Fast Analytical Tuning of Repetition-Based MAC for IoT in Sparse LEO Constellations

Benoit Coeugnet
Kinéis & CITI UR3720
Toulouse & Villeurbanne, France
bcoeugnet@kineis.com

André-Luc Beylot
IRIT, Univ. Toulouse, INP, CNRS
Toulouse, France
andre-luc.beylot@enseeiht.fr

Vincent Deslandes
Kinéis
Toulouse, France
vdeslandes@kineis.com

Juan A. Fraire
Inria, INSA Lyon, CITI UR3720
Villeurbanne, France
juan.fraire@inria.fr

***Abstract*—The growing demand for IoT deployments beyond cellular coverage is driving adoption of low-data-rate direct-to-satellite (DtS-IoT) communication, with sparse Low Earth Orbit (LEO) constellations emerging as a cost-effective way to achieve global coverage. In such systems, downlink bandwidth may be too limited to acknowledge every uplink packet, making non-slotted ALOHA with blind retransmissions the dominant MAC paradigm, and turning the choice of repetition count and spacing into a critical design decision. Existing analytical treatments either rely on idealized visibility formulas or require time-consuming event-driven simulation, neither of which is well-suited to per-device, per-location deployment tuning. We propose an analytical model, grounded in renewal theory and discrete-time Markov chains, that uses empirical visibility distributions derived from real Two-Line Element (TLE) sets to compute the full reception-probability distribution under both blind (*BlindTx*) and beacon-assisted (*SatDetect*) MAC schemes. For a given deployment, the model returns optimal retransmission parameters in seconds, versus minutes for event-driven simulation. Validated against the FLoRaSat simulator (RMSE $< 3\%$), which is itself calibrated to Kinéis on-field measurements, the model shows that *SatDetect* achieves almost $4\times$ better energy efficiency than *BlindTx* at low latitudes while meeting a 95% reception target.**

***Index Terms*—MAC Layer, LEO Constellation, Markov Chains**

I. Introduction

By the end of the decade, tens of billions of IoT devices are projected to be connected, with a fast-growing share deployed in maritime zones, polar regions, deserts, and remote agricultural land — areas where terrestrial networks do not reach and which collectively span the vast majority of the Earth's surface. For these deployments, low-data-rate direct-to-satellite (DtS-IoT) communication via sparse Low Earth Orbit (LEO) constellations has emerged as the dominant cost-effective option. Kinéis, a French Mobile Satellite Service (MSS) operator, exemplifies the trend: its 25-satellite constellation, launched in 2025, already provides global IoT services to tens of thousands of devices, with traffic dominated by uplink reports from energy-constrained nodes.

These systems must cope with two structural constraints [1]: first, coverage is intermittent and strongly latitude-dependent (availability ranges from a few percent for a one-satellite system up to roughly 40% at the Equator for Kinéis). Second, the per-satellite downlink budget is too limited to acknowledge every uplink, and a large fraction of deployed devices have no downlink capability at all. The dominant MAC paradigm in this regime is therefore non-slotted ALOHA with blind retransmissions, and reliability is bought entirely through redundancy: M repetitions of each message spaced by an interval T [2], where every additional repetition costs energy. Jointly optimizing M and T on a per-device, per-location basis is the critical MAC design challenge. Yet existing analytical treatments [3], [4] rely on idealized theoretical visibility formulas that miss the non-uniform, latitude-dependent coverage of real Walker Star deployments, while event-driven simulators [5], [6] reproduce reality at the cost of minutes per scenario — too slow for per-deployment tuning. The gap this paper closes sits in the middle: an analytical model fed with *empirical* visibility statistics, fast enough to act as a deployment and protocol design tool.

Accordingly, this paper (i) derives ON/OFF visibility distributions from real Kinéis Two-Line Element (TLE) data via Simplified General Perturbations 4 (SGP4) propagation, fitted with two-component log-normal mixtures validated by Kolmogorov–Smirnov testing; (ii) develops, from renewal theory and a discrete-time Markov chain, a modular framework that accommodates any fitted visibility law and supports both *BlindTx* (blind transmission) and *SatDetect* (beacon-assisted) MAC schemes; (iii) formulates parameter selection as a constrained optimization solved in under a second per configuration, enabling per-location deployment tuning; and (iv) validates the model against the FLoRaSat simulator [5] with an RMSE below 3% across configurations and latitudes, showing that *SatDetect* cuts equivalent energy consumption by up to $4\times$ at low latitudes while *BlindTx* remains competitive at high latitudes. Section II covers the related work; Section III covers the visibility model, the Markov chain, and the optimization; Section IV covers the validation and numerical results; Section V concludes.

II. Background and Related Work

We use *message* to denote the unit of information sent from a device to the service center, and *repetition* to denote a MAC-layer retransmission of that message.

A. Constellation

Several operators run sparse LEO IoT systems today (Kinéis, Astrocast and Sateliot among them) with constellation sizes from a handful to a few tens of satellites; Kinéis serves

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

as the running exemplar and the basis for all validation results, while the fitting methodology is exercised on all three in §IV-A. The Kinéis constellation is a near-polar Walker Star of 25 nanosatellites at 650 km altitude in five planes of five, of which 20 currently carry the IoT mission. Coverage is global but strongly non-uniform: the planes converge towards the poles, giving near-continuous coverage at high latitudes and intermittent coverage below 40% at the Equator, further reduced by the 4-of-5 satellite selection per plane. Real orbits also deviate from the nominal pattern under perturbations and operational constraints. Given the rotational symmetry of the planes, longitude has only a marginal effect, so characterizing visibility at a set of positive latitudes suffices to cover all deployment locations.

B. MAC Schemes

BlindTx is a non-slotted repetition-based ALOHA scheme analogous to LoRaWAN Class A. When a message is scheduled for transmission, the device sends M identical copies, each separated by a fixed interval T, without any knowledge of satellite presence. Reception reliability is achieved solely through redundancy: at least one repetition must fall within a visibility window. This simplicity makes *BlindTx* well-suited to severely energy-constrained or low-complexity devices.

Non-slotted and slotted ALOHA in the DtS-IoT setting have been studied from a throughput and collision standpoint [2], [7], showing that sparse-constellation dynamics, short visibility windows, and satellite geometry all shape the achievable rate. We absorb these effects into the per-repetition success probability P_{succ} of §III-C.

SatDetect is the beacon-assisted counterpart to *BlindTx*, operating analogously to LoRaWAN Class B. It differs from *BlindTx* in two ways. First, before each scheduled transmission, the device opens a short listening window to detect a downlink beacon from a passing satellite; the uplink is sent only if a beacon is successfully received, which serves as a proxy for satellite visibility and avoids unnecessary transmissions when no satellite is overhead. Second, the scheme stops after N_{sd} successful detections, on the assumption that enough copies have been delivered. Together, these mechanisms, used operationally at Kinéis, reduce both device energy consumption and uplink-channel load.

Beacon-assisted DtS-IoT is well represented in the literature: Class B-inspired schemes synchronize ground uplinks on periodic satellite beacons [8], beacon-triggered frameworks adapt spreading-factor allocation to high satellite dynamics [9], and hybrid MACs trade energy against delay [10]. *SatDetect* differs mainly in using a Kinéis-specific beacon protocol rather than LoRaWAN Class B framing.

C. Related Work

Alvarez et al. [3] studied uplink transmission policies for LoRa-based DtS-IoT using a theoretical visibility model, providing the closest prior analytical treatment to the present work. Their model, however, assumes simplified contact distributions that do not reflect the non-uniform Walker Star coverage of a real deployment, a limitation this work directly addresses. Le et al. [4] analyzed NB-IoT Non-Terrestrial Network performance under discontinuous LEO coverage using a hybrid analytical framework, also relying on theoretical visibility assumptions.

A complementary line of work models DtS-IoT access using stochastic geometry [11], whose homogeneity assumptions break down for the sparse, near-polar Walker Star footprints studied here; ON/OFF channel modeling more broadly draws on renewal theory [12]. None of these works addresses the per-device retransmission-parameter optimization problem under the coverage distribution of a real constellation, which is the focus of the present paper.

III. System modeling

We denote by M the maximum number of repetitions per message and by T the spacing between two consecutive repetitions. The goal of the model is to compute the distribution Π_i of the number of received repetitions, for $i \in [\![0, M]\!]$, under both *BlindTx* and *SatDetect*. The system is formulated as a discrete-time Markov chain sampled at transmission instants, with transition probabilities derived from the visibility statistics of the constellation under study.

A. Visibility Modeling

The Markov-chain transitions are governed by the durations of satellite visibility (ON) and non-visibility (OFF) windows seen from a given ground location. We sample these durations by propagating the constellation's publicly available TLEs with the SGP4 model and recording all contact intervals at elevation $\geq 5^\circ$. While the static link budget permits transmissions below this angle, severe masking and shadow fading make such links practically unreliable; we therefore adopt 5° as a conservative hard threshold, a standard modeling simplification of what is, in reality, a gradual degradation in propagation.

To be plugged into the analytical framework developed below, these empirical CDFs are fitted with parametric continuous distributions, selecting the best fit for each latitude via the Kolmogorov–Smirnov (KS) statistic. The full comparison across candidate laws, constellations, and latitudes is reported in §IV-A; the two-component log-normal mixture is consistently selected and is used as the running example. If both ON and OFF durations are taken to be exponential, memorylessness makes the residual life exponential as well, and the chain admits a clean closed form; that special case fits acceptably at mid-latitudes, but the KS statistic exceeds 0.3 at low and high latitudes (cf. §IV-A), which motivates the general treatment below. The formulation is therefore developed for an arbitrary continuous random variable L.

B. General Distributions

We now drop the exponential assumption and develop the framework for an arbitrary continuous random variable L. The construction below is analytical where possible and numerical otherwise, depending on the closed-form tractability of L.

1) Method: Rather than tracking continuous time, we evaluate how many transmissions fall into each successive visibility/non-visibility window using Probability Mass Functions (PMFs) that give the probability of exactly j transmissions in a window. The resulting discrete-time Markov chain has states

$$S = (n, v, i, k), \tag{1}$$

with n the period index, $v \in \{0, 1\}$ the visibility status, $i \leq M$ the attempts made so far and $k \leq i$ the number that fell within an ON window. There are therefore $(M+1)(M+2) = O(M^2)$ states per iteration, and only non-zero entries are stored.

2) Residual life of the first window: Devices lack knowledge of satellite orbital mechanics and therefore initiate transmissions independently of the satellite's presence. The first sampled window is thus a *partial* window whose duration follows the residual-life distribution of the underlying period law L, with CDF F_L and mean $\mathbb{E}[L]$. The Probability Density Function (PDF) and CDF of the residual life are

$$f_Z(x) = \frac{1 - F_L(x)}{\mathbb{E}[L]}, \tag{2}$$

$$F_Z(x) = \int_0^x f_Z(t)\, dt = 1 - \frac{1}{\mathbb{E}[L]} \int_x^\infty \big(1 - F_L(t)\big) dt. \tag{3}$$

The probability that the device wakes up in an ON state is weighted by the expected ON/OFF durations:

$$\tau = \frac{\mathbb{E}[\mathrm{ON}]}{\mathbb{E}[\mathrm{ON}] + \mathbb{E}[\mathrm{OFF}]}. \tag{4}$$

3) Residual life of subsequent windows: Handling the transition between periods is delicate: the time offset of the first transmission inside the current period depends on the timestamp of the last transmission in the previous period. Since transmissions are spaced by T, this offset lies in $[0, T]$, and we approximate it by a uniform random variable

$$Y \sim \mathcal{U}(0, T), \tag{5}$$

independent of the period duration L. The total effective duration before the next transition is then $V = L + Y$, with PDF given by the convolution

$$f_V(x) = \frac{1}{T} \int_0^{\min(T,x)} f_L(x - y)\, dy. \tag{6}$$

The $\mathcal{U}(0, T)$ assumption is not exact for arbitrary L, but the validation in §IV-B confirms that the resulting error remains below 3% across the tested configurations. The CDF of V then discretizes each subsequent window into probabilities of containing j transmissions, effectively decoupling the timing across consecutive periods.

4) Chain evaluation: The chain is evaluated iteratively over successive windows, in the base form (i.e., assuming $P_{\mathrm{succ}} = 1$), with the binomial weighting of §III-C applied at the end: from the residual-life PMFs ϕ (ON) and ψ (OFF), each iteration advances every state by the number of transmissions falling in the next window, drawn from α (ON) or β (OFF), until the mass still holding $i < M$ attempts drops below ϵ_{th}; $\gamma_r(M)$ is then read off the absorbing states. The PMFs are obtained by discretizing the corresponding CDFs on a uniform grid of slot width T up to a maximum window index $J_{\max} = 40$, the residual tail being assigned to the boundary bin so that each PMF sums to 1; this retains more than 99.99% of the probability mass for the parameters considered.

C. Weighting of the reception rate

The chain above assumes that any repetition falling in a visibility window is received successfully. In practice, fading, interference, and intra-fleet collisions all reduce that probability. We fold these effects into a single per-repetition success probability P_{succ}, conditioned on coverage.

Let $\gamma_r(M)$ denote the distribution of the number of *covered* repetitions (from the chain evaluation above), and $\xi_k(M)$ the distribution of *effective* successes after M repetitions. Given r covered repetitions, the number of successes K follows a binomial distribution $K \sim \mathcal{B}(r, P_{\mathrm{succ}})$ — assuming successes are i.i.d. across repetitions and uncorrelated with the visibility state. By the law of total probability,

$$\xi_k(M) = \sum_{r=k}^{M} \binom{r}{k} P_{\mathrm{succ}}^k (1 - P_{\mathrm{succ}})^{r-k} \gamma_r(M). \tag{7}$$

The parameter P_{succ} looks simple but encapsulates considerable physical complexity: it can be obtained from a PHY-level simulator when the expected network load is known, or fitted directly from field measurements. For the numerical results in §IV-C, we further decompose P_{succ} into a line-of-sight (geometric masking) factor P_{los} and an in-visibility reception factor P_r, so that the cost split between the two MAC schemes becomes explicit. Modeling endogenous collisions (load-dependent P_{succ}) is left outside the present scope.

D. Key Performance Indicators

1) Reception Rate: R_{rx} is the ratio of messages received to messages sent. In our model,

$$R_{\mathrm{rx}} = 1 - \xi_0(M). \tag{8}$$

2) Delay: We define the delay D as the time from the first transmission attempt to the *first* successful reception at any satellite The satellite-to-ground-station segment is outside the scope of this analysis. Because D is meaningful only conditional on eventual delivery, we report the conditional expectation

$$\mathbb{E}[D] \;=\; \frac{\sum_{m=1}^{M} (m-1)\, T \cdot p_m}{\sum_{m=1}^{M} p_m}, \tag{9}$$

where p_m is the probability that the first success happens at attempt m.

3) Energy Consumption: Let E_{tx} be the energy consumed by a single transmission, and E_l be the average energy consumed during a listening period. Based on typical IoT device specifications, opening a listening slot consumes, on average, about 10 times less energy than sending a message, yielding the relationship $E_{\mathrm{tx}} \approx 10 E_l$.

For the *BlindTx* scheme, the total energy consumption increases linearly with the number of transmissions, N_{tx}. However, the *SatDetect* scheme introduces N_l listening periods. To provide a fair comparison, we normalize the energy consumption of *SatDetect* into an equivalent number of transmissions, denoted as N_{eq}:

$$N_{\mathrm{eq}} = N_{\mathrm{tx}} + \frac{1}{10} N_l. \tag{10}$$

E. Parameter Optimization

We formulate parameter selection as a constrained optimization over a discrete set Θ ($\Theta_{\mathrm{BlindTx}} = \{M\}$, $\Theta_{\mathrm{SatDetect}} = \{M, N_{\mathrm{sd}}\}$) and the continuous spacing T, restricted to $M \in [\![2, 12]\!]$, $N_{\mathrm{sd}} \in [\![1, M]\!]$ and $T \leq T_{\max} = 1500\,\mathrm{s}$ — ranges in which our verified solutions consistently sit. Depending on the deployment constraint, we either maximize $R_{\mathrm{rx}}(\Theta, T)$ subject to $N_{\mathrm{eq}}(\Theta, T) \leq N_{\max}$, where a target battery life fixes the maximum equivalent consumption $N_{\max}$, or minimize $N_{\mathrm{eq}}(\Theta, T)$ subject to $R_{\mathrm{rx}}(\Theta, T) \geq R_{\mathrm{target}}$. Both forms exploit the smooth dependence of R_{rx} on T to solve the inner problem

$$T^{\star}(\Theta) = \arg \max_{T \leq T_{\max}} R_{\mathrm{rx}}(\Theta, T) \tag{11}$$

per discrete configuration using Brent's method, which converges well under a second. The grid contains 77 *SatDetect* and 11 *BlindTx* configurations; sweeping all 88 completes in about one minute on a laptop (Core Ultra 7 265H, 32 GB RAM), easily parallelizable across Θ.

IV. Evaluation

A. Visibility and non-visibility distributions

We first assess the fitting methodology across three real sparse LEO constellations with varying densities: Kinéis (20 IoT-payload satellites), Astrocast (20 satellites), and Sateliot (5 satellites). Among the candidate laws evaluated — Uniform, Weibull, Gamma, single log-normal — the two-component log-normal mixture consistently provided the best fit as measured by the Kolmogorov–Smirnov (KS) statistic. Excluding latitudes $\geq 80°$, which lie outside the scope due to near-continuous coverage, the KS statistic remains below 0.08 across all latitudes examined for Kinéis and Astrocast (most values in the range 0.02–0.05); Sateliot, due to its much sparser constellation, reaches up to 0.1, with most values between 0.05 and 0.08. These low and consistent KS values confirm that the two-component log-normal mixture is a suitable model for the ON/OFF durations across a broad range of constellation densities and latitudes.

B. Validation

FLoRaSat [5] is an event-driven, open-source[1] system-level simulator built on OMNeT++; its Kinéis-specific hardware and channel models were validated against on-field Kinéis measurements, which is why we use it as ground truth here. It is also far more compute-intensive than the analytical model: a single configuration can take up to a minute to simulate, and each scenario requires hundreds of runs to reach statistical significance, whereas the analytical model converges in well under a second. We therefore use a focused set of simulations to cross-check the model rather than to sweep configuration space. We anchor the validation at latitude $0°$, where coverage is sparsest and the reception rate lowest — the regime in which numerical disagreement with the simulator is most likely to surface — and additionally check two contrasting latitudes ($40°$ and $80°$) to confirm the agreement holds as coverage grows.

1) Reception rate across the T sweep: For each latitude, we use FLoRaSat to simulate the reception rate across a wide range of T values and compare against the analytical model, using $P_{\mathrm{succ}} = 0.665$ for *BlindTx* (a typical value for the Kinéis constellation). Across the four configurations $M \in \{2, 4, 6, 8\}$, the RMSE between simulated and analytical points stays below 2% at all three latitudes; at $80°$ both model and simulator saturate near 100% reception for all T, which compresses the absolute differences and partly explains the lower RMSE there.

2) Full distribution at the optimal $T^{\star}$: Evaluated at the optimal spacing $T^{\star}$ returned by the optimization, the RMSE on the *full* distribution $\xi_i(M)$ of received repetitions stays below 3% across the board, with the tightest agreement at latitude $0°$. The model therefore captures not only the mean reception rate but also the full repetition distribution.

C. Results

We now exercise the optimization on Kinéis. For both schemes, we use as a baseline that the probability that a device is masked by the environment is, on average, around 30%, and that the probability of reception, even with direct visibility between the satellite and the device, is only about 95%. Therefore, we denote $P_{\mathrm{los}} = 0.7$ and $P_r = 0.95$. For *BlindTx*, we then have $P_{\mathrm{succ}} = P_{\mathrm{los}} \cdot P_r = 0.665$, and for *SatDetect*, $P_{\mathrm{det}} = P_{\mathrm{los}} = 0.7$ (masking is symmetric between uplink and downlink) and $P_{\mathrm{succ}} = P_r = 0.95$. Therefore, in both cases, the effective probability of reception is the same, but the split of consumption greatly differs, which is the main advantage of the *SatDetect* scheme.

1) Impact of N_{sd}: The detection quota N_{sd} trades reception against energy in a non-trivial way, since it interacts with both P_{det} and P_{succ}. Raising N_{sd} increases consumption (additional listening + transmission slots) but yields only marginal reception gains once the cumulative success probability under coverage exceeds $\sim$0.9. The repetition budget M behaves similarly: a larger M improves the odds of catching a visibility window but also increases the number of listening slots, hence the energy cost. To reach $R_{\mathrm{target}} = 0.95$, the best Kinéis configuration is $(M, N_{\mathrm{sd}}) = (9, 2)$, with $R^{\star}_{\mathrm{rx}} = 0.95$, $N^{\star}_{\mathrm{eq}} = 2.38$, $T^{\star} = 650\,\mathrm{s}$, and $D^{\star} = 1314\,\mathrm{s}$ — a delay expected in a region where coverage is below 40%.

2) BlindTx vs SatDetect: BlindTx trades cleanly: each additional repetition increases N_{eq} by exactly 1, so its energy footprint scales linearly with M. Matching $R_{\mathrm{target}} = 0.95$ with *BlindTx* requires $M \geq 9$, yielding $N^{\star}_{\mathrm{eq}} = 9$ — almost $3\times$ the equivalent *SatDetect* configuration — and a delay of $D^{\star} = 1317\,\mathrm{s}$, the same as for *SatDetect*. At the opposite extreme, *BlindTx* with $M = 3$ keeps $N^{\star}_{\mathrm{eq}} = 3$ but reaches only $R^{\star}_{\mathrm{rx}} = 0.65$ at $T^{\star} = 610\,\mathrm{s}$; its low delay $D^{\star} = 520\,\mathrm{s}$ is misleading, since it averages only over the (few) messages that do get through.

[1] https://gitlab.inria.fr/jfraire/florasat

The empirical ON/OFF distributions are nearly identical between 0° and 40°, then shift towards longer ON periods — up to visibility windows of several tens of minutes — as latitude rises towards 80°.

TABLE I
BEST CONFIGURATIONS UNDER EACH CONSTRAINT

Lat	Config (M, N_{sd})	$T^\star$(s)	$R^\star_{\mathrm{rx}}$	$D^\star$(s)	$N^\star_{\mathrm{eq}}$
Reception target $R_{\mathrm{target}} = 0.95$					
10	SD (9, 2)	660	0.95	1330	2.38
	Btx (9,−)	668	0.95	1333	9
70	SD (4, 2)	1100	0.96	721	2.06
	Btx (4,−)	1010	0.96	727	4
Energy budget $N_{\mathrm{eq}} \leq 2.5$					
10	SD (10, 2)	687	0.96	1420	2.46
	Btx (2,−)	690	0.5	330	2.00
70	SD (12, 2)	970	0.99	1625	2.34
	Btx (2,−)	980	0.80	314	2.00

Table I (top) reports the optimal configurations reaching $R_{\mathrm{target}} = 0.95$ at the extremes of the studied range; the intermediate latitudes 30° and 50° interpolate monotonically between them. As coverage increases with latitude, the required energy budget drops and the gap between *SatDetect* and *BlindTx* narrows: at 10°, *SatDetect* consumes almost $4\times$ less than *BlindTx*, but at 70° the ratio shrinks to 2. Two cross-cutting observations stand out. First, as the latitude rises, M decreases while $T^\star$ grows — yet the delay does not grow proportionally, meaning that most messages are now received on the *first* repetition; this is confirmed by $\mathrm{Btx}(1, -)$ at 70°, which alone delivers $R^\star_{\mathrm{rx}} = 0.57$. Second, $N_{\mathrm{sd}} = 2$ is consistently the sweet spot across latitudes, because $P_{\mathrm{r}} = 0.95$ leaves a very small residual loss after two successful detections.

Table I (bottom) reports the dual problem: best configurations under an energy budget $N_{\mathrm{eq}} \leq 2.5$ (chosen to bracket the consumption of a baseline two-repetition *BlindTx*). At low latitudes, *SatDetect* delivers up to 92% more messages than *BlindTx* while staying within the same budget. At high latitudes, the reception-rate gap shrinks to $\sim 24\%$ while *SatDetect* consumes about 17% more energy, so the best choice becomes constraint-dependent. Pushing M beyond 12 yields only a marginal reception gain ($\sim 0.1\%$ per additional repetition) at substantial delay cost, which is why we set $M_{\max} = 12$.

Overall, *SatDetect* delivers a strong return on investment whenever beacon detection and uplink reception are both reasonably reliable.

V. CONCLUSION

We presented an analytical model for tuning the retransmission parameters of non-slotted ALOHA in sparse LEO IoT constellations, grounded in renewal theory and a discrete-time Markov chain fed with empirical ON/OFF visibility distributions. This enables, for the first time, deployment-time MAC tuning rather than design-time guesswork.

Two takeaways follow. *First*, the beacon-assisted *SatDetect* scheme cuts equivalent energy consumption by almost $4\times$ versus *BlindTx* at low latitudes while meeting a 95% reception target, or delivers up to 92% more messages at a fixed energy budget. *Second*, at high latitudes, where coverage is near-continuous, *BlindTx* becomes competitive and is preferable on simplicity and reliability grounds. This is a useful design rule for heterogeneous fleets.

The most consequential next step is to model collisions endogenously rather than absorb them into a static P_{succ}: that would let the optimization respond to fleet-level load and would naturally support *schedule randomization* (per-device T jitter) as a first-class MAC-design variable, instead of a deployment-time workaround. Additional avenues include integrating PHY-level effects (fading, shadowing) into a load-aware P_{succ}.

ACKNOWLEDGMENT

This work was supported by Kinéis through a CIFRE doctoral fellowship, by the DORSAL-IoT project funded by Inria Equipes Associées (EA) and the STIC-AmSud DORSAL project (AMSUD240007), by the French National Research Agency (ANR) under grant ANR-22-CE25-0014-01, and by the PEPR Networks of the Future under grant ANR-24-PEFT-0001.

REFERENCES

[1] G. M. Capez, S. Henn, J. A. Fraire, and R. Garello, "Sparse Satellite Constellation Design for Global and Regional Direct-to-Satellite IoT Services," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 3786–3801, Oct. 2022.

[2] S. Herrería-Alonso, M. Rodríguez-Pérez, R. F. Rodríguez-Rubio, and F. Pérez-Fontán, "Improving Uplink Scalability of LoRa-Based Direct-to-Satellite IoT Networks," *IEEE Internet of Things Journal*, vol. 11, no. 7, pp. 12 526–12 535, Apr. 2024.

[3] G. Álvarez, J. A. Fraire, K. A. Hassan, S. Céspedes, and D. Pesch, "Uplink Transmission Policies for LoRa-Based Direct-to-Satellite IoT," *IEEE Access*, vol. 10, pp. 72 687–72 701, 2022.

[4] T. D. Le, M. Stusek, D. Moltchanov, P. Palurik, P. Masek, and J. Hosek, "A Hybrid Analytical Framework for NB-IoT NTN Performance Under Discontinuous LEO Coverage," *IEEE Open Journal of the Communications Society*, vol. 7, pp. 1793–1811, 2026.

[5] A. Y. Choquenaira-Florez, B. Coeugnet, R. Ohs, J. A. Fraire, and H. Rivano, "FLoRaSat 2: Simulating Cross-Linked Direct-to-Satellite IoT LEO Constellations with Joint Access and Routing Evaluation," *International Journal of Satellite Communications and Networking*, 2025. [Online]. Available: https://hal.science/hal-05424077

[6] A. Traspadini, M. Zorzi, and M. Giordani, "Performance Evaluation of LoRa for IoT Applications in Non-Terrestrial Networks via ns-3," in *GLOBECOM 2025 - 2025 IEEE Global Communications Conference*, Dec. 2025, pp. 3249–3254, arXiv:2509.02811 [cs.NI].

[7] E. Testi and E. Paolini, "Packet Collision Probability of Direct-to-Satellite IoT Systems," *IEEE Internet of Things Journal*, vol. 12, no. 2, pp. 1843–1855, Jan. 2025.

[8] M. A. Mojamed, "Beacon-Based Uplink Transmission for LoRaWAN Direct to LEO Satellite Internet of Things," *International journal of Computer Networks & Communications*, vol. 16, no. 5, pp. 43–58, Sep. 2024.

[9] C. Zhang, H. Peng, Y. Ji, T. Hong, and G. Zhang, "Adaptive Resource Optimization for LoRa-Enabled LEO Satellite IoT System in High-Dynamic Environments," *Sensors*, vol. 25, no. 11, p. 3318, Jan. 2025.

[10] H. Tao, Z. Zhiwei, W. Chuming, D. Xiaojin, and Z. Gengxin, "LoRa Energy-Delay-Tradeoff Hybrid MAC Protocol for LEO Satellite IoT Networks," *IEEE Sensors Journal*, vol. 25, no. 13, pp. 25 982–26 000, Jul. 2025.

[11] N. Okati, T. Riihonen, D. Korpi, I. Angervuori, and R. Wichman, "Downlink Coverage and Rate Analysis of Low Earth Orbit Satellite Constellations Using Stochastic Geometry," *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 5120–5134, Aug. 2020.

[12] D. R. Cox, *Renewal Theory, Methuen and Co*, ltd.,london ed., 1962.