POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Course: Internet of Things - 5 CFU

Professor: Matteo Cesana
Students: Bruno Arena, Matteo Cavalleri
Personal IDs: 10806605, 10768765

Challenge 3 – Exercise

POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Abstract

The part addresses a theoretical exercise on LoRaWAN network performance. Given a network of 40 sensor nodes transmitting 20-byte payloads at 868 MHz with 125 kHz bandwidth and a rate of 2 packets per minute each, the maximum spreading factor compatible with a packet success rate of at least 75% is determined using a Pure ALOHA collision model. Additionally, a practical deployment scenario is analyzed to identify the most appropriate corrective action when measured performance falls below expectations and exhibits significant node-to-node variability.

Question #1 – Determine the largest LoRa spreading factor (SF) that guarantees a packet success rate of at least 75%.

The scenario consists of a single-cell LoRaWAN deployment in which N = 40 end-devices transmit 20-byte application payloads on the 868 MHz band with a channel bandwidth of B = 125 kHz. Each node generates packets independently according to a Poisson process with mean rate λ = 2 packets per minute, i.e. λ = 1/30 packets per second. Because LoRaWAN end-devices in Class A access the channel without any form of carrier sensing or coordinated scheduling, the medium-access behaviour is correctly described by the Pure ALOHA model. Under this model the probability that a generic transmitted packet is received without collision — hereafter referred to as the packet success rate (PSR) — is given by

$$PSR = e^{-2G}$$

$$G = N \cdot \lambda \cdot T_{air}$$

where G is the normalised offered load and T_{air} is the on-air time required to transmit one packet at the chosen spreading factor (SF). The vulnerability window of Pure ALOHA equals $2 \cdot T_{air}$, which is the source of the factor of two in the exponent.

Imposing the design requirement $PSR \geq 0.75$ on the closed-form expression yields:

$$e\text{\^{}}(-2G) \geq 0.75 \quad \Leftrightarrow \quad 2G \leq -ln(0.75) \approx 0.2877$$

from which the maximum admissible airtime follows directly:

$$T_{air} \leq -\frac{ln(0.75)}{2 \cdot N \cdot \lambda} = \frac{0.2877}{2 \cdot 40 \cdot \frac{1}{30}} = \frac{0.2877}{2.667} \approx 0.108\, s \approx 108\, ms$$

The problem therefore reduces to selecting the largest spreading factor whose on-air time does not exceed approximately $108\, ms$, since increasing SF roughly doubles T_{air}at every step and therefore doubles the offered load G.

The airtime of a LoRa frame depends on SF, bandwidth, coding rate, payload length, header configuration and the low-data-rate optimisation flag. Using the standard Semtech formulation (as implemented by the well-known TTN airtime calculator) with the parameters of the exercise — $BW = 125\,kHz$, $payload = 20\,B$, coding rate $CR = \frac{4}{5}$, explicit header, CRC enabled, $preamble\ length = 8$ symbols, and low-data-rate optimisation activated automatically for SF11 and SF12 — the resulting airtimes are:

SF7 : $T_{air} \approx 56.6\,ms$
SF8 : $T_{air} \approx 102.7\,ms$
SF9 : $T_{air} \approx 185.3\,ms$
SF10 : $T_{air} \approx 329.7\,ms$
SF11 : $T_{air} \approx 660.5\,ms$
SF12 : $T_{air} \approx 1318.9\,ms$

Comparison with the threshold $T_{air} \leq 108\,ms$ shows that SF7 and SF8 satisfy the bound, whereas SF9 and higher exceed it. The two boundary cases can be checked numerically by substituting the corresponding airtimes back into the PSR expression.
For SF8, the offered load is $G = 1.333 \cdot 0.1027 \approx 0.137$, leading to $PSR = e^{-0.274} \approx 0.761$, which meets the 75% requirement.
For SF9, the offered load increases to $G = 1.333 \cdot 0.1853 \approx 0.247$, which yields $PSR = e^{-0.494} \approx 0.610$, clearly below the threshold.

Answer
In conclusion, the largest spreading factor that meets the design requirement is SF8, with a packet success rate of about 76% — just above the 75% target.
SF9 already fails (T_air ≈ 185 ms > 108 ms, PSR ≈ 61%): each SF step roughly doubles the airtime and therefore the offered load G, which makes the collision probability grow too quickly for the constraint to remain feasible.

Question #2 – Performance

The deployed network operates at SF8, i.e. at the operating point identified in Question #1, where Pure ALOHA theory predicts an average PSR of approximately 76%. In practice, however, two qualitatively different deviations from the theoretical behaviour are reported: the measured average PSR is markedly lower than expected, and — more importantly — the per-node performance is far from uniform, with some nodes performing close to the expected value while others perform considerably worse.

The two symptoms must be examined jointly. Channel congestion, modelled by the Pure ALOHA collision term, would affect every node in essentially the same way, since all nodes share the same spreading factor, the same packet generation rate λ, and the same channel. Collisions alone are therefore consistent with a uniform degradation of the PSR, but cannot account for the observed inter-node variability. The heterogeneity must instead originate in the physical layer, and specifically in the link budget of the individual nodes: when the path loss to the gateway is large, the received signal-to-noise ratio drops below the demodulation threshold corresponding to SF8, and packets from those nodes are lost regardless of whether or not they collide. Conversely, nodes that are close to the gateway, or in line-of-sight conditions, retain a healthy SNR margin and continue to perform near the theoretical value. The dominant impairment is therefore a coverage-limited regime affecting a subset of the nodes, superimposed on the baseline collision loss.

Three possible interventions are considered.

(a) Reducing the number of nodes. Decreasing N would lower the aggregate offered load and therefore reduce the collision probability uniformly across the network. However, this action does not modify the per-link path loss of the disadvantaged nodes, and the SNR of those nodes would remain below the SF8 sensitivity threshold. The inter-node variability — the dominant feature of the observed performance — would therefore persist. The action treats a secondary effect (collisions) while leaving the root cause (coverage) untouched.

(b) Moving the nodes closer to the gateway. Reducing the average distance to the gateway directly attacks the root cause identified above: the path loss decreases, the received SNR at the gateway rises above the SF8 sensitivity threshold, and the previously disadvantaged nodes recover. Crucially, this intervention does not alter the airtime of any transmission and therefore does not increase the offered load G; the collision component of the loss is left at the value already accounted for by the design at SF8. Both observed symptoms — the low average PSR and the inter-node variability — are addressed simultaneously.

(c) Increasing the spreading factor. A higher SF improves receiver sensitivity by approximately 2.5 dB per step, which would partially compensate the link-budget deficit of the weakest nodes. The drawback is that, as shown in Question #1, the airtime roughly doubles at every SF increment. With N and λ unchanged, the offered load G doubles as well, and so does the exponent of the collision term in the Pure ALOHA expression. The PSR of every node — including those that currently perform well — is therefore degraded by additional collision loss. The action is in fact a trade-off between coverage and capacity, and on a network already operating near the 75% PSR threshold it is likely to worsen, rather than improve, the average performance.

Answer

Among the three options the only one that simultaneously eliminates the dominant impairment and preserves the channel load is option (b): the most appropriate corrective action is to move the nodes closer to the gateway. This intervention restores an adequate link budget for the disadvantaged nodes, improves both the mean and the uniformity of the packet success rate across the 40-node population, and does so without sacrificing the capacity margin established by the SF8 design point.