

A Third of the Nodes, Two Thirds of the Airtime: Spreading-Factor Inequity and the Real Capacity Limit of a LoRaWAN Cell

Dr. P. Sivasubramanian

Associate Professor, Kalasalingam Academy of Research and Education, Srivilliputhur, Tamil Nadu, India.

Email: sivaresearch948@gmail.com

ABSTRACT

LoRaWAN deployments are commonly sized by a headline node count per gateway. That figure conceals a structural asymmetry: because the spreading factor is raised to reach distant nodes and each increment roughly doubles the time on air, the nodes furthest from the gateway occupy the channel for far longer per message than nearby ones. The consequence for capacity and for service equity is rarely quantified during planning. This study quantifies how channel airtime is distributed across spreading factors in a realistic cell, how packet delivery degrades with node count for near and far nodes separately, and which of the available mitigations actually recovers capacity. Time on air was computed from the standard LoRa symbol formulation for a 20-byte payload at 125 kHz bandwidth and a coding rate of 4/5. Nodes were distributed uniformly over a 6 km disc and assigned spreading factors by range ring. Per-spreading-factor channel load was converted to delivery probability using the pure-ALOHA relation, treating spreading factors as orthogonal, across eight uplink channels. Node counts from 200 to 8,000 and five mitigations were evaluated, and the results were compared against the 1 % regulatory duty cycle. Time on air ranged from 56.6 ms at SF7 to 1,318.9 ms at SF12, a factor of 23.3. Nodes at SF12 constituted 36.0 % of the population and consumed 63.4 % of channel airtime, while SF7 nodes were 4.0 % of the population and consumed 0.3 %. At 1,000 nodes reporting six times per hour, network delivery was 91.2 % but ranged from 99.9 % at SF7 to 82.1 % at SF12; at 8,000 nodes it was 56.1 % overall, 99.2 % at SF7 and 20.5 % at SF12. Among mitigations at 2,000 nodes, a second gateway raised delivery from 83.7 % to 91.2 % and halving the payload to 87.2 %, while capping the spreading factor at SF10 slightly reduced it to 83.0 %. Reducing the reporting rate from six to two per hour raised delivery to 94.0 % while cutting delivered messages from 10,047 to 3,759 per hour. The 1 % duty cycle permits only 27 messages per hour at SF12 against 636 at SF7. Capacity in a LoRaWAN cell is consumed by the range, not by the node count, and the degradation falls almost entirely on the far nodes that the technology is chosen to reach. Spreading-factor capping does not recover capacity because it redistributes load rather than removing it; gateway density and payload discipline do.

Keywords: LoRaWAN; low-power wide-area network; spreading factor; network capacity; Internet of Things; duty cycle.

1. INTRODUCTION

Low-power wide-area networking has made it practical to connect large numbers of battery-powered sensors over kilometres without cellular subscriptions, and LoRaWAN is the most widely deployed of the unlicensed variants [1-3,16]. Its appeal to a system designer is a simple one: a single gateway is advertised as supporting thousands of nodes over a radius of

several kilometres, which makes the economics of a municipal or agricultural deployment work.

The mechanism that provides the range is also the one that constrains the capacity. LoRa achieves sensitivity by spreading each symbol over a longer interval, and the spreading factor is raised for nodes with weaker links. Each increment approximately doubles the symbol duration, so a message from the edge of the cell occupies the shared channel for an order of magnitude longer than the same message from near the gateway [1,9].

This creates a coupling that a node count does not express. Two deployments with identical numbers of sensors will have very different capacity if one is compact and the other is spread to the limit of range, and within a single deployment the nodes competing hardest for the channel are the ones contributing least to the total [4-6].

The question matters because the far nodes are usually the reason the technology was chosen. A designer selecting LoRaWAN over a short-range mesh is buying range, and this analysis asks what that range costs in delivered messages and who pays for it. Three questions are addressed: how airtime is actually distributed across the cell; how delivery degrades with scale for near and far nodes separately; and which mitigations recover capacity rather than merely moving the problem.

2. MODEL

2.1 Time on Air

Time on air was computed from the standard LoRa formulation [1,2]. With SF the spreading factor, BW the bandwidth and PL the payload in bytes, the symbol period and packet duration are

$$T_{sym} = \frac{2^{SF}}{BW}, T_{preamble} = (n_{pre} + 4.25) T_{sym} \quad (1)$$

$$n_{payload} = 8 + \max \left(\left\lceil \frac{8 PL - 4 SF + 28 + 16 CRC - 20 IH}{4 (SF - 2 DE)} \right\rceil (CR + 4), 0 \right) \quad (2)$$

with DE the low-data-rate optimisation flag, enabled at SF11 and SF12. The 8-symbol preamble, explicit header, CRC and a coding rate of 4/5 were used throughout, with a 20-byte application payload representative of a sensor reading with metadata.

2.2 Cell Geometry and Channel Access

Nodes were distributed uniformly over a disc of 6 km radius and assigned spreading factors by range ring, which approximates the outcome of adaptive data rate under a monotonic path-loss model [2,14]. Because area grows with the square of radius, the outer rings contain disproportionately many nodes even though they are geometrically narrow.

Channel access is unslotted and uncoordinated, so the pure-ALOHA relation applies within each spreading factor [11,17]. Treating spreading factors as orthogonal — which is an idealisation discussed in Section 4.3 — the success probability for a message at spreading factor s is

$$P_{success}(s) = \exp(-2 G_s), G_s = \frac{\lambda_s T_{air}(s)}{N_{ch}} \quad (3)$$

where λ is the message arrival rate at that spreading factor and N the number of uplink channels. The factor of two in the exponent is the vulnerable period of unslotted ALOHA: a message is lost if any other transmission on the same spreading factor and channel begins within one packet duration either side of it [11].

Table 1: Model parameters.

Parameter	Value
Bandwidth	125 kHz
Coding rate	4/5
Application payload	20 bytes
Preamble	8 symbols, explicit header, CRC enabled
Cell radius	6 km, uniform node density
Spreading factor rings	SF7 to 1.2 km, SF8 to 2.0, SF9 to 2.8, SF10 to 3.8, SF11 to 4.8, SF12 to 6.0
Uplink channels	8 (16 in the second-gateway scenario)
Reporting rate	6 messages per hour per node
Node count	200 to 8,000 per gateway
Regulatory duty cycle	1 % per sub-band

Note: Ring boundaries are representative of a suburban path-loss environment and are stated assumptions rather than measurements. Absolute delivery ratios are specific to them; the structural findings follow from the time-on-air ratios and the area weighting.

3. RESULTS

3.1 Time on Air and Who Consumes It

Table 2 and Figure 1 give the time on air and the node distribution for each spreading factor.




Fig.1: Time on air for a 20-byte payload at each spreading factor.

Figure 2 states the structural problem. Nodes at SF12 are 36.0 % of the population and consume 63.4 % of the airtime; nodes at SF7 are 4.0 % of the population and consume 0.3 %. The ratio of airtime consumed per node between the two extremes is over two hundredfold,

combining the 23.3-fold time-on-air penalty with the area weighting that places most nodes in the outer rings.

Table 2: Time on air and node distribution by spreading factor.

Spreading factor	Time on air (ms)	Relative to SF7	Share of nodes	Share of airtime
SF7	56.6	1.0×	4.0 %	0.3 %
SF8	102.9	1.8×	7.1 %	1.0 %
SF9	185.3	3.3×	10.7 %	2.6 %
SF10	370.7	6.6×	18.3 %	9.1 %
SF11	741.4	13.1×	23.9 %	23.6 %
SF12	1,318.9	23.3×	36.0 %	63.4 %




Fig.2: Share of nodes against share of channel airtime, by spreading factor.

The consequence for planning is that a node count is not a capacity statement. Two deployments of a thousand nodes, one confined to 2 km and one spread to 6 km, present the gateway with channel loads differing by roughly an order of magnitude, and nothing in a datasheet node figure distinguishes them [4].

3.2 Scaling and Who Pays for It

Table 3 and Figure 3 give the packet delivery ratio as the number of nodes per gateway increases.

Table 3: Packet delivery ratio against node count.

Nodes per gateway	Network average	SF7 nodes	SF12 nodes	Messages delivered per hour
200	98.1 %	100.0 %	96.1 %	1,178
500	95.4 %	100.0 %	90.6 %	2,863
1,000	91.2 %	99.9 %	82.1 %	5,473
2,000	83.7 %	99.8 %	67.3 %	10,047
4,000	71.8 %	99.6 %	45.3 %	17,232
8,000	56.1 %	99.2 %	20.5 %	26,950




Fig.3: Packet delivery ratio against node count for the network and for the two extreme spreading factors.

The average delivery ratio falls from 98.1 % to 56.1 % across the range, which is the figure a network operator would report [12,13]. The two columns either side of it show that the average describes almost nobody.

Nodes at SF7 deliver 99.2 % of their messages even at 8,000 nodes, because they occupy the channel for 56.6 ms and compete only with the small population that shares their spreading factor. Nodes at SF12 fall to 20.5 %: four messages in five are lost. The network does not degrade uniformly under load; it degrades at the edge, and the edge is where the nodes that justified the technology are.

This also means that the deployments most exposed to the problem are those that are working hardest at range — agricultural, environmental and utility monitoring spread over open country — while dense urban deployments with short links scale comfortably [13,16]. The two look identical in a node-count specification.

3.3 Which Mitigations Work

Table 4 and Figure 4 compare the mitigations evaluated at 2,000 nodes per gateway.

Table 4: Mitigations evaluated at 2,000 nodes.

Mitigation	Delivery ratio	SF12 delivery	Messages delivered per hour	Assessment
Baseline	83.7 %	67.3 %	10,047	—
Payload 20 → 10 bytes	87.2 %	74.3 %	10,462	Effective and cheap
Reporting 6 → 2 per hour	94.0 %	87.6 %	3,759	Better ratio, far less data
Cap spreading factor at SF10	83.0 %	n/a	9,958	No benefit; see text
Cap at SF10 with 10-byte payload	86.4 %	n/a	10,369	Benefit comes from the payload
Second gateway, 16 channels	91.2 %	82.1 %	10,946	Most effective

Spreading-factor capping does not help. This is the result most likely to be counterintuitive and it is worth stating carefully. Capping at SF10 reduced delivery marginally, from 83.7 % to 83.0 %, because the traffic previously carried at SF11 and SF12 does not disappear — it is reassigned to SF10, where it competes with the existing SF10 population. Total airtime falls, since SF10 messages are shorter, but it is concentrated into fewer spreading factors, and the two effects nearly cancel.




Fig.4: Delivery ratio under each mitigation.

In a real deployment the outcome is worse than modelled. The model assumes a capped node still reaches the gateway at SF10, whereas the reason it was assigned SF12 is that its link budget requires the extra sensitivity. Capping in practice means those nodes are lost rather than reassigned, so the measured benefit would be a delivery ratio computed over a smaller population — an improvement in the statistic and a reduction in the service.

Payload reduction works. Halving the payload raised delivery from 83.7 % to 87.2 % and raised delivered volume as well, because time on air falls with payload while the message count is unchanged. It is the cheapest mitigation available and is entirely within the application developer's control, which makes it the first thing to examine when a deployment underperforms.

Gateway density works best. A second gateway, modelled as doubling the available channels, raised delivery to 91.2 % and delivered volume to 10,946 messages per hour. It is also the only mitigation that improves service for SF12 nodes without reducing what the network carries, at the cost of the capital and backhaul of a second site.




Fig.5: Delivery ratio against messages delivered per hour for each mitigation.

Figure 5 separates two objectives that are frequently conflated. Reducing the reporting rate produces the best delivery ratio of any single mitigation, at 94.0 %, and delivers 3,759 messages per hour against the baseline 10,047. A network reported as performing well on delivery ratio may simply have been asked to carry less, and a deployment specification should state both quantities.

3.4 The Duty Cycle Binds First

Table 5 and Figure 6 give the number of messages per hour that the regulatory duty cycle permits at each spreading factor.

Table 5: Messages per hour permitted by a 1 % duty cycle.

Spreading factor	SF7	SF8	SF9	SF10	SF11	SF12
Messages per hour	636	350	194	97	49	27
Headroom at 6 per hour	106×	58×	32×	16×	8×	4.5×




Fig.6: Messages per hour permitted by the regulatory duty cycle, on a logarithmic scale.

The 1 % duty cycle applied in the European sub-bands permits 636 messages per hour at SF7 and 27 at SF12 [15]. At the modelled reporting rate of six per hour, an SF7 node uses under 1 % of its allowance and an SF12 node uses 22 %.

The practical implication concerns retransmission and confirmed messaging. A far node that loses two thirds of its uplinks, as it does at 2,000 nodes, cannot simply retransmit its way to reliability: three attempts per message at SF12 consumes two thirds of the regulatory allowance, and confirmed downlink acknowledgements consume the gateway's far tighter downlink budget as well [10]. The mitigation most naturally reached for by an application developer is the one least available to the nodes that need it.

4. DISCUSSION

4.1 Principal Findings

Four results follow. Airtime consumption is grossly unequal, with 36.0 % of nodes consuming 63.4 % of the channel. Degradation under load falls almost entirely on the far nodes, whose delivery ratio reaches 20.5 % at 8,000 nodes while near nodes remain above 99 %. Spreading-factor capping does not recover capacity because it redistributes load rather than removing it, whereas payload reduction and gateway density do. And the regulatory duty cycle

leaves SF12 nodes only a 4.5-fold headroom, which forecloses retransmission as a remedy for exactly the nodes that lose most messages.

4.2 Implications for Deployment

Table 6 summarises each finding, its implication and the corresponding recommendation for practice.

Table 6: Implications.

Finding	Implication	Practice
SF12 airtime is 23.3× that of SF7	Range is bought with shared capacity	Size the cell by airtime, not by node count
36 % of nodes use 63 % of airtime	Node count conceals the load	Report the spreading-factor distribution in any deployment plan
SF12 delivery falls to 20.5 % at scale	The far nodes fail first	Monitor delivery ratio per spreading factor, not network-wide
Capping SF gave no benefit	Redistribution is not reduction	Do not cap the spreading factor expecting capacity; it removes nodes
Halving payload gave +3.5 points	Airtime scales with payload	Audit application payloads before adding infrastructure
Second gateway gave +7.5 points	Channel capacity is the binding constraint	Plan gateway density from the SF distribution, not the radius
Duty cycle leaves SF12 only 4.5× headroom	Retransmission is unavailable where it is needed	Do not rely on confirmed messaging at the cell edge

The third row is the cheapest recommendation and would change how these networks are operated. Delivery ratio is almost always reported as a network aggregate, and the aggregate is dominated by the near nodes that never have difficulty. A per-spreading-factor breakdown requires no additional instrumentation — the server already records the data rate of every uplink — and would reveal edge degradation long before the average moves.

The fourth row is worth emphasising because spreading-factor capping is a commonly offered configuration option and appears to be a capacity control. On this model it is not. Where it appears to help in a real deployment, the mechanism is usually that the nodes which could not close the link at the capped rate have stopped reporting altogether, which improves every statistic computed over the remainder.

4.3 Limitations

The collision model treats spreading factors as perfectly orthogonal, so a transmission at one never interferes with another. Imperfect orthogonality is well documented, particularly for a strong near-field signal at one spreading factor against a weak edge signal at another [7,8], and including it would reduce delivery further and would do so disproportionately at the high spreading factors already identified as worst served. The model is therefore optimistic in the direction of its own conclusion.

The capture effect is also omitted. Where two colliding signals differ sufficiently in power the stronger is often decoded successfully, so pure ALOHA underestimates delivery [5,12]; this acts in the opposite direction to the orthogonality simplification and favours near nodes, again widening the inequity rather than narrowing it. Modelling both effects together requires a link-level simulation with a propagation model and is the natural extension of this analysis.

Spreading factors are assigned by deterministic range rings, whereas adaptive data rate assigns them from measured link margin and responds to fading, shadowing and node mobility;

real distributions are broader and less clean [14]. Uniform node density and a circular cell are idealisations, and a real deployment concentrated along a road or a river would produce a different distribution. The second-gateway scenario is modelled simply as doubled channel capacity and ignores the diversity gain from a node being heard by two gateways, which would make additional gateways more effective than reported here.

Finally, no downlink traffic, no join procedure, no acknowledgement and no retransmission behaviour is modelled, and all four consume airtime [2,10]. Absolute delivery ratios are consequently optimistic, while the comparisons between spreading factors and between mitigations follow from the time-on-air ratios and are robust to these omissions.

5. CONCLUSION

A LoRaWAN cell is limited by airtime rather than by node count, and airtime is distributed extremely unevenly. In the cell modelled here, nodes at SF12 constituted 36.0 % of the population and consumed 63.4 % of the channel, because a 20-byte message takes 1,318.9 ms at SF12 against 56.6 ms at SF7. The consequence under load is not a uniform degradation but a selective one: at 8,000 nodes the network delivered 56.1 % of messages overall, 99.2 % from near nodes and 20.5 % from far ones.

That is an uncomfortable result for a technology chosen for its range, and it shapes what can be done about it. Capping the spreading factor does not recover capacity, because the load is reassigned rather than removed, and in practice it works by silently dropping the nodes that cannot close the link at the capped rate. Halving the application payload recovered 3.5 points of delivery at no cost, and a second gateway recovered 7.5 while increasing the volume carried.

The recommendation for practice is a measurement one before it is an engineering one. Deployment plans and operational dashboards should carry the spreading-factor distribution and the delivery ratio per spreading factor rather than network aggregates, because the aggregate is dominated by the nodes that were never at risk and moves only after the edge of the cell has already failed.

DECLARATIONS

Ethical approval. Not applicable. This is a computational study; no human participants, personal data or field measurements were involved.

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Conflicts of interest. The author declares no conflicts of interest. The author has no financial or advisory relationship with any network operator, gateway vendor or semiconductor supplier.

Scope disclaimer. No commercial network, gateway, device or operator is evaluated, and no result describes a deployed system. Regional regulations differ; the duty-cycle figures apply to the European sub-bands and not to jurisdictions using other access rules.

Model disclaimer. Delivery ratios are obtained from an analytical ALOHA model with orthogonal spreading factors and no capture effect, not from link-level simulation or measurement. Absolute values are optimistic; see Section 4.3.

Data and code availability. This is a computational study; no field measurements or external datasets were used. All results are derived from Equations (1) to (3) using the parameters listed in Table 1, and every reported value can be reproduced from them. The analysis code is available from the corresponding author upon reasonable request.

Author's contribution. Dr. P. Sivasubramanian is the sole author of this work and was responsible for the conception and design of the study, development of the analytical model, analysis and interpretation of the results, preparation of the figures and tables, and drafting,

revision and final approval of the manuscript. The author is accountable for all aspects of the work.

REFERENCES

1. Semtech Corporation. LoRa modulation basics. Application Note AN1200.22. Camarillo: Semtech; 2015.
2. LoRa Alliance. LoRaWAN 1.0.4 specification. Fremont: LoRa Alliance; 2020.
3. Raza U, Kulkarni P, Sooriyabandara M. Low power wide area networks: an overview. *IEEE Commun Surv Tutor*. 2017;19(2):855–873.
4. Adelantado F, Vilajosana X, Tuset-Peiro P, Martinez B, Melia-Segui J, Watteyne T. Understanding the limits of LoRaWAN. *IEEE Commun Mag*. 2017;55(9):34–40.
5. Bor MC, Roedig U, Voigt T, Alonso JM. Do LoRa low-power wide-area networks scale? *Proc 19th ACM Int Conf Model Anal Simul Wirel Mob Syst*. 2016:59–67.
6. Georgiou O, Raza U. Low power wide area network analysis: can LoRa scale? *IEEE Wirel Commun Lett*. 2017;6(2):162–165.
7. Mahmood A, Sisinni E, Guntupalli L, Rondón R, Hassan SA, Gidlund M. Scalability analysis of a LoRa network under imperfect orthogonality. *IEEE Trans Ind Inform*. 2019;15(3):1425–1436.
8. Croce D, Gucciardo M, Mangione S, Santaromita G, Tinnirello I. Impact of LoRa imperfect orthogonality: analysis of link-level performance. *IEEE Commun Lett*. 2018;22(4):796–799.
9. Reynders B, Meert W, Pollin S. Range and coexistence analysis of long range unlicensed communication. *Proc 23rd Int Conf Telecommun*. 2016:1–6.
10. Sorensen RB, Kim DM, Nielsen JJ, Popovski P. Analysis of latency and MAC-layer performance for class A LoRaWAN. *IEEE Wirel Commun Lett*. 2017;6(5):566–569.
11. Abramson N. The ALOHA system: another alternative for computer communications. *Proc Fall Joint Comput Conf*. 1970:281–285.
12. Van den Abeele F, Haxhibeqiri J, Moerman I, Hoebeke J. Scalability analysis of large-scale LoRaWAN networks in ns-3. *IEEE Internet Things J*. 2017;4(6):2186–2198.
13. Magrin D, Centenaro M, Vangelista L. Performance evaluation of LoRa networks in a smart city scenario. *Proc IEEE Int Conf Commun*. 2017:1–7.
14. Haxhibeqiri J, De Poorter E, Moerman I, Hoebeke J. A survey of LoRaWAN for IoT: from technology to application. *Sensors*. 2018;18(11):3995.
15. European Telecommunications Standards Institute. ETSI EN 300 220-2: short range devices operating in the frequency range 25 MHz to 1000 MHz. Sophia Antipolis: ETSI; 2018.
16. Centenaro M, Vangelista L, Zanella A, Zorzi M. Long-range communications in unlicensed bands: the rising stars in the IoT and smart city scenarios. *IEEE Wirel Commun*. 2016;23(5):60–67.
17. Zorbas D, Abdelfadeel K, Kotzanikolaou P, Pesch D. TS-LoRa: time-slotted LoRaWAN for the industrial Internet of Things. *Comput Commun*. 2020;153:1–10.