

LoRaWAN Scalability Under Worst-Case ALOHA: A Single-Channel Stress Test for Dense IoT Networks

Izzah Nilamsyukriyah Buang

24010024@siswa.unimas.my

Universiti Malaysia Sarawak

Kartinah Zen

Universiti Malaysia Sarawak

Adnan Shahid Khan

Universiti Malaysia Sarawak

Research Article

Keywords: LoRaWAN, ALOHA protocol stress test, IoT scalability, energy-reliability trade-off, high-throughput IoT, fragmented data transfer

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11198550/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

LoRaWAN Scalability Under Worst-Case ALOHA: A Single-Channel Stress Test for Dense IoT Networks

Izzah Nilamsyukriyah Buang^{a,*}, Kartinah Zen^a, Adnan Shahid Khan^a

^a Faculty of Computer Science and Information Technology, Universiti Malaysia Sarawak, 94300 Kota Samarahan, Sarawak, Malaysia

* Corresponding Author. Email: 24010024@siswa.unimas.my

ABSTRACT

As Internet of Things deployments move beyond periodic telemetry toward high-throughput applications, LoRaWAN's scalability limits remain insufficiently quantified under worst-case conditions. Multi-channel systems report capacities of hundreds to thousands of nodes, but confounding factors obscure attribution of these gains. This study isolates single-channel gateway performance under pure ALOHA access, deliberately pushing traffic demand to characterise its worst-case degradation envelope. Experiments span node densities from 1 to 200, mobility speeds from 0 to 72 km/h, and confirmed and unconfirmed transmission modes, across a three-phase, 60-scenario design comprising 1,682,057 packet transmissions. We identify a scenario-specific critical threshold at approximately 50 to 60 nodes, where packet delivery ratio (PDR) drops from 97.9% (95% CI 97.7–98.1%) at 50 nodes to 77.9% (95% CI 76.9–79.0%) at 60 nodes. Confirmed transmission mode improves reliability by 33.3 percentage points at threshold density but incurs 139% energy overhead relative to unconfirmed mode; this advantage inverts by 200 nodes, where confirmed mode delivers 3.8 percentage points worse PDR whilst consuming 706% more energy, a critical limitation for streamed, fragmented payloads such as multi-chunk images. Node mobility shows negligible impact under the evaluated collision-dominated regime ($F(4,20) = 0.0005$, $p \approx 1.0$), as collisions dominate over mobility-induced variation beyond 20% contention. These findings motivate a capacity decomposition framework separating protocol-level contention limits from hardware- and architecture-driven capacity multipliers, providing network designers a worst-case reference point for dense, mobile LoRaWAN deployments and a benchmark for future medium-access-control enhancements.

Keywords: LoRaWAN, ALOHA protocol stress test, IoT scalability, energy-reliability trade-off, high-throughput IoT, fragmented data transfer

1. INTRODUCTION

Long Range Wide Area Network (LoRaWAN) has emerged as a dominant low-power wide-area network (LPWAN) technology with over 200 million devices deployed globally (LoRa Alliance

2023). Its use is increasingly extending beyond simple, sparsely reported telemetry toward high-throughput, non-traditional applications such as mobile device tracking, disaster-response telemetry and chunked image or file transfer, and the fundamental scalability limits of LoRaWAN under such worst-case, high-demand conditions remain insufficiently quantified. Streaming a single image over LoRaWAN, for example, requires dividing the file into dozens of small packets transmitted in rapid succession (Jebril et al. 2018; Chaparro et al. 2022), and this behaviour creates severe, localised channel contention that differs qualitatively from the sparse, independent transmissions assumed in conventional LoRaWAN scalability studies. Reported capacity estimates vary dramatically from 50 to 6000 nodes per gateway (Figueiredo et al. 2022; Fragkopoulos et al. 2023), reflecting methodological heterogeneity where multi-channel and multi-gateway configurations introduce multiple confounding factors that obscure attribution of performance gains. LoRaWAN employs an unslotted ALOHA-based medium access control (MAC) protocol, where end devices transmit whenever data is available without coordination or carrier sensing (Khan et al. 2021). This pure random access approach inherits the fundamental scalability constraints of ALOHA whilst benefiting from LoRa's spreading factor (SF) diversity, which creates quasi-orthogonal transmission channels (Gamage et al. 2022; Chasserat et al. 2024).

Scalability challenges in dense LoRaWAN deployments are inherently multi-dimensional, encompassing not only MAC efficiency but also connectivity management, energy optimisation, and mobility support across network layers (Buang et al. 2025; Jouhari et al. 2023). Classical unslotted ALOHA achieves maximum throughput of approximately 18% at optimal offered load $G \approx 0.5$ (Abramson 1970), with rapid throughput degradation beyond this point. Whilst LoRaWAN's quasi-orthogonal spreading factors (SFs) and multi-channel architectures modify this theoretical limit, the relationship between classical ALOHA theory and practical LoRaWAN performance under controlled single-channel conditions has not been rigorously characterised.

Three recurring limitations emerge consistently across the literature. First, most studies employ multi-channel gateways, which prevents establishment of per-channel baselines (Almuhaya et al. 2025; Yu et al. 2024). Recent work by Almuhaya et al. (2024, 2025) examined LoRa scalability through zone-based and multi-band architectures, demonstrating performance improvements in dense scenarios though evaluating scalability under mixed conditions rather than isolating single-channel ALOHA behaviour. Similarly, Yu et al. (2024) investigated concurrent transmission across massive logical channels, yet primarily within multi-channel contexts where hardware parallelism confounds protocol-level analysis. Second, energy and reliability are studied independently rather than as coupled trade-offs under varying contention levels (Maudet et al. 2021). Third, mobility studies focus on low-density scenarios (Anwar et al. 2021; Gorrela et al. 2026), leaving unresolved whether mobility-induced channel variations matter when collision dominates. Previous work (Buang et al. 2025) examined connectivity and energy optimisation in LoRaWAN routing protocols for the mobile Internet of Things (IoT), highlighting the multi-dimensional nature of scalability challenges. However, single-channel ALOHA baselines necessary for isolating protocol-level effects from architectural enhancements remained unestablished. Unlike prior studies that implicitly embed ALOHA behaviour within multi-channel or multi-gateway systems, this work explicitly isolates and quantifies the fundamental single-channel ALOHA limit under

controlled worst-case conditions. While prior studies have examined LoRaWAN scalability in both broad multi-channel or multi-gateway configurations (Almuhaya et al. 2024, 2025; Yu et al. 2024) and single-gateway deployments (Haxhibeqiri et al. 2017), this study specifically isolates the single-channel, single-gateway ALOHA case under a controlled stress test to resolve its critical transition zone at single-node granularity. Prior work evaluates scalability under mixed architectural and protocol conditions, whereas this study isolates the protocol-level limit under controlled worst-case assumptions.

In this paper, we present three empirically grounded contributions. First, under a deliberately pushed stress scenario, we identify a scenario-specific critical threshold at approximately 50 to 60 nodes, under the evaluated traffic pattern, SF distribution and single-channel assumptions, identified through progressively targeted confirmed-mode density sampling first across 50 to 100 nodes and then within the 50 to 60 node transition zone at the fixed 5 m/s mobility condition used throughout the Phase 2 and Phase 3 drill-down. At this fixed mobility condition, packet delivery ratio (PDR) exhibits sharp nonlinear degradation from 97.4% at 50 nodes to 77.9% at 60 nodes, a fixed-mobility figure that differs slightly from the 97.9% value averaged across all five mobility conditions reported for 50 nodes elsewhere in the paper, and continues to decline to 33.96% by 100 nodes. Piecewise regression identifies the onset of sharp PDR degradation at 52 nodes (97.9% PDR), with the transition complete by 75 nodes (47.6% PDR). This transition regime exhibits degradation rates approximately 49 times steeper than the low-density flat regime, confirming a distinct three-regime structure, namely stable operation (≤ 52 nodes), sharp collapse (52 to 75 nodes), and saturated gradual decline (> 75 nodes). This threshold is consistent with classical ALOHA predictions when parameterised with LoRaWAN-specific airtime and traffic characteristics. Second, we demonstrate that node mobility exhibits negligible effect under the evaluated collision-dominated, periodic-traffic regime ($F(4,20) = 0.0005$, $p \approx 1.0$), calling into question assumptions in parts of the literature motivating significant research attention to mobility-specific optimisation in dense deployments. Third, we quantify confirmed-mode energy-reliability trade-off collapse, with efficiency declining from 0.545 J per percentage point at threshold to -24.5 J per percentage point at 200 nodes, a transition through 47.3 J per percentage point at 100 nodes, revealing conditions under which retransmission mechanisms provide diminishing or genuinely negative returns.

These findings establish a conceptual framework distinguishing protocol-level capacity from system-level capacity through the following decomposition:

$$C_{system} \approx C_o \times N_{channels} \times \eta_{MAC} \times G_{diversity} \quad (\text{Equation 1})$$

where C_o represents fundamental single-channel ALOHA performance (approximately 50 nodes under evaluated conditions), $N_{channels}$ denotes hardware parallelism, η_{MAC} captures protocol enhancements (typically 2-5 times the protocol-level reference), and $G_{diversity}$ reflects spatial redundancy (typically 1.5-3.0 times). This framework is intended as an interpretive tool for normalising reported capacities to stress-test-derived multipliers, thereby supporting more consistent comparison across studies by separating hardware scaling from algorithmic improvements.

2. RELATED WORK

Recent LoRaWAN scalability research has pursued enhanced MAC protocols, multi-channel architectures, and mobility optimisation, yet a theoretical gap persists between pure ALOHA theory (Abramson 1970) and LoRaWAN practice. This gap has been noted implicitly in survey literature (Jouhari et al. 2023), yet not addressed through systematic single-channel evaluation under worst-case assumptions.

Enhanced protocols attempt to overcome ALOHA limitations through coordination mechanisms, reporting substantial gains that often represent 2 to 5 times improvement over baseline performance. Alahmadi et al. (2024) propose adaptive time-slotted schemes achieving significant throughput gains, Li et al. (2024) introduce grouped bit-slot protocols supporting 50 to 1000 devices, and Benserir et al. (2025) demonstrate successive interference cancellation. Whilst these approaches report substantial gains, they fundamentally alter MAC behaviour through coordination and signal processing unavailable in standard implementations, preventing isolation and direct quantification of intrinsic ALOHA limitations. Consequently, enhancement multipliers relative to baseline ALOHA remain unquantified, limiting objective assessment of whether observed gains justify implementation complexity.

Architectural enhancements through multi-channel and multi-gateway designs address scalability through hardware parallelism. Almuhaya et al. (2024, 2025) propose zone-based and multi-band designs supporting thousands of nodes through spatial and spectral diversity, whilst Yu et al. (2023) model multi-gateway throughput optimisation. However, these architectures distribute contention across spatially and spectrally separated resources, masking per-channel dynamics and preventing identification of worst-case limits representing fundamental protocol constraints.

Energy modelling advances have been substantial but predominantly treat energy as isolated from reliability dynamics under varying contention. Maudet et al. (2021), Loh et al. (2024), and Lorincz et al. (2025) develop detailed consumption models deriving primarily from low-density conditions. When collision rates exceed 30 to 40%, confirmed mode energy escalates through retries whilst reliability simultaneously diminishes, creating dynamics unexplored in existing literature.

Mobile deployments receive increasing attention driven by applications requiring spatial dynamics. Anwar et al. (2021), Gorrela et al. (2026), and Neto et al. (2024, 2025) develop adaptive data rate optimisations for mobile environments. However, it remains empirically unvalidated whether mobility-induced signal strength fluctuations matter when collision probability dominates packet loss mechanisms.

Collectively, existing work improves system capacity but does not establish a lower-bound reference point against which such improvements can be normalised. This study addresses these gaps through systematic single-channel evaluation spanning densities from 1 to 200 nodes, mobility from 0 to 72 km/h, and both confirmed and unconfirmed transmission modes. Confirmed mode employs acknowledgement-based retransmission with up to 8 retry attempts, providing

reliability enhancement at the cost of increased energy consumption and network overhead. Unconfirmed mode transmits once without acknowledgement or retries, minimising energy and airtime usage but sacrificing delivery guarantees. We intentionally adopt worst-case assumptions, defined here as the combination of single-channel operation, absence of capture effect, and lack of spatial diversity. Single-channel operation isolates per-channel ALOHA constraints, the collision model assumes perfect orthogonality where transmissions on different SFs do not interfere whilst same-spreading-factor temporal overlap causes mutual failure (excluding capture effects, which represents a conservative approximation wherein any temporal overlap causes mutual failure), periodic traffic represents steady-state behaviour, and single-gateway configuration eliminates spatial diversity benefits. This abstraction is analogous to normalising multi-channel systems to per-channel contention, enabling fair comparison across heterogeneous deployments. Across these directions, improvements are demonstrated at the system level. However, the absence of a normalised reference point prevents quantification of how much gain arises from protocol innovation versus resource scaling.

Recent literature has increasingly explored extending LoRaWAN beyond static, low-rate telemetry toward non-traditional, high-throughput applications, including low-resolution image transmission, camera-based sensing and large-file transfer (Jebail et al. 2018; Zhang et al. 2021; Chaparro et al. 2022; Keshmiri et al. 2025; Kim et al. 2024). Because LoRaWAN strictly limits payload size, streaming such data requires slicing large files into dozens of fragmented packets transmitted in rapid succession (Sánchez Gómez et al. 2020; Muñoz Lara et al. 2022). Existing studies that evaluate image fragmentation largely under isolated, single-node or low-density conditions (Croce et al. 2020) overlook the mutual interference generated when multiple active devices simultaneously stream chunked payloads or continuous mobility updates, an interaction that has been directly observed between periodic scalar traffic and image transmission in hybrid multimedia deployments (Ta et al. 2024). Parallel work on single-channel ALOHA saturation and collision-driven retransmission in LoRaWAN (Polonelli et al. 2019; Farhad et al. 2020; Petroni and Biagi 2022; Wu et al. 2023) establishes the protocol-level mechanisms that this study extends to a high-transmission-interval regime. Unlike prior work that evaluates image transfer or mobility support largely in isolation and under comparatively low contention, this study establishes the single-channel stress ceiling that applies once such high-transmission-interval workloads are deployed at scale.

3. METHODOLOGY

We employ discrete-event simulation under intentionally conservative collision-model assumptions combined with a deliberately elevated traffic-demand stress test, providing worst-case degradation estimates that characterise a stress ceiling rather than a sizing target for typical network planning. Simulation configurations and parameters are fully specified to enable reproducibility. Figure 1 illustrates the network topology employed throughout this study.

2 km × 2 km deployment area (50 nodes)




Figure 1. Experiment network topology with a single-channel gateway at centre

3.1 Experimental Design

The study employs a three-phase experimental design. Phase 1 applies a full factorial across five density levels, 1, 10, 50, 100 and 200 nodes, and five mobility conditions, a static baseline of 0 m/s plus four non-zero speeds of 0.5, 5.0, 10.0 and 20.0 m/s corresponding to approximately 1.8, 18, 36 and 72 km/h under the Random Waypoint model. Combining these with the two transmission modes of confirmed and unconfirmed yields $5 \times 5 \times 2$, or 50 scenarios. Confirmed mode employs acknowledgement-based retransmission with up to 8 retry attempts, the maximum number of transmissions permitted under the LoRaWAN 1.0.3 specification, while unconfirmed mode transmits once without acknowledgement. Controlled parameters include a 2 × 2 km deployment area with a central gateway providing direct access without routing, which eliminates link loss due to topology changes, together with a nominal 5-second periodic transmission interval subject to per-transmission duty-cycle compliance. The physical layer uses the AS923 MHz frequency band with 125 kHz channel bandwidth and the standard LoRaWAN adaptive data rate (ADR) algorithm. Each scenario simulates a 500-second window with a 1-second warmup discarded, yielding approximately 100 transmitted packets per node per scenario at the nominal 5-second interval. The 5-second interval is the application-layer request rather than guaranteed wire-level spacing. Each transmission remains within the 1% duty-cycle limit, with excess demand deferred until compliant, as illustrated in Figure 2, while confirmed-mode retransmissions extend this delay. At 50 nodes, requested demand reaches approximately 1.98%, roughly twice the compliant budget, consistent with the mechanism reported by Haxhibeqiri et al. (2017). We deliberately exceed this margin whilst maintaining duty-cycle compliance in order to characterise network behaviour when compliant scheduling becomes an active constraint. This deliberately

aggressive transmission interval represents a conservative periodic-traffic proxy for the aggregate channel demand created when multiple active nodes simultaneously transmit fragmented image chunks or frequent real-time mobility updates, capturing the sustained high-intensity channel load such applications generate without simulating their specific bursty, on-off arrival pattern, and it represents the stress scenario examined throughout this paper.




Figure 2. Duty-Cycle Compliant Transmission Scheduling with Retransmissions

Motivated by Phase 1's observation of a sharp PDR decline between 50 and 100 nodes under confirmed mode, Phase 2 investigates this range using confirmed mode only, holding mobility fixed at 5 m/s and adding density levels at 60, 70, 80, and 90 nodes, yielding 4 additional scenarios. Phase 3 further narrows the search into the 50 to 60 node transition zone, again confirmed mode only at 5 m/s, adding density levels at 53, 55, 56, 57, 58, and 59 nodes, yielding 6 additional scenarios. Combined, the three phases span 15 density levels and 60 scenarios in total, and summing node counts across all 15 density levels yields exactly 999 node-level records.

3.2 Simulation Framework

Simulations use the FLoRa framework (Slabicki et al. 2018) implemented in OMNeT++ version 5.6.2 (Varga 2010), configured to follow LoRaWAN Specification v1.0.3 ADR and retransmission mechanisms. The physical layer employs the log-distance path loss model:

$$PL(d) = PL(d_0) + 10n\log_{10}(d/d_0) + X_\sigma \quad (\text{Equation 2})$$

with path loss exponent $n = 2.75$, shadowing standard deviation $\sigma = 3.57$ dB, and reference distance $d_0 = 40$ metres following González-Palacio et al. (2023). Gateway sensitivity follows Semtech SX1301 specifications ranging from -137 dBm (SF12) to -124 dBm (SF7). The collision model implements perfect SF orthogonality with ideal collision detection, wherein transmissions on different SFs do not interfere whilst same-spreading-factor temporal overlap causes mutual failure. This assumption aligns with prior analytical models and intentionally removes inter-SF interference to isolate same-SF contention effects, ensuring that observed degradation is attributable to temporal collision rather than cross-SF coupling. The model excludes capture

effects, which are 10 to 30% probable in real systems, representing a conservative approximation that establishes lower-bound performance.

3.3 Metrics and Statistical Analysis

The evaluation employs several key performance metrics to characterise network behaviour. The primary metric is PDR, defined as the ratio of successfully received packets to transmitted packets:

$$PDR = \frac{\text{Successfully received packets}}{\text{Transmitted packets}} \quad (\text{Equation 3})$$

This metric quantifies end-to-end reliability from the perspective of end devices. We additionally measure energy consumption per node (joules including transmission, reception, and idle listening) and network capacity (maximum node count achieving $PDR \geq 90\%$). Secondary metrics capture collision rate, retransmission count, and acknowledgement (ACK) success rate. The complete dataset, spanning all three phases, comprises 1,682,057 packet transmissions, comprising 1,503,719 confirmed-mode transmissions (241,267 original attempts and 1,262,452 retransmissions, plus 104,707 acknowledgements counted separately as downlink signalling) and 178,338 unconfirmed-mode transmissions. Within confirmed mode, original transmissions, retransmissions, and acknowledgements represent 15.0%, 78.5%, and 6.5% of the confirmed-mode total respectively. Collision rate is defined as final unresolved packet loss after all retry attempts are exhausted. Statistical methods employ polynomial regression for nonlinear trends, piecewise regression with breakpoint detection (Muggeo 2003), Pearson correlation with 95% confidence intervals, one-way Analysis of Variance (ANOVA) treating each density-mobility scenario as an independent observation, independent samples t-tests with Cohen's d effect sizes, and ratio analysis for trade-off efficiency. Precision of the reported packet delivery ratio and collision rate figures is quantified with a Wilson score 95% confidence interval computed from the original transmission attempts and successful receptions recorded within each scenario, and precision of the reported energy figures is quantified with a 95% confidence interval obtained through node-level bootstrap resampling. Precision of the reported packet delivery ratio and collision rate figures is quantified with a Wilson score 95% confidence interval computed from the original transmission attempts and successful receptions recorded within each scenario, and precision of the reported energy figures is quantified with a 95% confidence interval obtained through node-level bootstrap resampling.


4. RESULTS

All 60 scenarios across the three phases completed successfully, generating 1,682,057 packet transmissions (1,503,719 confirmed-mode and 178,338 unconfirmed-mode). Scenario-level descriptive statistics for the Phase 1 confirmed-mode factorial dataset, comprising 25 scenarios formed from 5 densities across 5 mobility conditions, reveal substantial heterogeneity across the evaluated conditions. Gateway PDR averages 68.85% with a standard deviation of 38.70% and a range of 11.64% to 100.00%, collision rate averages 31.17% with a standard deviation of 38.72% and a range of 0.00% to 88.41%, and energy per node averages 50.90 J with a standard


deviation of 40.08 J and a range of 13.02 J to 105.48 J. These figures reflect a sharp nonlinear transition across densities, with near-perfect delivery at low density, where PDR approaches 100% and energy remains near 13 J, collapsing to severe degradation at high density, where PDR falls to between 12% and 34% and energy rises to between 92 J and 105 J. This heterogeneity reflects genuine variation across the twenty five density and mobility conditions evaluated.

4.1 Threshold Characterisation


PDR declines nonlinearly with increasing node density whilst remaining invariant to mobility speed, as demonstrated in Figure 3. Figure 3(a) presents PDR versus density across all mobility conditions (static, 0.5, 5, 10, 20 m/s), showing that all mobility speeds converge to identical degradation trajectories. This visual independence confirms that density, rather than mobility, governs performance degradation. Figure 3(b) demonstrates the overall degradation trajectory across the full density range (1 to 200 nodes) with shaded critical zones highlighting the 50 to 60 node transition region where performance shifts from stable to degraded operation. The mobility-independence phenomenon occurs because direct gateway access eliminates routing-related link failures that would otherwise manifest during topology changes, so packet loss stems primarily from temporal collisions rather than connectivity disruption. Figure 3(c) provides a high-resolution examination of the critical transition zone, revealing detailed degradation dynamics across eight measurement points spanning 50 to 60 nodes, drawn from the Phase 2 and Phase 3 confirmed-mode drill-down conducted at a fixed mobility speed of 5 m/s.



(a) PDR versus density across five mobility conditions



(b) PDR across full density range with logistic fit and critical transition region



(c) Detailed view of 50-60 node transition region in fixed 5 m/s mobility condition

Figure 3. PDR as a function of node density demonstrating mobility independence and critical threshold.

Three-regime piecewise linear regression identifies two critical breakpoints, at 52.0 nodes (marking onset of sharp transition, PDR = 97.9%) and 75.0 nodes (marking transition to gradual decline, PDR = 47.6%). The model achieves an overall $R^2 = 0.9945$, with strong individual segment fits ($R^2 = 0.91, 0.98, 0.95$ for the flat, transition, and decline regimes respectively). The sharp-transition regime (52 to 75 nodes) exhibits degradation 48.7 times steeper than the flat regime and 7.6 times steeper than the subsequent decline regime. At the fixed 5 m/s mobility condition used for this drill-down, the transition region reveals rapid deterioration from 97.4% at 50 nodes to 77.9% at 60 nodes, corresponding to approximately 20 percentage points loss over a 10-node increase. This fixed-mobility value at 50 nodes reflects the same distinction between the fixed and mobility-averaged conditions already introduced in Section 1. Quadratic regression outperforms linear specification ($R^2 = 0.947$ versus 0.812), confirming systematic nonlinearity. By

50 nodes, PDR has already fallen to between 97.4% and 97.9% depending on which of the two conditions introduced in the Introduction is used, and in both cases the value falls just below the 98% threshold. For network planning purposes where $PDR \geq 90\%$ defines acceptable performance, the maximum sustainable node count ranges from 50 to 55 nodes per channel under these conditions. Table 1 stratifies performance across representative densities, with the 1, 10, 100 and 200 node rows averaged across the five mobility conditions evaluated in Phase 1 and the 60 node row taken from the fixed 5 m/s condition used in the Phase 2 and Phase 3 drill-down, and each figure is reported with its own 95% confidence interval using the same pooled estimator applied consistently throughout this analysis. The table demonstrates sustained near-perfect reliability above 99% at low density, a mobility-averaged threshold of 97.902% PDR, with a 95% confidence interval of 97.716% to 98.074%, and 2.098% collision at 50 nodes, and severe degradation at high density where collision saturation overwhelms the network.

Table 1: Performance Stratification Across Density Levels


Density (nodes)	PDR (%)	PDR 95% CI	Collision (%)	Collision 95% CI	Energy (J)	Energy 95% CI	Retrans/Node
1	100.000	99.227% to 100.000%	0.000	0.000% to 0.773%	13.069	13.017 to 13.174 J	0.4
10	99.980	99.885% to 99.996%	0.020	0.004% to 0.115%	14.194	14.019 to 14.374 J	9.0
50	97.900	97.716% to 98.074%	2.100	1.926% to 2.284%	31.145	30.614 to 31.684 J	138.0
60	77.922	76.853% to 78.962%	22.078	21.038% to 23.147%	58.439	56.493 to 60.366 J	346.4
100	33.902	33.486% to 34.321%	66.098	65.679% to 66.514%	91.180	89.987 to 92.342 J	594.8
200	12.377	12.173% to 12.584%	87.623	87.416% to 87.827%	104.913	104.262 to 105.547 J	699.0

Collision analysis (Figure 4) reveals that SF diversity reduces effective collision probability relative to pure ALOHA under moderate contention. For 50 nodes with offered load $G = 0.992$ (airtime = 99.157 ms, nominal 5-second transmission interval), classical ALOHA collision probability is calculated using:


$$P_{collision} = 1 - e^{-2G} \quad (\text{Equation 4})$$

which predicts a collision probability of approximately 86.2%, whereas the observed collision rate of 2.6% in the fixed 5 m/s mobility condition demonstrates an approximately 33.5-fold reduction attributable to spreading factor diversity. This behaviour is consistent with SF diversity creating

quasi-parallel collision domains. Beyond approximately 60 nodes, collision rate accelerates to 22.1%, narrowing the gap with theoretical predictions as contention forces SF concentration. At 100 nodes, observed collisions reach 64.5%, approaching conditions where SF diversity provides limited benefit.



(a) Observed collision rates versus theoretical ALOHA model across density range



(b) Detailed collision dynamics in 50-60 node transition region

Figure 4. Collision rate versus node density compared with theoretical ALOHA prediction.

4.2 Mobility Impact Analysis

One-way ANOVA, treating each density-mobility scenario as an independent observation ($n = 25$: 5 densities $\times$ 5 mobility conditions), reveals no significant effect of mobility on PDR ($F(4,20) = 0.0005$, $p \approx 1.0$) or energy consumption ($F(4,20) = 0.0004$, $p = 1.0$). Mean PDR varies by less than 1 percentage point across the 0 to 72 km/h range (static: 69.27%, 20 m/s: 68.35%), whilst energy consumption remains constant at approximately 50 to 51 J across all mobility conditions.

This negligible mobility impact is attributable to the model structure: a single gateway with full-area coverage (−137 dBm sensitivity across the 2×2 km deployment), absence of routing, log-distance path loss with static (non-fading) shadowing, no Doppler modelling, and ADR averaging over 20 packets together leave little mechanism through which mobility could affect PDR or energy under collision-dominated conditions. Maximum deviation from the static baseline across all speeds is approximately 0.92 percentage points, well within measurement noise.

This behaviour is consistent with ALOHA's position-independent collision dynamics, wherein temporal overlap determines failure regardless of spatial configuration. When collision rates exceed 20%, packet losses stem predominantly from contention rather than insufficient signal strength, rendering mobility-induced variations (typically ±3 dB) negligible. Additionally, ADR averaging over 20 packets (approximately 100 seconds, or 1.7 minutes, at the 5-second transmission interval) filters short-term mobility effects. With direct gateway access and no multi-hop routing, topology changes do not cause link failures, and therefore mobility has no substantial effect on connectivity.




Figure 5. Impact of mobility on network performance.

4.3 Energy-Reliability Trade-off Dynamics

Confirmed transmission mode provides substantial reliability improvement at threshold but exhibits rapid efficiency collapse beyond this point, ultimately inverting entirely at extreme density (Figure 6). Transmission mode comparison reveals three distinct regimes as density increases. At 50 nodes, confirmed mode achieves 97.9% PDR versus unconfirmed mode's 64.6% PDR, a +33.3 percentage point advantage ($t(8) = 184.79$, $p < 0.000001$, $d = 116.9$), requiring 139% energy overhead. This overhead represents the additional energy consumed by confirmed mode compared with unconfirmed mode at 50 nodes, calculated as:

$$Energy_{overhead} = \frac{(Energy_{confirmed} - Energy_{unconfirmed})}{Energy_{unconfirmed}} \times 100\% \quad (\text{Equation 5})$$

This is driven primarily by retransmissions, since 18.127 J of the 18.128 J energy difference is attributable to retransmission energy alone, given that base transmission, reception and idle

energy are essentially identical between modes at this density. Trade-off efficiency at threshold equals 0.545 J per percentage point of PDR improvement.

For scenarios beyond 50 nodes, efficiency deteriorates substantially. At 100 nodes, confirmed mode achieves 33.902% PDR compared with unconfirmed mode's 32.311% PDR, a near-convergence with only a +1.652 percentage point advantage, whilst energy overhead has risen to 600.5%. Trade-off efficiency correspondingly collapses to 47.3 J per percentage point, an 87-fold deterioration from the threshold value.

At 200 nodes, the trade-off inverts entirely. Confirmed mode achieves 12.377% PDR versus unconfirmed mode's 16.156% PDR, a -3.750 percentage point disadvantage, whilst energy overhead reaches 706.0%. Applying Equation 5's efficiency ratio to a negative PDR gain yields -24.5 J per percentage point, meaning confirmed mode is now spending energy to achieve a worse outcome than the simpler unconfirmed scheme, a genuinely negative return. This demonstrates that retransmission mechanisms become actively counterproductive under extreme high-contention conditions, where retransmissions themselves exacerbate collision probability rather than resolving it.






Figure 6. Energy consumption and PDR for confirmed and unconfirmed modes.

Retransmission dynamics (Figure 7) reveal growth from 0.0 retransmissions per node at 1 node to 695.9 at 200 nodes (averaging 7.1 retries per packet, within the 8-retry limit) for a fixed mobility speed of 5 m/s. A strong correlation between collision rate and retransmissions ($R^2 = 0.909$, $p < 0.001$) confirms that collision-driven failures trigger cascading retry attempts, creating self-reinforcing degradation. ACK success rate parallels uplink PDR, maintaining 100% at low density (1 to 10 nodes) but declining to 97.4% at 50 nodes and 12.6% at 200 nodes, demonstrating bidirectional contention saturation.



(a) Average retransmissions per packet versus node density.



(b) Relationship between collision rate and retransmissions

Figure 7. Retransmission dynamics and collision interaction.

4.4 Statistical Precision of Reported Estimates

Every packet delivery ratio and collision rate figure reported in this paper is accompanied by a Wilson score 95% confidence interval, computed from the original transmission attempts and successful receptions pooled across the relevant scenario or set of scenarios, a method chosen for its established accuracy for proportions near the 0% and 100% boundaries where the more common normal approximation is known to perform poorly. Every reported energy figure is accompanied by a 95% confidence interval obtained through node-level bootstrap resampling with ten thousand iterations, pooling the per-node energy values across the five Phase 1 mobility conditions where a figure is a mobility average, or drawn directly from the fixed-mobility scenario where a figure is reported at the fixed 5 m/s condition used in the Phase 2 and Phase 3 drill-down. At the fifty node threshold, the mobility-averaged packet delivery ratio of 97.902% carries a 95%

confidence interval of 97.716% to 98.074%, based on 24,694 original transmission attempts pooled across the five Phase 1 mobility conditions, a pooling justified directly by the null mobility effect established in Section 4.2. At the fixed 5 m/s condition used for the Phase 2 and Phase 3 drill-down, packet delivery ratio at fifty nodes is 97.429%, with a 95% confidence interval of 96.949% to 97.835% based on 4,939 attempts, falling to 77.926% at sixty nodes, with a 95% confidence interval of 76.853% to 78.962% based on 5,939 attempts. The mobility-averaged mean energy per node is 13.069 J at 1 node, 14.194 J at 10 nodes, 31.145 J at 50 nodes, 91.180 J at 100 nodes and 104.913 J at 200 nodes, each accompanied by its node-level bootstrap 95% confidence interval reported in Table 1, while the fixed 5 m/s condition gives 32.303 J at fifty nodes, with a 95% confidence interval of 31.228 to 33.388 J, and 58.439 J at sixty nodes, with a 95% confidence interval of 56.493 to 60.366 J. The narrowness of these intervals reflects the substantial packet and node count generated within each evaluated configuration, and the mobility impact analysis in Section 4.2, which varies node speed across five conditions from 0 to 72 km/h while holding density fixed, provides a complementary and independent check, since packet delivery ratio and energy consumption both varied by less than one percentage point and one joule respectively across that range, indicating that the reported thresholds are stable under substantial variation in the stochastic mobility input.

5. DISCUSSION

5.1 Mechanistic Interpretation

The observed threshold extends classical ALOHA principles whilst revealing important deviations attributable to SF diversity. With $G = 0.992$ at 50 nodes (airtime = 99.157 ms, 5-second transmission interval, fixed 5 m/s mobility condition), pure ALOHA predicts an 86.2% collision probability ($1 - e^{(-2G)}$), yet the observed collision rate is merely 2.6%, an approximately 33.5-fold reduction enabled by orthogonal SFs effectively partitioning the single physical channel into quasi-independent logical channels. ADR dynamically allocates SFs based on measured link quality, effectively partitioning offered load. Under moderate contention (≤ 50 nodes), SF distribution remains relatively balanced, maximising orthogonality benefits. Beyond approximately 60 nodes, collision acceleration narrows the gap with theoretical predictions as high contention drives ADR to concentrate traffic on higher SFs, reducing effective parallelism.

The 50 to 55 node range represents the optimal operating point for single-channel LoRaWAN under these conditions, maintaining $PDR \geq 90\%$ whilst maximising network utilisation. Beyond 55 nodes, PDR falls below this threshold, even though the paper's broader 50 to 60 node critical zone, defined by the sharp-transition regression segment, extends further. Beyond this threshold, SF diversity benefits become insufficient to prevent collision cascade, marking the transition from stable to degraded operation. Figure 3(a) provides visual confirmation that this threshold operates identically across all mobility speeds, with all trajectories converging onto the same degradation curve regardless of whether nodes are static or moving at 72 km/h. Beyond threshold, retransmission cascade amplifies baseline collision through positive feedback, representing contention amplification analogous to thrashing in shared-resource systems.

The negligible mobility impact is a direct consequence of the deployment conditions this study is designed to represent, and it demonstrates that once density-driven collision dominates the channel, mobility becomes structurally irrelevant to reliability. A single gateway with full-area coverage across the 2×2 km deployment (sensitivity of −137 dBm) ensures every node remains within range regardless of movement, the absence of routing removes the multi-hop link failures through which mobility typically degrades performance, and log-distance path loss with static shadowing isolates collision as the governing failure mechanism. This is consistent with the paper's broader single-channel stress-test methodology, which deliberately controls for confounding factors to isolate protocol-level effects. The result is a practically useful finding for network designers evaluating dense single-gateway deployments. Below the density threshold identified in this study, mobility-aware optimisation offers little benefit, since collision dynamics, not channel variation, govern reliability. Extending this stress-test framework to conditions with fast fading and Doppler effects is a natural direction for follow-up work and would establish whether density-driven collision dominance continues to outweigh mobility effects under more complex channel conditions.

The energy-reliability trade-off itself exhibits a critical transition mirroring the PDR threshold. Below 50 nodes, confirmed mode's retransmission mechanism operates as designed, with a modest energy cost purchasing a large reliability gain. Between 50 and 100 nodes, retransmissions increasingly compete with new traffic for channel access, and the marginal reliability gain per unit energy collapses by roughly two orders of magnitude. Beyond 100 nodes, the mechanism becomes actively counterproductive, with confirmed mode's own retry traffic the dominant contributor to channel saturation, producing worse delivery outcomes than the simpler unconfirmed scheme despite far higher energy expenditure. This finding suggests that confirmed-mode acknowledgement schemes require an admission-control or adaptive-retry mechanism to remain beneficial as density approaches and exceeds the single-channel capacity limit identified in this study. The implication is particularly acute for applications that stream fragmented payloads, such as chunked image transfer, under the periodic-traffic proxy described in Section 3.1. For illustration, if an image is divided into approximately fifty fragments for transmission, the loss of even a small number of those fragments to retransmission-induced congestion can prevent successful reconstruction of the complete image (Kim et al. 2024), motivating the conclusion that standard LoRaWAN confirmed mode requires complementary fragmentation-aware MAC scheduling or forward error correction to support such high-throughput streaming reliably (Coutaud et al. 2020).

5.2 Capacity Decomposition Framework

Figure 8 represents the conceptual framework for LoRaWAN capacity decomposition, illustrating how total system capacity emerges from the multiplicative contributions of the single-channel stress-test capacity ceiling, hardware channel parallelism, MAC protocol enhancements, and spatial gateway diversity. These findings establish a conceptual distinction between protocol-level capacity and system-level capacity. Protocol capacity represents performance achievable under standard MAC with minimal hardware configuration. Our findings indicate approximately 50 nodes per channel under the evaluated single-channel configuration, though this benefits from SF diversity. The observed collision reduction at threshold demonstrates that even single-channel

LoRaWAN outperforms pure ALOHA through quasi-orthogonal modulation schemes. However, this benefit diminishes beyond threshold as contention forces SF concentration.

System capacity can be approximated through multiplicative enhancements formalised in Equation 1. This decomposition is conceptual and intended as an interpretive framework rather than a strict analytical model. For illustration, taking the stress-test reference $C_0 \approx 50$ nodes, an eight-channel gateway ($N_{channels} = 8$), a protocol enhancement factor of $\eta_{MAC} = 3.0$, and multi-gateway spatial diversity of $G_{diversity} = 2.0$, the framework yields $C_{sys} \approx 50 \times 8 \times 3.0 \times 2.0 = 2400$ nodes.




Figure 8. Conceptual framework for LoRaWAN capacity decomposition, illustrating contributions of the single-channel stress-test capacity ceiling, channel parallelism, MAC enhancements, and spatial diversity.

Unlike recent work by Almuhaaya et al. (2024, 2025) and Yu et al. (2024), which evaluated scalability under mixed multi-channel and multi-gateway conditions, this work isolates single-channel ALOHA behaviour under controlled assumptions. This approach aligns with recent multi-channel ALOHA analyses that decompose performance by channel resources rather than treating scalability as a single quantity (Zanella et al. 2025). This decomposition reconciles apparently contradictory capacity claims by explicitly attributing gains to hardware resources versus

algorithmic improvements, challenging commonly held assumptions that scalability improvements primarily reflect protocol-level innovation. The framework suggests that a substantial portion of multi-channel gains can be attributed to hardware parallelism distributing contention across independent resources, though empirical multi-channel comparison would be required to validate this interpretation definitively.

5.3 Threshold Sensitivity and Generalisation

The scenario-specific threshold at 50 to 60 nodes depends on five primary factors that modulate the precise location of this critical transition point. First, transmission interval directly determines offered load, with the threshold scaling approximately linearly as intervals increase or decrease. Second, airtime distribution modulates collision window duration, whereby low-path-loss scenarios may increase the threshold by approximately 20 to 40% due to reduced transmission durations at lower SFs. Third, the capture effect, which occurs with 10 to 30% probability in real systems, may increase the threshold by approximately 20 to 40% beyond the conservative estimates presented here, though this phenomenon was intentionally excluded from our model to establish worst-case bounds. Fourth, bursty event-driven traffic patterns, such as the back-to-back fragment transmission involved in chunked image transfer, may reduce the threshold by approximately 30 to 50% relative to periodic transmission, as temporal clustering of transmissions intensifies collision probability during active periods. This estimate reflects the general ALOHA burstiness literature and is identified as a candidate for direct empirical validation using an event-driven traffic generator. Fifth, this study demonstrates that duty-cycle enforcement itself can become a binding constraint on delivery once requested demand exceeds the regulatory budget, operating alongside and independently of collision-based limits. The 50 to 60 node threshold identified here directly characterises this combined effect.

These sensitivities highlight that the identified 50-node threshold represents a conservative lower bound specific to the evaluated parameter combinations rather than a universal constant. The threshold should be interpreted as a scenario-bounded reference point that network designers must adjust based on their specific operational conditions, traffic patterns, and regulatory constraints. Empirical validation using real-world deployments remains essential to confirm the generality of these findings beyond simulation assumptions.

5.4 Practical Implications

The stress scenario evaluated establishes a worst-case reference point for single-channel gateway sizing. Up to around 50 nodes, PDR remains at 95% or higher. Beyond this point it falls sharply as requested demand approaches the duty-cycle budget, reaching around 78% by 60 nodes. This gives network designers, particularly those running bursty or high-rate applications that trade duty-cycle margin for lower latency, a practical ceiling for assessing collision and retransmission risks in dense deployments. For multi-channel gateways under comparable stress conditions, capacity can be estimated using the approximation $50 \times N_{\text{channels}} \times 0.90$, accounting for imperfect channel orthogonality and residual inter-channel interference. Under this model, eight-channel gateways yield approximately 360 nodes at $\text{PDR} \geq 90\%$, a sevenfold increase over the single-channel stress-test reference. Enhanced MAC protocols incorporating

time-slotted access, coordinated transmission scheduling, or successive interference cancellation may provide additional gains of two to five times the single-channel stress-test reference capacity, though such improvements require validation against the established single-channel reference point. Establishing this single-channel stress ceiling, at approximately 50 nodes under the periodic high-transmission-interval conditions evaluated here, therefore provides a conservative reference benchmark for designing and evaluating future MAC enhancements, such as time-slotted image-fragment scheduling, mobility-aware slot allocation and adaptive chunk-retry mechanisms (Sánchez Gómez et al. 2020; Muñoz Lara et al. 2022; Coutaud et al. 2020), for high-throughput and mobile LPWAN applications.

Transmission mode selection should adapt strategically to network density conditions. At densities of 50 nodes per channel or below, confirmed mode provides cost-effective reliability enhancement with an efficiency of 0.545 joules per percentage point of PDR improvement, representing a favourable trade-off for applications requiring delivery guarantees. In the intermediate density range of 50 to 100 nodes per channel, hybrid strategies that selectively employ confirmed mode for high-priority traffic whilst using unconfirmed mode for routine telemetry may optimise overall system efficiency. Beyond 100 nodes per channel, confirmed mode should generally be avoided given the efficiency collapse to 47.3 joules per percentage point at 100 nodes, worsening to a genuinely negative return of -24.5 J per percentage point by 200 nodes, where retransmission overhead exceeds any reliability benefit and retransmission mechanisms actively contribute to network congestion rather than alleviating it. Traffic engineering that maintains densities below approximately 50 nodes per channel (corresponding to offered load $G \approx 1.0$ under the evaluated conditions of 5-second intervals and 99 ms airtime) provides a conservative safety margin, keeping the network in the stable pre-threshold operating regime.

Research priorities should align with the dominant performance-limiting mechanisms at different density regimes. Under collision-dominated conditions exceeding 50 nodes per channel, research effort may benefit from rebalancing toward collision reduction mechanisms, including intelligent admission control that regulates network entry, load-aware spreading factor allocation that dynamically adjusts SF distribution to balance contention, and time-slotted access protocols that coordinate transmission timing to eliminate random collisions.

5.5 Limitations

The ideal collision model employed in this study excludes capture effects and therefore represents a conservative worst-case scenario. Collision rate throughout this study is defined as final unresolved packet loss after all retry attempts are exhausted, rather than the incidence of any single colliding transmission attempt. Retransmissions in confirmed mode can therefore recover from individual collisions, meaning the reported collision rate is a conservative (lower) estimate of true channel-level contention. Real hardware implementations achieve capture effect probability of 10 to 30% under favourable power imbalance conditions, potentially sustaining node densities 20 to 40% higher than the thresholds identified here. This conservative modelling approach, combined with the deliberately pushed traffic-demand stress test, was intentionally adopted to establish worst-case degradation estimates that define a clear stress ceiling. Capture effects and

lighter duty-cycle margins are known to improve outcomes beyond this ceiling, marking a concrete direction for empirical follow-up validation.

Results apply specifically to single or limited multi-channel gateway configurations with 2 to 4 channels, to the stress scenario established in Section 3.1, to suburban and rural environments with moderate path loss characteristics, and to standard LoRaWAN 1.0.x protocol implementations without proprietary enhancements. Results may not generalise to several important deployment scenarios, including enhanced MAC protocols that fundamentally alter medium access behaviour, bursty traffic patterns with event-driven transmission clustering such as the chunked image or file transfer that motivates this study, dense urban environments with high path loss and multipath propagation, wide-area deployments spanning tens of kilometres with heterogeneous propagation conditions, or multi-gateway cooperative reception systems that exploit spatial diversity through packet combining or selection diversity. The periodic-traffic model evaluated here provides a conservative proxy for the aggregate channel demand of these bursty applications, and implementing and validating an event-driven traffic generator against this baseline is identified as a specific and scoped priority for future empirical work.

The findings derive from validated discrete-event simulation under controlled parameter configurations, providing systematic characterisation of fundamental performance limits whilst abstracting implementation-specific details and environmental variability. Confidence intervals for every reported figure, detailed in Section 4.4, quantify the precision of each estimate given the substantial packet count generated within its evaluated configuration, and the mobility impact analysis in Section 4.2 provides an independent indication that the reported thresholds hold under substantial variation in the stochastic mobility input. Extending this analysis with topology-level replication across independently seeded node placements and mobility traces at the critical fifty to sixty node transition zone is identified as a natural next step for future empirical work, offering a further and complementary measure of robustness alongside the confidence intervals already reported here.

6. CONCLUSION

This study establishes a fundamental worst-case stress-test reference for single-channel LoRaWAN networks through controlled analysis of 1,682,057 packet transmissions across a three-phase, 60-scenario design conducted under a deliberately pushed traffic-demand stress test, revealing a scenario-specific critical threshold at 50 to 60 nodes under this stress condition, where performance transitions from a stable, mobility-averaged 97.9% PDR at 50 nodes to a markedly degraded 77.9% PDR at 60 nodes under the fixed 5 m/s condition used for the Phase 2 and Phase 3 drill-down, with further decline to 33.96% by 100 nodes. Beyond this threshold, confirmed mode's reliability advantage collapses and ultimately inverts. At 200 nodes, confirmed mode delivers worse packet delivery than unconfirmed mode (12.4% vs. 16.2%) while consuming approximately eight times more energy, demonstrating a fundamental limit of retransmission-based reliability under extreme contention. The findings demonstrate that density-driven collision dominates network behaviour, with post-threshold degradation approximately 49 times steeper than pre-threshold conditions, whilst node mobility exhibits negligible impact ($\eta^2 < 0.001$) under the collision-dominated, full-coverage single-gateway conditions evaluated. Energy-reliability

trade-offs collapse beyond threshold, with confirmed transmission mode efficiency degrading from 0.545 J per percentage point at threshold to -24.5 J per percentage point at 200 nodes, yielding negative reliability returns despite 706% energy overhead. The proposed capacity decomposition framework (Equation 1) establishes a conceptual basis for normalising heterogeneous scalability claims by separating protocol-level capacity (~50 nodes per channel) from system-level enhancements through hardware parallelism, MAC improvements, and spatial diversity. These conservative lower-bound estimates, derived under worst-case single-channel ALOHA assumptions excluding capture effects, provide network designers with scenario-bounded reference points that scale with transmission interval, propagation environment, and traffic patterns, whilst highlighting that substantial portions of reported multi-channel capacity gains may be attributable to hardware resource distribution rather than protocol innovation alone.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

During the preparation of this work, the author(s) used Google NotebookLM (Gemini) and Claude (Anthropic) in order to review and improve language, grammar, wording and clarity. After using these tools/services, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

Data Availability

The simulation configurations, network-level outputs, time-series data, and transmission logs for all 60 scenarios evaluated in this study are publicly available at the GitHub repository: <https://github.com/izzahns/lorawan-stress-test-confirmed-unconfirmed>.

Funding Sources

This work was supported by Universiti Malaysia Sarawak under Mentor Grant (UNI/F08/MENTOR/86070/2023).

Ethics, Consent to Participate, and Consent to Publish Declarations

Not applicable.

Acknowledgements

This work was supported by Universiti Malaysia Sarawak under Mentor Grant (UNI/F08/MENTOR/86070/2023). The authors express their gratitude to the open-source OMNeT++ community for providing the network simulation framework used in this evaluation.

REFERENCES

1. Abramson N (1970) The ALOHA system: another alternative for computer communications. In: Proceedings of the Fall Joint Computer Conference (AFIPS). ACM, pp 281–285

2. Alahmadi H, Bouabdallah F, Al-Dubai A, Ghaleb B (2024) A novel autonomous adaptive frame size for time-slotted LoRa MAC protocol. *IEEE Trans Ind Inform* 20(10):12284–12293
3. Almuahaya M, Al-Hadhrami T, Brown DJ, Qasem SN (2025) ZBMG-LoRa: a novel zone-based multi-gateway approach towards scalable LoRaWANs for Internet of Things. *Sensors* 25(17):5457. <https://doi.org/10.3390/s25175457>
4. Almuahaya M, Al-Hadhrami T, Kaiwartya O, Brown DJ (2024) MBMD-LoRa: scalable LoRaWAN for Internet of Things - a multi-band multi-data rate approach. *Lecture Notes in Computer Science*. Springer, pp 54–67
5. Anwar K, Rahman T, Zeb A, Khan I, Zareei M, Vargas-Rosales C (2021) RM-ADR: resource management adaptive data rate for mobile application in LoRaWAN. *Sensors* 21(23):7980
6. Buang IN, Zen K, Junaini SN (2025) Optimising connectivity and energy: the future of LoRaWAN routing protocols for mobile IoT applications. *Appl Model Simul* 9:67–91
7. Benserir N, Etiabi Y, Sabir E, Amhoud E, Elbiaze H, Diallo A (2025) IRSA over spreading factors for spatio-temporal SIC in scalable LoRaWAN IoT networks. *IEEE Symposium on Computers and Communications (ISCC)*. IEEE, pp 1–6
8. Bor MC, Rodig U, Voigt T, Alonso JM (2016) Do LoRa low-power wide-area networks scale? *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*. ACM, pp 59–67
9. Chaparro F, Pérez M, Mendez D (2022) A communication framework for image transmission through LPWAN technology. *Electronics* 11(11):1764. <https://doi.org/10.3390/electronics11111764>
10. Chasserat, L., Accettura, N., & Berthou, P. (2024). LoRaSync: Energy efficient synchronization for scalable LoRaWAN. *Transactions on Emerging Telecommunications Technologies*, 35. <https://doi.org/10.1002/ett.4940>.
11. Coutaud U, Heusse M, Tourancheau B (2020) Fragmentation and forward error correction for LoRaWAN small MTU networks. In: *Proceedings of the 16th International Conference on Mobility, Sensing and Networking (MSN)*. IEEE, pp 289–294. <https://doi.org/10.5555/3400306.3400363>
12. Croce D, Gucciardo M, Mangione S, Santaromita G, Tinnirello I (2020) LoRa technology demystified, from link behaviour to cell-level performance. *IEEE Trans Wirel Commun* 19(2):822–834. <https://doi.org/10.1109/twc.2019.2948872>
13. Farhad A, Kim DH, Pyun JY (2020) Resource allocation to massive Internet of Things in LoRaWANs. *Sensors* 20(9):2645. <https://doi.org/10.3390/s20092645>
14. Figueiredo A, Luís M, Zúquete A (2022) A fair channel hopping scheme for LoRa networks with multiple single-channel gateways. *Sensors* 22(14):5260
15. Fragkopoulos M, Panagiotakis S, Kostakis M, Markakis E, Astyrakakis N, Malamos A (2023) Experimental assessment of common crucial factors that affect LoRaWAN performance on suburban and rural area deployments. *Sensors* 23(3):1316
16. Gamage, A., Liando, J., Gu, C., Tan, R., Li, M., & Seller, O. (2022). LMAC: Efficient Carrier-Sense Multiple Access for LoRa. *ACM Transactions on Sensor Networks*, 19, 1 - 27. <https://doi.org/10.1145/3564530>.

17. González-Palacio M, Tobón-Vallejo D, Sepúlveda-Cano L, et al. (2023) Machine-learning-based combined path loss and shadowing model in LoRaWAN for energy efficiency enhancement. *IEEE Internet Things J* 10(12):10725–10739
18. Gorrela A, Choudhury N, Calafate C, Jiang W, Jamshed M, Kaushik A (2026) Adaptive data rate optimization in mobile and dense LoRaWAN IoT environments. *IEEE Internet Things J* 13(8):16566 to 16577. <https://doi.org/10.1109/JIOT.2026.3659938>
19. Haxhibeqiri J, Van den Abeele F, Moerman I, Hoebeke J (2017) LoRa scalability: a simulation model based on interference measurements. *Sensors* 17(6):1193.
20. Jebri EM, Sali A, Ismail A, Rasid MFA (2018) Overcoming limitations of LoRa physical layer in image transmission. *Sensors* 18(10):3257. <https://doi.org/10.3390/s18103257>
21. Jouhari M, Amhoud E, Saeed N, Alouini MS (2023) A survey on scalable LoRaWAN for massive IoT: recent advances, potentials, and challenges. *IEEE Commun Surv Tutor* 25(3):1841–1876
22. Keshmiri H, Rahman GME, Wahid KA (2025) LoRa resource allocation algorithm for higher data rates. *Sensors* 25(2):518. <https://doi.org/10.3390/s25020518>
23. Khan, M., , H., Aamir, S., & Jin, Y. (2021). Optimizing the Performance of Pure ALOHA for LoRa-Based ESL. *Sensors (Basel, Switzerland)*, 21. <https://doi.org/10.3390/s21155060>.
24. Kim J, Jenkins J, Seol J, Kwak M (2022) Extending data transmission in the multi-hop LoRa network. *Issues Inf Syst* 23(3). https://doi.org/10.48009/3_iis_2022_124
25. Kim J, Kwak M, Jenkins J (2024) Reliable image transmission over LoRa networks. *Issues Inf Syst* 25(1). https://doi.org/10.48009/1_iis_2024_117
26. Li X, Xu J, Li R, Jia L, You J (2024) Advancing performance in LoRaWAN networks: the circular region grouped bit-slot LoRa MAC protocol. *Electronics* 13(3):621
27. Loh F, Raffek S, Geissler S, Hossfeld T (2024) Plan the access? Generic hardware independent energy consumption and efficiency model for different LoRaWAN channel access approaches. *IEEE Internet Things J* 11(11):19678–19693
28. LoRa Alliance (2017) LoRaWAN® Specification v1.0.3. LoRa Alliance Technical Committee. LoRa Alliance. Available at: <https://resources.lora-alliance.org/document/lorawan-specification-v1-0-3>
29. LoRa Alliance (2023) LoRa Alliance Annual Report 2023: Global IoT Deployments
30. Lorincz J, Kusačić M, Čusto E, Blažević Z (2025) A comprehensive multiple linear regression modeling and analysis of LoRa user device energy consumption. *J Sens Actuator Netw* 15(1):5. <https://doi.org/10.3390/jsan15010005>
31. Maudet S, Andrieux G, Chevillon R, Diouris J (2021) Refined node energy consumption modeling in a LoRaWAN network. *Sensors* 21(19):6398
32. Muggeo VMR (2003) Estimating regression models with unknown break-points. *Stat Med* 22(19):3055–3071
33. Muñoz Lara R, Hidalgo JS, Canales F, Dujovne D, Céspedes Umaña S (2022) SCHC over LoRaWAN efficiency, evaluation and experimental performance of packet fragmentation. *Sensors* 22(4):1531. <https://doi.org/10.3390/s22041531>
34. Neto AG, Silva TC, Abreu PA, Veloso AFS, Mendes LDO, Soares A, Júnior JGR (2025) Evaluating a mobility-aware ADR scheme in urban and suburban LoRaWAN environments. *Ann Telecommun* 80(9-10):745–758

35. Neto AG, Silva TC, Abreu PA, Veloso AFS, De O. Mendes L, Reis JG (2024) Addressing mobility challenges in LoRaWAN through adaptive data rate: a statistical median-based approach. *FiCloud. IEEE*, pp 330–337
36. Petroni A, Biagi M (2022) Interference mitigation and decoding through gateway diversity in LoRaWAN. *IEEE Trans Wirel Commun.* <https://doi.org/10.1109/twc.2022.3172629>
37. Polonelli T, Brunelli D, Marzocchi A, Benini L (2019) Slotted ALOHA on LoRaWAN, design, analysis and deployment. *Sensors* 19(4):838. <https://doi.org/10.3390/s19040838>
38. Sánchez Gómez J, Gallego-Madrid J, Sanchez-Iborra R, Santa J, Gómez-Skarmeta A (2020) Impact of SCHC compression and fragmentation in LPWAN, a case study with LoRaWAN. *Sensors* 20(1):280. <https://doi.org/10.3390/s20010280>
39. Slabicki M, PremSankar G, Di Francesco M (2018) Adaptive configuration of LoRa networks for dense IoT deployments. *IEEE/IFIP Network Operations and Management Symposium (NOMS)*. IEEE, pp 1–9
40. Ta QH, Ta V, Nguyen T, Oh H (2024) A WSN and LoRa hybrid multimedia transmission protocol for scalar data and image transmission. *Sensors* 24(24):8165. <https://doi.org/10.3390/s24248165>
41. Varga A (2010) OMNeT++. In: Wehrle K, Güneş M, Gross J (eds) *Modeling and Tools for Network Simulation*. Springer, Berlin, pp 35–59
42. Wu W, Wang H, Cheng Z (2023) ReLoRaWAN, reliable data delivery in LoRaWAN networks with multiple gateways. *Ad Hoc Netw* 147:103203. <https://doi.org/10.1016/j.adhoc.2023.103203>
43. Yu B, Bao Y, Huang Y, Zhan W, Liu P (2023) Modeling and throughput optimization of multi-gateway LoRaWAN. *IEEE Access* 11:142940–142950
44. Yu S, Xia X, Hou N, Zheng Y, Gu T (2024) Revolutionizing LoRa gateway with XGate: scalable concurrent transmission across massive logical channels. *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. ACM, pp 1–15
45. Zanella A, Baiocchi A, Rolich A (2025) Analysis of a persistent multi-channel slotted ALOHA protocol without acknowledgment. *2025 International Conference on Computing, Networking and Communications (ICNC)*, pp 635–641
46. Zhang JYQ, Yeung BL, Wong JCY, Cheung RCC, Lam AHF (2021) LoRaWAN-based camera with CIRA compression and image recovery algorithm. *2021 IEEE 7th World Forum on Internet of Things (WF-IoT)*. IEEE, pp 136–141. <https://doi.org/10.1109/wf-iot51360.2021.9595674>