Khaire Hussain
Department of Medical Instruments Techniques, Mosul Polytechnic College, Northern Technical University, Mosul, Nineveh, Iraq.

Received Date: 01 October 2026 Revised Date: 15 October 2026 Accepted Date: 02 November 2026

Abstract: *LoRa devices send payloads of a few bytes, so every authentication-tag byte is a large share of each frame. NIST SP 800-232 standardises Ascon-AEAD128 for constrained devices, permits tags shorter than 64 bits only after a careful risk analysis, and bounds the verification failures tolerated per key. This study performs that analysis for LoRa uplinks. An analytical model links the tag length to time on air, forgery probability against a channel-time-limited attacker, the attacker airtime needed to exhaust the failure budget and force a key update, benign budget consumption by corrupted frames that pass the 16-bit LoRa CRC, and packet errors. It is evaluated for SF7–SF12 with a bit-exact LoRa physical-layer simulator verified against 240 NIST test vectors and the exact symbol-error rate. A 64-bit tag shortens time on air by 4.8–18.7% relative to a 128-bit tag, and about 61% of 8-bit tag increments cost no airtime. Longer frames cost at most 0.23 dB in AWGN. Under the strict failure limit, tags shorter than 64 bits let one forged frame force a key update. The minimum admissible tag is therefore 64 bits under the strict limit and 48–56 bits under the relaxed limit.*

Keywords: *Ascon, Authenticated Encryption, Internet of Things Security, Lora, Message Authentication Tag Length, Physical-Layer Simulation.*

I. INTRODUCTION

Low-power wide-area networks (LPWANs) connect battery-powered sensors over kilometre ranges at very low data rates. LoRa, with the LoRaWAN medium-access layer, is one of the most widely deployed LPWAN technologies. It achieves its range through chirp spread spectrum with spreading factors SF7–SF12. The price is that time on air (ToA) roughly doubles with every increase in SF. The European 868 MHz band additionally limits each device to a 1% duty cycle. Consequently, every byte added to a frame costs energy, channel occupancy, and allowed transmissions per hour [1], [2].

Authentication is one of the bytes that cannot be removed. LoRaWAN appends a 4-byte message integrity code (MIC), computed with AES-CMAC and truncated to 32 bits, to every data frame [6]. Applications that need end-to-end confidentiality and integrity beyond the network server usually add an authenticated-encryption layer [6], [12]. With a full 128-bit tag, that layer adds 16 bytes to payloads that are often only 4–24 bytes long.

The August 2025 NIST lightweight-cryptography standard, SP 800-232 [16], makes this trade-off explicit. It standardises Ascon-AEAD128 [17], selected after a multi-round public evaluation [19], a permutation-based authenticated-encryption scheme designed for constrained devices. It also offers tag truncation as an implementation option, but attaches two requirements:

- R4: truncated tags must be at least 32 bits long, and lengths below 64 bits may be chosen only after a careful risk analysis;
- R5: the number of decryption (verification) failures under one key is bounded. The bound is $2^{\lambda-32}$ for $64 \le \lambda \le 128$. For shorter tags it should be a single failure, relaxable up to $2^{\lambda-32}$ if a risk analysis shows the security goal is still met.

The standard also requires the key to be updated when the failure limit is reached (R7) [16].

The standard does not say what such a risk analysis looks like for any particular link. For LoRa the question is not only how likely a forgery is. Three LoRa-specific effects interact with the tag length:

- Symbol granularity. ToA grows in blocks of symbols, not bytes, so some tag bytes are free and others are expensive.
- Budget exhaustion by attackers. The R5 failure counter can be exhausted by an adversary. Under the strict limit, a single forged frame forces a key update, which on LoRaWAN means a new join procedure.
- Budget exhaustion by benign errors. The 16-bit LoRa CRC occasionally accepts corrupted frames. Such frames fail tag verification and consume the same budget even when no attacker is present.


This is an open access article under the CC BY-NC-ND license (https://creativecommons.org/licenses/by-nc-nd/2.0/)

To the best of our knowledge, none of these interactions has been quantified. Existing LoRa security studies treat the MIC or tag as a fixed overhead [9]–[13]. Studies of short tags address other links or modes: GCM [20]–[22], Sigfox [23], and generic messages [25].

This paper provides the missing risk analysis. Its contributions are:

- An airtime cost model of the tag length for SF7–SF12 and payloads of 4–51 bytes. It shows that about 61% of 8-bit tag increments are free because of symbol-block granularity, and that a 64-bit tag saves 4.8–18.7% ToA relative to a 128-bit tag.
- A threat and failure model aligned with SP 800-232. It jointly quantifies forgery probability under a channel-time-limited attacker, the attacker airtime needed to force a key update through R5 budget exhaustion, and the benign consumption of that budget by CRC-passing corrupted frames.
- A bit-exact LoRa physical-layer (PHY) Monte Carlo simulator, verified against NIST test vectors, the Semtech ToA formula and the exact symbol-error rate. It measures the packet-error penalty of longer tags and the undetected-error probability of the payload CRC.
- A design rule and design table giving the minimum admissible tag length under strict and relaxed interpretations of R5, with a sensitivity analysis over ten scenario variants.

Section II reviews related work. Section III presents the system and threat model. Section IV describes the simulation framework and its verification. Section V reports and discusses the results. Section VI states the limitations, and Section VII concludes.

II. RELATED WORK

LoRa performance. Adelantado et al. analysed the capacity and duty-cycle limits of LoRaWAN [1]. Bor et al. quantified collision-driven loss in dense deployments [2]. Vangelista formalised LoRa as frequency-shift chirp modulation and established the orthogonality of its symbols [3]. Elshabrawy and Robert derived a closed-form approximation of its bit-error rate [4]. Tapparel et al. released an open-source implementation of the full LoRa PHY chain in GNU Radio [5]. The ToA formula used throughout this paper is Semtech's [8], and frame formats follow the LoRaWAN L2 1.0.4 specification [6] and regional parameters [7]. None of these works considers the security overhead.

LoRa security and overhead. Yang et al. demonstrated practical attacks on LoRaWAN key and counter management [9]. Butun et al. performed a threat-oriented risk analysis of LoRaWAN 1.1 [10]. Several recent works measure the cost of cryptography on LoRa-class devices:

- Thaenkaew et al. evaluated AES-192/256 in LoRaWAN and found the time and energy increase negligible relative to ToA [11];
- Simin et al. compared AES-256-GCM, ASCON-128, SPECK, XTEA and a commercial scheme on a LoRaWAN testbed [12];
- Shkuratskyy et al. estimated the airtime cost of lightweight and post-quantum encryption for compressed images over LoRa [13];
- Amirkhanova et al. built an AES-GCM pipeline on ESP32 nodes [14];
- Atul et al. proposed a reproducible benchmark protocol for Ascon-AEAD128, AES-GCM and ChaCha20-Poly1305 on IoT nodes [15].

In all of these works the tag length is fixed (32-bit MIC or 128-bit tag). The variable we study is not considered.

Short tags and truncated MACs. Truncation of message authentication codes has long been known to enable attacks when the truncated length is small [22]. For GCM, Ferguson showed that 32-bit tags leak information about the authentication key after successful forgeries [21]. NIST SP 800-38D therefore allows 32- and 64-bit GCM tags only under strict usage limits [20]. Ferreira analysed Sigfox, whose uplink MAC is 16–40 bits, and demonstrated forgeries [23]. Andrushchak and Koloshko proposed an algorithmic selection of MAC length that balances forgery risk against delay and throughput for short messages. Their model has no PHY, packet-error or LPWAN dimension [25]. Wagner et al. studied MAC aggregation over lossy channels as an alternative way to reduce per-packet overhead [24]. For Ascon, Lefevre and Mennink surveyed the security of the Ascon modes, including truncated tags, in single- and multi-user settings [18]. An IETF draft defines 128-, 64- and 32-bit Ascon-AEAD128 variants for COSE and JOSE [26]. Daidone et al. analysed the performance impact of the IEEE 802.15.4 security sub-layer [28].

Undetected channel errors. The probability that a cyclic redundancy check accepts a corrupted frame depends on the polynomial and the error pattern. It approaches 2^{-r} for an r-bit CRC under random errors [27]. We found no work that connects these undetected errors to AEAD failure counters or key-update triggers.

Gap. No prior work performs the SP 800-232 R4/R5 risk analysis for an LPWAN link. In particular, none jointly quantifies how the tag length determines LoRa airtime, forgery probability, attacker cost to exhaust the failure budget, benign budget consumption, and packet error rate. Table 1 summarises the positioning.

Table 1: Positioning of This Work with Respect to the Closest Related Studies.

Work	Link	Tag length varied	Forgery model	Failure budget / key update	Airtime cost	PHY packet errors
[11], [12], [14]	LoRaWAN / ESP32	No	No	No	Partly	No
[13]	LoRa	No	No	No	Yes	No
[20], [21]	GCM (generic)	32/64/≥96	Yes	Usage limits	No	No
[23]	Sigfox	Fixed per mode	Yes (attack)	No	No	No
[25]	Generic	Yes	Yes	No	Throughput	No
[16], [18], [26]	Ascon (generic)	Yes	Bounds	R5 (stated only)	No	No
This work	LoRa	32–128 bits	Yes	Yes (strict and relaxed R5)	Yes	Yes (bit-exact)

III. SYSTEM AND THREAT MODEL

A. Frame Format

The system is a LoRa uplink from an end device to a network server (Figure 1). The frame keeps the LoRaWAN data-frame header [6]:

- MHDR (1 byte);
- DevAddr (4);
- FCtrl (1);
- FCnt (2);
- FPort (1).

The application payload of P bytes is encrypted with Ascon-AEAD128 [16], and the 4-byte MIC is replaced by the Ascon tag truncated to λ bits, λ ∈ {32, 40, ..., 128}. The 128-bit nonce is not transmitted. It is derived from the device address, the 32-bit frame counter and the direction, so the frame counter guarantees nonce uniqueness (R3 of [16]) as long as it never repeats under one key. The PHY payload length is therefore

$$B(P,\lambda) = 9 + P + \frac{\lambda}{8} \quad \text{bytes.} \tag{1}$$

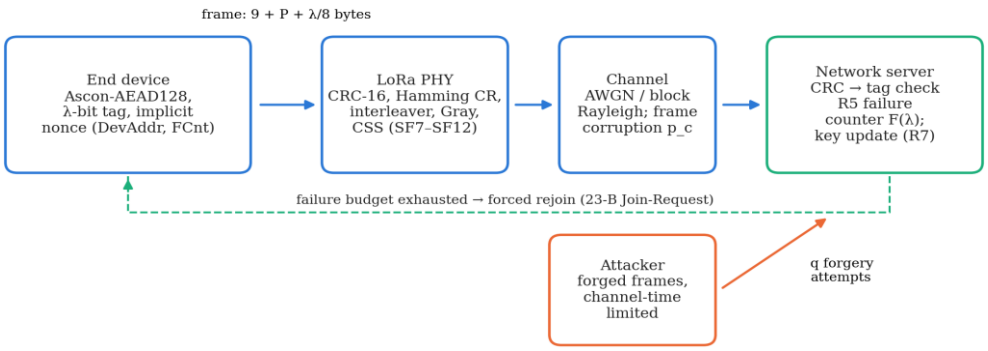


Figure 1: System And Threat Model. The Network Server Checks the Lora CRC, Then the Truncated Ascon Tag. It Counts Verification Failures Against the SP 800-232 R5 Budget F(Λ) And Forces a Key Update (A New Join Procedure) When the Budget Is Exhausted.

B. Time on Air

With bandwidth BW, symbol duration $T_s = 2^{SF}/BW$ and an n_p-symbol preamble, the ToA of a B-byte PHY payload is [8]

$$T_{oa}(B) = \left(n_p + 4.25 + N_{pl}\right) T_s, \tag{2}$$

$$N_{\mathrm{pl}} = 8 + \max\left(\left\lceil \frac{8B - 4SF + 28 + 16\mathrm{CRC} - 20\mathrm{IH}}{4(SF - 2\mathrm{DE})} \right\rceil (CR+4),\ 0 \right), \tag{3}$$

where:

- CR ∈ {1,...,4} encodes the coding rates 4/5–4/8;
- CRC = 1 when the payload CRC is enabled;
- IH = 1 for implicit-header mode;
- DE = 1 when low-data-rate optimisation is enabled, which is required when T_s exceeds 16 ms.

The ceiling in (3) makes ToA a step function of B. A tag increment either leaves ToA unchanged or adds a whole block of CR + 4 symbols. At a fixed transmit power, the radio energy per frame is proportional to T_{oa}. We therefore report ToA and treat it as the energy proxy. Under a duty cycle d, a device can send at most 3600 d/T_{oa} frames per hour.

C. Threat Model and Failure Budget

The attacker can transmit arbitrary frames with the victim's header fields and any frame-counter value. It aims either to have a forged frame accepted or to force the victim to update its key. The attacker is assumed **not** to respect duty-cycle regulation. It is limited only by channel time on n_{ch} channels. In an attack of duration D it can therefore attempt

$$q = \left\lfloor \frac{D n_{\mathrm{ch}}}{T_{\mathrm{oa}}(B)} \right\rfloor \tag{4}$$

forgeries. For a λ-bit Ascon tag, each attempt succeeds with probability $2^{-\lambda}$ [16, Sec. 4.4.1], so the probability of at least one successful forgery is

$$P_f = 1 - \left(1 - 2^{-\lambda}\right)^q. \tag{5}$$

The server maintains a per-key failure counter with the budget of R5 [16]:

$$F(\lambda) = \begin{cases} 2^{\lambda-32}, & 64 \le \lambda \le 128, \\ 1, & 32 \le \lambda < 64 \text{ (strict)}, \\ 2^{\lambda-32}, & 32 \le \lambda < 64 \text{ (relaxed, after risk analysis)}. \end{cases} \tag{6}$$

When the counter reaches F(λ), the key is updated (R7). On LoRaWAN this means a new over-the-air activation. Its uplink is a 23-byte Join-Request [6], which we use as the victim-side cost. With the relaxed budget, the per-key success probability is at most F·2−λ = 2−32 for every λ.

Two metrics describe the denial-of-service (DoS) exposure. The first is the attacker airtime needed to force one key update,

$$A(\lambda) = \frac{F(\lambda)\, T_{\mathrm{oa}}(B)}{n_{\mathrm{ch}}}, \tag{7}$$

and the second is the DoS amplification $T_{\mathrm{oa}}(23)/(F\ T_{\mathrm{oa}}(B))$, i.e., the ratio of victim airtime to attacker airtime per forced update.

D. Benign Budget Consumption

A frame corrupted by noise or a collision is normally discarded by the 16-bit payload CRC. With probability p_{ud}, however, the CRC accepts it. The tag check then fails with probability $1 - 2^{-\lambda} \approx 1$, and the failure is counted exactly like a forgery attempt. For a device sending N frames over its lifetime, a fraction p_c of which are corrupted, the expected number of benign failures is

$$\mathbb{E}[N_b] = N\, p_c\, p_{ud}, \qquad p_{ud} \approx 2^{-16}, \tag{8}$$

and the expected number of benign key updates is $\mathbb{E}[N_b]/F(\lambda)$. The approximation of p_{ud} is checked by simulation in Section V-C.

E. Design Rule

Given a target forgery probability ε, a minimum attacker cost $D_{\min}$ per forced key update, and a tolerance of 10^{-3} benign key updates per lifetime, the minimum admissible tag length is

$$\lambda_{\min} = \min\left\{\lambda \in \{32,40,\ldots,128\} :\ P_f \le \varepsilon,\ A(\lambda) \ge D_{\min},\ \frac{\mathbb{E}[N_b]}{F(\lambda)} \le 10^{-3}\right\}. \tag{9}$$

IV. MATERIALS AND METHODS

A. Simulation Framework

All models were implemented in Python 3.11 with NumPy 2.4.4, SciPy 1.17.1 and pandas. The framework has three modules:

- a pure-Python Ascon-AEAD128 implementation following SP 800-232, including truncation and nonce masking;
- the ToA model of (2)–(3);
- a bit-exact LoRa PHY packet simulator.

The PHY simulator performs the following steps:

- Appends a CRC-16 (polynomial 0x1021) to a random payload.
- Splits the frame into nibbles.
- Encodes the nibbles with a Hamming-type code of rate 4/(4+CR). Rates 4/5 and 4/6 detect errors only; rates 4/7 and 4/8 correct single errors by minimum-distance decoding.
- Applies diagonal interleaving over blocks of SF (or SF−2) codewords.
- Gray-maps the interleaved bits to chirp symbols.

The explicit header occupies the first eight symbols with rate 4/8 and SF−2 bits per symbol. Low-data-rate optimisation is applied for SF11 and SF12. Whitening is omitted because it does not alter error statistics for random payloads.

Because LoRa chirps are orthogonal [3], a dechirp-and-FFT receiver sees a correct FFT bin √Es h + n0 and 2SF − 1 noise-only bins, with n ∼ $\mathcal{CN}$(0, 1). Here Es = 2SF·SNR and the SNR is measured over the channel bandwidth. The maximum noise power over the wrong bins is sampled exactly by inverse transform: if U ∼ $\mathcal{U}$(0, 1), then −ln[1 − U1/(M−1)] has the distribution of that maximum. A symbol error occurs when this maximum exceeds the correct-bin power, and the erroneous symbol is then uniform over the other bins. Block Rayleigh fading draws one |h|2 ∼ Exp(1) per packet.

A packet is counted as correctly delivered only if the header, the payload and the CRC all decode correctly. It is counted as an undetected error if the header decodes and the CRC passes but the payload differs from the one transmitted.

B. Verification

Before any experiment, the framework was verified in five ways:

- The Ascon implementation was run against all 240 NIST ACVP vectors for Ascon-AEAD128 [29]. These vectors cover bit-length inputs, truncated tags and nonce masking.
- 500 random byte-level encryption/decryption round trips with truncated tags were checked, each followed by a one-bit tag modification.
- The ToA model was compared with published reference values.
- The number of symbols produced by the PHY encoder was checked against (3) for all SF, CR and payload combinations.
- The Monte Carlo symbol-error rate (SER) was compared with the exact SER of noncoherent orthogonal M-ary detection in AWGN [30], computed by numerical integration:

$$P_s = 1 - \int_0^\infty f_X(x)\left(1 - e^{-x}\right)^{M-1} dx, \quad f_X(x) = e^{-(x+E_s)} I_0\left(2\sqrt{E_s x}\right), \tag{10}$$

where $M = 2^{SF}$ and I_0 is the modified Bessel function of the first kind. The results of these checks are reported in Section V-A.

C. Experimental Design

Table 2 lists the parameters. Four experiment groups were run:

- E1–E2 evaluate (1)–(3) over the full grid of SF, payload and tag length.
- E3 estimates the packet error rate (PER) for P = 12 bytes and λ ∈ {32, 64, 128}, i.e., frames of 25, 29 and 37 bytes, at SF7, SF10 and SF12 in AWGN and block Rayleigh fading. Each point ran until at least 200 packet errors or 40,000 packets, and 95% Clopper–Pearson intervals [31] were computed.
- E3b estimates pud with 2×10^6 packets per point. It uses the CRC-16 and, as a check of the 2−r scaling, an 8-bit CRC surrogate (polynomial 0x07).
- E4–E5 evaluate (4)–(9), and E7 varies the scenario parameters one at a time.

All random numbers derive from the base seed 20261001. The complete run takes about 150 s on one CPU core, and every table and figure is regenerated from the saved raw outputs.

Table 2: Simulation and Scenario Parameters.

Parameter	Value(s)
Bandwidth, coding rate, preamble	125 kHz, CR 4/5, 8 symbols
Header, payload CRC, LDRO	explicit, CRC-16 on, LDRO for SF11–SF12
Spreading factor	SF7–SF12
Application payload P	4, 12, 24, 51 bytes
Tag length λ	0 (none), 32, 40, 48, 56, 64, 80, 96, 128 bits
Channel	AWGN; block Rayleigh (one fade per packet)
PER stopping rule	≥ 200 packet errors or 40,000 packets per point
Device lifetime; report period	10 years; 600 s (525,960 frames)
Attack duration D; channels n_ch	1 year; 1
Forgery target ε; minimum attacker cost D_min	10^{-6}; 24 h
Corrupted-frame fraction p_c	0.1
Benign key-update tolerance	≤ 10^{-3} per lifetime
Random seed	20261001

V. RESULTS AND DISCUSSION

A. Verification

Table 3 summarises the verification. The Ascon implementation reproduced all 120 encryption vectors and all 60 valid decryption vectors, and it rejected all 60 vectors with a modified tag. The ToA model gives 61.70 ms (SF7) and 1482.75 ms (SF12) for a 23-byte PHY payload, matching the values produced by AN1200.13-based calculators. The PHY encoder emitted exactly the number of symbols predicted by (3) in all 1,320 configurations tested. The Monte Carlo SER lay within its 95% confidence interval of the exact value (10) at all nine SNR points tested. The simulator can therefore be trusted in both its framing and its detection statistics.

Table 3: Verification Results

Test	Result
NIST ACVP Ascon-AEAD128 vectors (240)	encrypt 120/120; valid decrypt 60/60; modified tag rejected 60/60
Random round trips with truncated tags; 1-bit tag tamper	500/500 recovered; 500/500 rejected
ToA of a 23-byte PHY payload (SF7 / SF12)	61.70 ms / 1482.75 ms
Encoder symbol count vs (3), SF7–12 × CR 4/5–4/8 × 5–59 B	1,320 configurations, 0 mismatches
Noise-free round trip of the PHY chain	72 configurations, 0 packet errors
Monte Carlo SER vs exact SER (10), SF7/SF9/SF12	exact value inside the 95% CI at 9/9 points

B. Airtime Cost of the Tag

Figure 2 and Table 4 show ToA as a function of tag length for a 12-byte payload. Four observations follow.

- ToA is a staircase, not a line. At SF10, tags of 32, 40, 48, 56 and 64 bits all give 411.6 ms. At SF7, 48–64 bits all give 66.8 ms. Across all SF and payload combinations, 60.7% of 8-bit increments above 32 bits leave ToA unchanged. The reason is (3): the payload is carried in blocks of SF (or SF−2) codewords, each block costs CR + 4 symbols, and a tag increment only costs airtime when it crosses a block boundary.
- The full tag is expensive for short payloads. A 128-bit tag costs 1.10–1.37 times the airtime of a 32-bit tag, and the ratio is largest for the shortest payloads. Relative to 128 bits, a 64-bit tag saves 4.8–18.7% of ToA and radio energy.
- The same costs apply to throughput. Under the 1% EU868 duty cycle, the maximum goodput of a 12-byte payload at SF7 falls from 15.6 bit/s with λ = 32 to 11.7 bit/s with λ = 128.
- The consequences at high SF are larger. At SF12 the difference between 32 and 128 bits is almost 0.5 s per frame, which is significant for battery life at slow data rates.

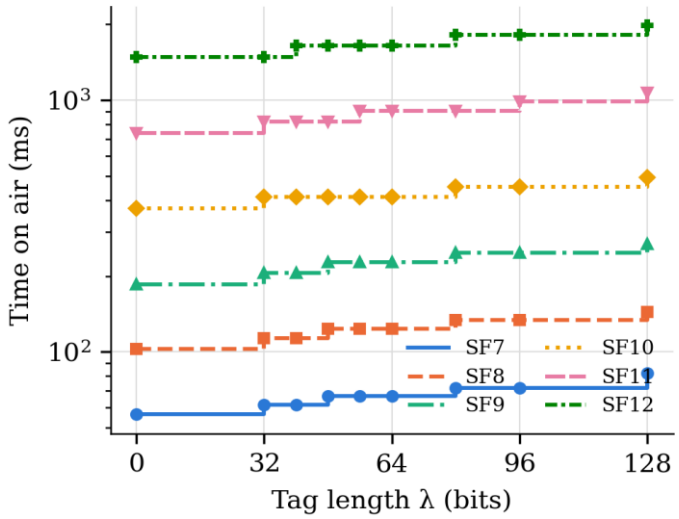


Figure 2: Time On Air Versus Tag Length for A 12-Byte Payload (CR 4/5, 125 Khz, Explicit Header, CRC On). Steps Are Caused by The Symbol-Block Granularity Of (3).

Table 4: Time on air (ms) for a 12-byte payload.

λ (bits)	SF7	SF8	SF9	SF10	SF11	SF12
32	61.7	113.2	205.8	411.6	823.3	1482.8
48	66.8	123.4	226.3	411.6	823.3	1646.6
64	66.8	123.4	226.3	411.6	905.2	1646.6
96	71.9	133.6	246.8	452.6	987.1	1810.4
128	82.2	143.9	267.3	493.6	1069.1	1974.3

C. Packet-Error Penalty and Undetected CRC Errors

Figure 3 shows the PER in AWGN. Longer frames contain more symbols that can fail, so their PER is higher at a given SNR. The effect is small, however. At PER = 10%, the SNR requirement increases from λ = 32 to λ = 128 by:

- 0.20 dB at SF7;
- 0.15 dB at SF10;
- 0.23 dB at SF12.

From λ = 32 to λ = 64 the increase is only 0.03–0.09 dB. Under block Rayleigh fading (Figure 4), the 95% confidence intervals of the three tag lengths overlap at 20 of 21 SNR points. The one exception (SF10, −7 dB) is not ordered by λ. When the entire frame experiences one fade, packet loss is governed by the probability of a deep fade (outage), which does not depend on frame length. In practice, the reliability cost of a longer tag is therefore dominated by its airtime cost, not by its error-rate cost.

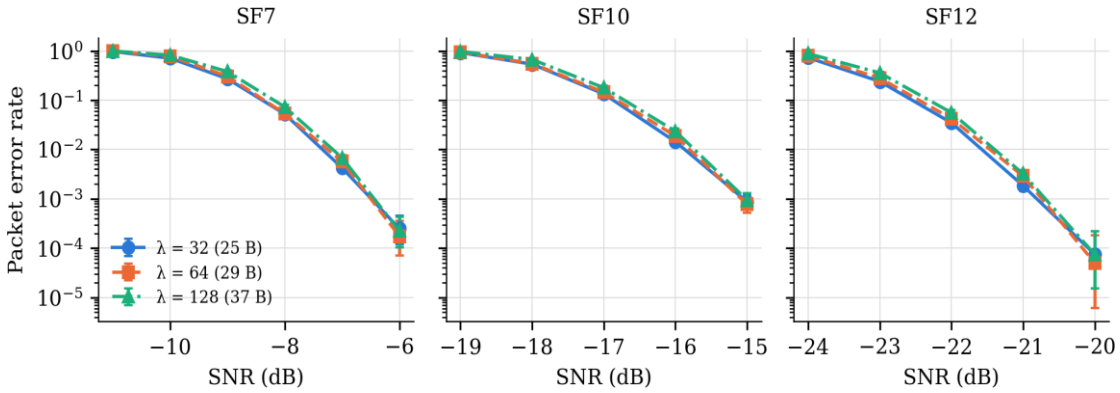


Figure 3: Packet Error Rate in AWGN For A 12-Byte Payload With Λ = 32, 64 and 128 Bits (Frames Of 25, 29 And 37 Bytes). Error Bars Are 95% Clopper–Pearson Intervals.

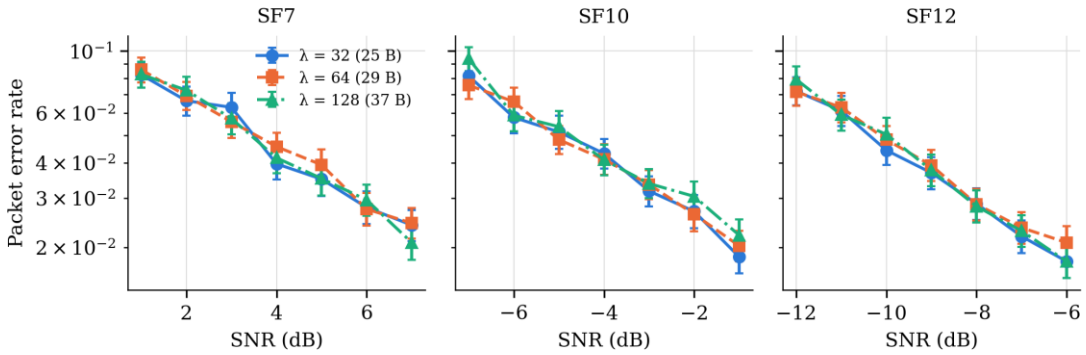


Figure 4: Packet Error Rate Under Block Rayleigh Fading; Parameters as in Figure 3.

Table 5 and Figure 5 report the probability that a corrupted payload passes the CRC. For the 8-bit surrogate, the measured values (3.77×10^{-3} and 3.93×10^{-3}) are close to $2^{-8} = 3.91 \times 10^{-3}$, which confirms the 2^{-r} scaling assumed in (8). For the 16-bit CRC at −14 dB, the measurement agrees with 2^{-16}. At −12 dB, however, the measured rate of 2.47×10^{-5} is about 1.6 times higher than 2^{-16}, and its confidence interval excludes 2^{-16}. At that SNR most corrupted frames carry only a few symbol errors, and such structured error patterns are not equivalent to random bit errors. Equation (8) should therefore be read as an order-of-magnitude model that can underestimate p_{ud} by a factor of about 2. Because p_c is varied by a factor of 30 in the sensitivity study (Section V-F), this uncertainty does not change any conclusion.

Table 5: Probability That a Corrupted Payload Passes the CRC (SF7, 21-Byte Payload, 2×10^6 Packets Per Point).

CRC	SNR (dB)	Corrupted frames with valid header	Undetected	P(CRC pass | corrupted)	95% CI	2^{-r}
8-bit surrogate	−12	1,213,715	4,580	3.77×10^{-3}	$[3.67, 3.88]\times10^{-3}$	3.91×10^{-3}
8-bit surrogate	−14	239,663	943	3.93×10^{-3}	$[3.69, 4.19]\times10^{-3}$	3.91×10^{-3}
CRC-16	−12	1,215,060	30	2.47×10^{-5}	$[1.67, 3.53]\times10^{-5}$	1.53×10^{-5}
CRC-16	−14	239,758	2	0.83×10^{-5}	$[0.10, 3.01]\times10^{-5}$	1.53×10^{-5}

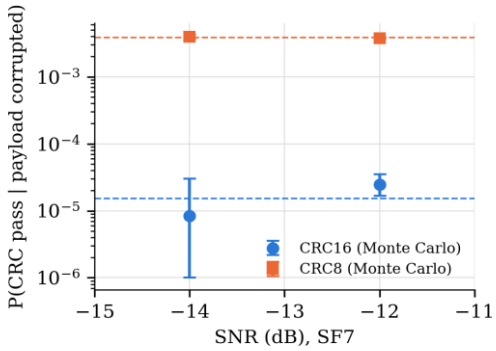


Figure 5: Monte Carlo Probability That a Corrupted Payload Passes The CRC, Compared With 2^{-R} (Dashed). Error Bars Are 95% Clopper–Pearson Intervals.

D. Forgery Probability

Figure 6 plots (5) for a one-year attack on a single channel with P = 12 bytes. Because the attacker is limited by channel time, a slower SF gives it fewer attempts:

- 5.1×10^8 attempts at SF7;
- 2.1×10^7 attempts at SF12.

A 32-bit tag without a failure limit (as used by the LoRaWAN MIC) yields a forgery probability of 0.112 at SF7, 0.018 at SF10 and 0.0049 at SF12. A 32-bit tag is therefore unacceptable against a determined attacker unless failures are limited.

With λ = 48, the forgery probability drops to 1.7×10^{-6} (SF7) and 6.8×10^{-8} (SF12). Each additional 8 bits reduces it by a factor of 256.

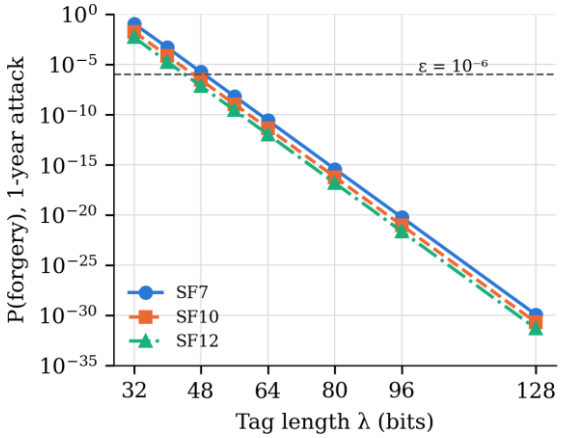


Figure 6: Probability of at Least One Successful Forgery During a One-Year, Single-Channel Attack (P = 12 Bytes, No Failure Limit). The Dashed Line Marks E = 10^{-6}.

E. Failure Budget: Denial of Service and Benign Exhaustion

Considering forgery probability alone would suggest that a 48-bit tag is sufficient. The R5 failure budget changes this conclusion (Figure 7, Table 6).

Strict R5 (F = 1 for λ < 64). Every forged frame forces a key update.

- The attacker spends only one frame of airtime per forced update: 0.06–0.07 s at SF7 and 1.5–1.6 s at SF12.
- The victim then spends at least the airtime of a 23-byte Join-Request. The DoS amplification is 0.52–1.25 across all SF and payload cases, so the attack costs the attacker about as much airtime as it costs the victim.
- Such an attack is cheap, repeatable and hard to distinguish from interference.
- Even without an attacker, a device reporting every 10 minutes for 10 years, with 10% of its frames corrupted, is expected to experience 0.80 CRC-passing corrupted frames. Under the strict limit each of them forces a key update, so the probability of at least one benign forced update is about 55%.

64-bit tags. For λ = 64, the budget 232 raises the attacker's cost per forced update to 2.9×10^{8} s (about 9 years) at SF7 and 7.1×10^{9} s at SF12. Benign exhaustion becomes negligible (1.9×10^{-10} expected updates per lifetime).

Relaxed R5. The relaxed budget makes intermediate lengths viable:

- λ = 48 requires 4.4×10^{3} s (73 min) of continuous attacker airtime per forced update at SF7 and 1.1×10^{5} s (30 h) at SF12;
- λ = 56 requires 1.1×10^{6} s (13 days) at SF7 and 2.8×10^{7} s (320 days) at SF12;
- the per-key forgery probability is capped at 2−32.

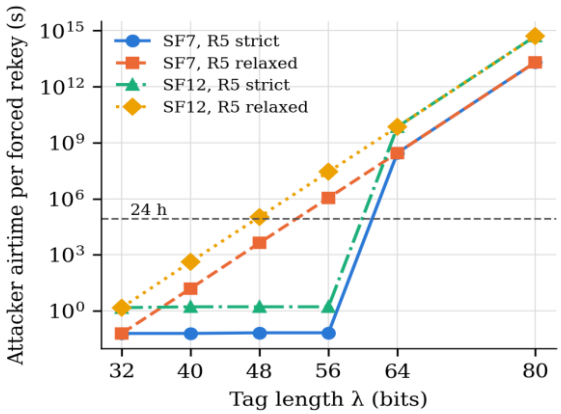


Figure 7: Continuous Attacker Airtime Needed to Force One Key Update, For Strict and Relaxed R5 (P = 12 Bytes). The Dashed Line Marks D_Min = 24 H.

Table 6: Security Quantities for A 12-Byte Payload (One-Year Single-Channel Attack; 10-Year Lifetime, 600 S Reports, P_C = 0.1).

SF	λ	ToA (ms)	P_f (1 year, no limit)	Attacker airtime per forced update, strict R5 (s)	Same, relaxed R5 (s)	Benign updates per lifetime, strict R5
7	32	61.7	1.1×10^{-1}	0.062	0.062	0.80
7	48	66.8	1.7×10^{-6}	0.067	4.4×10^{3}	0.80
7	56	66.8	6.6×10^{-9}	0.067	1.1×10^{6}	0.80
7	64	66.8	2.6×10^{-11}	2.9×10^{8}	2.9×10^{8}	1.9×10^{-10}
7	128	82.2	1.1×10^{-30}	6.5×10^{27}	6.5×10^{27}	≈0
12	32	1482.8	4.9×10^{-3}	1.48	1.48	0.80
12	48	1646.6	6.8×10^{-8}	1.65	1.1×10^{5}	0.80
12	56	1646.6	2.7×10^{-10}	1.65	2.8×10^{7}	0.80
12	64	1646.6	1.0×10^{-12}	7.1×10^{9}	7.1×10^{9}	1.9×10^{-10}
12	128	1974.3	4.7×10^{-32}	1.6×10^{29}	1.6×10^{29}	≈0

F. Design Rule, Design Table and Sensitivity

Applying (9) with the baseline parameters gives the design table in Table 7.

- Strict R5. The minimum admissible tag is 64 bits for all 24 combinations of SF and payload. The binding constraint is not forgery probability. It is the DoS and benign-exhaustion exposure created by a budget of one failure. The 64-bit tag saves 4.8–18.7% of ToA relative to 128 bits (mean 12.2%). Because of the staircase effect, in several cases it costs no more airtime than a 48-bit tag, or even a 32-bit tag (e.g., SF10 with P = 12 bytes).
- Relaxed R5. The minimum tag becomes 56 bits for SF7–SF10 and 48 bits at SF12. At SF12 each forged frame occupies the channel for 1.6 s, so the attacker needs more airtime to exhaust a given budget. The corresponding ToA savings relative to 128 bits are 7.1–19.9% (mean 13.6%).

Table 7: Minimum Admissible Tag Length Λ_Min (Bits), Baseline Scenario. Under Strict R5, Λ_Min = 64 Bits for Every Entry.

SF	P = 4 B	P = 12 B	P = 24 B	P = 51 B
7–10	56	56	56	56
11	56	56	56	48
12	48	48	48	48

Table 8 reports the sensitivity analysis for P = 12 bytes. Under strict R5 the result is 64 bits in every variant, including a 10-year attack, a 30-fold change in the corrupted-frame fraction, a one-minute report period and an eight-channel attacker. The strict-R5 recommendation is therefore robust to the scenario assumptions. Under relaxed R5, λ_min ranges from 48 to 64 bits. The most demanding variants are a forgery target of 10^{-9} and an attacker-cost requirement of 30 days. Relaxation is therefore justified only when the scenario parameters are well known, which is precisely the "careful risk analysis" that R4 and R5 require.

Table 8: Sensitivity of Λ_Min (Bits) for P = 12 Bytes (Strict / Relaxed R5).

Variant	SF7	SF10	SF12
Baseline	64 / 56	64 / 56	64 / 48
$\varepsilon = 10^{-9}$	64 / 64	64 / 64	64 / 56
Attack duration 1 day	64 / 56	64 / 56	64 / 48
Attack duration 10 years	64 / 56	64 / 56	64 / 48
D_min = 1 h	64 / 56	64 / 48	64 / 48
D_min = 30 days	64 / 64	64 / 56	64 / 56
p_c = 0.01	64 / 56	64 / 56	64 / 48
p_c = 0.3	64 / 56	64 / 56	64 / 48
Report period 60 s	64 / 56	64 / 56	64 / 48
Eight-channel attacker	64 / 56	64 / 56	64 / 56

G. Comparison and Engineering Implications

The results refine current practice in three ways.

- The LoRaWAN MIC length. The 32-bit MIC of LoRaWAN [6] is adequate only if the network strictly limits failed verifications and treats them as a signal of attack. Without such limits, a year of forged traffic gives a forgery probability of up to 11% at SF7.
- The fixed 128-bit tag. Using the full 128-bit tag, as in [12]–[14], wastes up to about 19% of airtime and energy relative to 64 bits without a measurable security benefit under R5.
- The value of failure budgets. Generic tag-length selection [25] and GCM short-tag limits [20] do not capture the LoRa-specific staircase of airtime or the interaction between the failure budget and the CRC. Our results show that this interaction, not forgery probability, determines the design.

For practitioners, the recommendation is:

- Use Ascon-AEAD128 with a 64-bit tag and an implicit counter-derived nonce.
- Size the payload so that the tag does not cross a symbol-block boundary.
- Use tags shorter than 64 bits only with a documented relaxed-R5 analysis of the kind in Section V-F.

VI. LIMITATIONS

The PHY model assumes ideal frame synchronisation, no carrier-frequency or timing offset, noncoherent orthogonal detection, and block fading. Collisions and interference are represented only through the corrupted-frame fraction p_c, not simulated explicitly. The Hamming parity equations and the CRC initialisation follow a generic implementation, and bit-level equivalence with commercial LoRa transceivers is not claimed. Energy is represented by ToA at a fixed transmit power, and the energy of the cryptographic computation on the end device was not measured. The attacker model ignores the receiver diversity of multi-gateway networks except through the eight-channel variant. Finally, the relaxed interpretation of R5 is one admissible reading of the standard's risk-analysis clause, not guidance issued by NIST. Hardware measurements on LoRa transceivers, explicit collision models, and downlink and multicast traffic are left for future work.

VII. CONCLUSIONS

This paper performed the risk analysis that NIST SP 800-232 requires before truncated Ascon-AEAD128 tags can be used, and did so for LoRa uplinks. A verified bit-exact PHY simulator and an analytical threat model showed four things:

- The airtime cost of the tag is a staircase: about 61% of 8-bit tag increments are free, and a 64-bit tag saves 4.8–18.7% of ToA relative to 128 bits.
- The packet-error penalty of longer tags is at most 0.23 dB in AWGN and is not resolvable under block fading.
- A 32-bit tag without failure limits gives a forgery probability of up to 0.11 in a one-year attack.
- Under the strict failure budget, tags shorter than 64 bits let a single forged frame, and on average 0.8 benign CRC-passing errors per ten-year lifetime, force a key update.

The minimum admissible tag is therefore 64 bits under strict R5, and 48–56 bits under a relaxed R5 whose parameters are well characterised. These results give LoRa system designers a quantitative basis for choosing tag lengths in line with the new standard. Future work will validate the model on transceiver hardware and extend it to downlink, multicast and collision-dominated scenarios.

VIII. REFERENCES

[1] F. Adelantado, X. Vilajosana, P. Tuset-Peiro, B. Martinez, J. Melia-Segui and T. Watteyne, Understanding the limits of LoRaWAN, IEEE Commun. Mag. 55(9) (2017) 34–40. https://doi.org/10.1109/MCOM.2017.1600613

[2] M. Bor, U. Roedig, T. Voigt and J. Alonso, Do LoRa low-power wide-area networks scale?, Proc. 19th ACM Int. Conf. Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM) (2016) 59–67. https://doi.org/10.1145/2988287.2989163

[3] L. Vangelista, Frequency shift chirp modulation: The LoRa modulation, IEEE Signal Process. Lett. 24(12) (2017) 1818–1821. https://doi.org/10.1109/LSP.2017.2762960

[4] T. Elshabrawy and J. Robert, Closed-form approximation of LoRa modulation BER performance, IEEE Commun. Lett. 22(9) (2018) 1778–1781. https://doi.org/10.1109/LCOMM.2018.2849718

[5] J. Tapparel, O. Afisiadis, P. Mayoraz, A. Balatsoukas-Stimming and A. Burg, An open-source LoRa physical layer prototype on GNU Radio, Proc. IEEE 21st Int. Workshop on Signal Processing Advances in Wireless Communications (SPAWC) (2020). https://doi.org/10.1109/SPAWC48557.2020.9154273

[6] LoRa Alliance Technical Committee, LoRaWAN L2 1.0.4 Specification (TS001-1.0.4), LoRa Alliance, Technical Specification (2020).

[7] LoRa Alliance Technical Committee, RP002-1.0.3 LoRaWAN Regional Parameters, LoRa Alliance, Technical Specification (2021).

[8] Semtech Corporation, SX1272/3/6/7/8: LoRa Modem Designer's Guide, Application Note AN1200.13, Rev. 1 (2013).

[9] X. Yang, E. Karampatzakis, C. Doerr and F. Kuipers, Security vulnerabilities in LoRaWAN, Proc. IEEE/ACM 3rd Int. Conf. Internet-of-Things Design and Implementation (IoTDI) (2018) 129–140. https://doi.org/10.1109/IoTDI.2018.00022

[10] I. Butun, N. Pereira and M. Gidlund, Security risk analysis of LoRaWAN and future directions, Future Internet 11(1) (2019) 3. https://doi.org/10.3390/fi11010003

[11] P. Thaenkaew, B. Quoitin and A. Meddahi, Leveraging larger AES keys in LoRaWAN: A practical evaluation of energy and time costs, Sensors 23(22) (2023) 9172. https://doi.org/10.3390/s23229172
[12] Z. A. Simin, S. Aslan, M. M. Carvalho and D. Valles, Enhanced security of bidirectional communication in IoT-driven utility networks using Sertainty UXP and LoRaWAN, Sensors 26(6) (2026) 1752. https://doi.org/10.3390/s26061752
[13] V. Shkuratskyy, A. B. Usman, H. Bagheri, S. Hill and S. K. Memon, Evaluating lightweight and post-quantum encryption for compressed image transmission over LoRa in the era of edge AI, Discov. Electron. 3(1) (2026) 114. https://doi.org/10.1007/s44291-026-00261-4
[14] G. Amirkhanova, S. Ismailov, A. Amirkhanov, S. Adilzhanova, M. Zhasuzakova and S. Chen, A lightweight, end-to-end encrypted data pipeline for IIoT: An AES-GCM implementation for ESP32, MQTT, and Raspberry Pi, Information 17(1) (2026) 33. https://doi.org/10.3390/info17010033
[15] Atul, V. Shukla, M. K. Das, V. Shukla and D. Mishra, A reproducible security–performance benchmark protocol for lightweight AEAD on resource-constrained IoT nodes, IACR Cryptology ePrint Archive, Paper 2026/1882 (2026).
[16] M. Sönmez Turan, K. A. McKay, D. Chang, J. Kang and J. Kelsey, Ascon-based lightweight cryptography standards for constrained devices: Authenticated encryption, hash, and extendable output functions, NIST Special Publication 800-232 (2025). https://doi.org/10.6028/NIST.SP.800-232
[17] C. Dobraunig, M. Eichlseder, F. Mendel and M. Schläffer, Ascon v1.2: Lightweight authenticated encryption and hashing, J. Cryptol. 34(3) (2021) 33. https://doi.org/10.1007/s00145-021-09398-9
[18] C. Lefevre and B. Mennink, SoK: Security of the Ascon modes, IACR Trans. Symmetric Cryptol. 2025(1) (2025) 138–210. https://doi.org/10.46586/tosc.v2025.i1.138-210
[19] M. Sönmez Turan, K. A. McKay, D. Chang, L. Bassham, J. Kang, N. Waller, J. Kelsey and D. Hong, Status report on the final round of the NIST lightweight cryptography standardization process, NIST Internal Report 8454 (2023). https://doi.org/10.6028/NIST.IR.8454
[20] M. Dworkin, Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC, NIST Special Publication 800-38D (2007). https://doi.org/10.6028/NIST.SP.800-38D
[21] N. Ferguson, Authentication weaknesses in GCM, Public comment submitted to NIST (2005).
[22] C. J. Mitchell, Truncation attacks on MACs, Electron. Lett. 39(20) (2003) 1439–1440. https://doi.org/10.1049/el:20030921
[23] L. Ferreira, Sigforgery: Breaking and fixing data authenticity in Sigfox, Financial Cryptography and Data Security (FC 2021), Lecture Notes in Computer Science, Springer (2021) 331–350. https://doi.org/10.1007/978-3-662-64322-8_16
[24] E. Wagner, M. Serror, K. Wehrle and M. Henze, When and how to aggregate message authentication codes on lossy channels?, Applied Cryptography and Network Security (ACNS 2024), Lecture Notes in Computer Science 14584 (2024) 241–264. https://doi.org/10.1007/978-3-031-54773-7_10
[25] I. Andrushchak and O. Koloshko, Algorithmic selection of MAC tag length for short messages: Balancing the risks of substitution, delay, and throughput, Int. Sci. J. Eng. Agric. 4(6) (2025) 15–27. https://doi.org/10.46299/j.isjea.20250406.02
[26] D. Ochkas, H. Le Bouder and A. Pelov, Ascon-AEAD128 for COSE and JOSE, IETF Internet-Draft draft-ochkas-cose-ascon-03, work in progress (2026).
[27] P. Koopman and T. Chakravarty, Cyclic redundancy code (CRC) polynomial selection for embedded networks, Proc. Int. Conf. Dependable Systems and Networks (DSN) (2004) 145–154. https://doi.org/10.1109/DSN.2004.1311885
[28] R. Daidone, G. Dini and G. Anastasi, On evaluating the performance impact of the IEEE 802.15.4 security sub-layer, Comput. Commun. 47 (2014) 65–76. https://doi.org/10.1016/j.comcom.2014.03.025
[29] National Institute of Standards and Technology, ACVP-Server: Automated Cryptographic Validation Test System — generation/validation JSON files, GitHub repository usnistgov/ACVP-Server, accessed September 2026.
[30] J. G. Proakis and M. Salehi, Digital Communications, 5th ed., McGraw-Hill, New York, NY, USA (2008).
[31] C. J. Clopper and E. S. Pearson, The use of confidence or fiducial limits illustrated in the case of the binomial, Biometrika 26(4) (1934) 404–413. https://doi.org/10.1093/biomet/26.4.404