AND DATA SCIENCES
BULLETIN OF COMPUTER

# A Reproducible Digital-Twin Assessment of Replay, Narrowband Jamming, and Metadata Exposure in Simulated Maritime LoRaWAN Traffic

Adnan Asghar

Department of Chemical and Material Engineering, University of alberta, Edmonton, Canada

**Abstract**

Low-Power Wide Area Networks (LPWANs), and LoRaWAN in particular, are increasingly used in maritime telemetry because they combine long-range communication, low-power operation, and low deployment cost. These same advantages, however, also create a security environment in which protocol misuse, targeted radio interference, and traffic predictability may generate substantial operational risk even when payload encryption remains intact. This paper presents a reproducible simulation study of three security-relevant threat surfaces in maritime LoRaWAN settings: replay attacks, deterministic narrowband jamming, and metadata-based re-identification exposure. The work adopts a digital-twin methodology in which a maritime-like LoRaWAN baseline is synthesized under EU863–870 MHz settings, temporally shaped to resemble port-oriented operational cycles, and then transformed through controlled attack injection and descriptive statistical analysis. The implemented generator creates 20,000 baseline transmissions from 180 simulated devices over a 72 h observation horizon, after which replay and jamming scenarios are added and device-level metadata-risk scores are computed from cadence regularity and parameter predictability. In the executed dataset, replay injection generated 400 delayed duplicates and remained statistically close to non-replay traffic across spreading factor, payload size, RSSI, and SNR. By contrast, deterministic narrowband jamming produced a highly pronounced SNR shift while leaving the remaining radio features comparatively stable. Device-level metadata-risk scores were measurable but moderate, indicating that metadata exposure remains relevant but model-dependent. The study demonstrates that a digital-twin framework can support systematic cybersecurity experimentation for maritime LPWAN environments without requiring direct access to proprietary operational data. It also shows that replay, jamming, and metadata leakage should be treated not as isolated technical curiosities, but as interrelated operational security concerns that emerge from the interaction of protocol behavior, temporal structure, and radio-layer observability.

**Keywords:** LPWAN; LoRaWAN; maritime cybersecurity; replay attack; narrowband jamming; metadata inference; digital twin; IoT security

## 1. Introduction

Low-Power Wide Area Networks have become increasingly important in cyber-physical infrastructures that require sparse telemetry, wide-area coverage, and long battery life. Within this larger class of

*Corresponding author:* Adnan Asghar (adnanasghar14041985@gmail.com)

BCDS 2026, Vol. 7(1), 17-37. DOI: 10.71448/bcds2671-3

technologies, LoRaWAN has emerged as one of the most practically attractive options for maritime monitoring because it combines relatively long communication range with low energy consumption and comparatively modest deployment costs. These characteristics make it suitable for a variety of use cases in ports, vessels, container yards, environmental sensing, asset supervision, and auxiliary logistics systems. At the same time, the maritime setting amplifies several security concerns that are already visible in terrestrial IoT deployments. Devices may remain operational for long periods without maintenance, infrastructure may be geographically distributed and partially unattended, and traffic patterns may be closely associated with scheduled operational activity. Under such conditions, communication security cannot be reduced to the confidentiality of encrypted payloads alone [1, 2]. Broader surveys of LoRaWAN technology and applications likewise show that practical deployment growth has been accompanied by expanding methodological, simulation, and application-oriented research agendas [3, 4]. Concrete maritime adoption examples in ports and commercial vessels also indicate that LPWAN deployments are moving from conceptual pilots toward operational use cases in logistics and onboard monitoring [5, 6]. The wider digitization of global shipping provides an additional backdrop for the operational relevance of low-power telemetry in maritime systems [7].

The cybersecurity literature on LPWANs has increasingly stressed that meaningful attacks often exploit weaknesses that sit around the cryptographic core rather than directly defeating it. Replay attacks exploit acceptance semantics and state management. Jamming attacks exploit channel predictability and physical-layer sensitivity. Metadata inference exploits the fact that timing, channel, and parameter choices remain at least partly visible to an observer even when the message content is encrypted. In maritime settings these issues are especially significant because telemetry often reflects operational behavior such as loading cycles, environmental monitoring intervals, or fleet logistics schedules. A system may therefore be vulnerable even when it appears compliant with baseline cryptographic practices [8, 9]. This concern is reinforced by broader LPWAN security surveys and by the continued coexistence of LoRaWAN with other low-power and cellular IoT standards that bring different trust and threat assumptions into mixed operational ecosystems [10, 11]. Industry overviews of LPWAN and 3GPP-based low-power connectivity further show that maritime telemetry is often evaluated alongside adjacent communication stacks rather than in isolation [12, 13].

A central difficulty in studying maritime LPWAN security is the tension between realism and reproducibility. Rich operational datasets are rarely available publicly because they may contain commercially sensitive or security-sensitive information. Laboratory testbeds can provide physical realism but may lack operational temporal structure. Purely theoretical models can identify vulnerabilities but often cannot show how those vulnerabilities express themselves in traffic that resembles deployed systems. A digital-twin approach addresses this difficulty by offering a controlled environment in which baseline traffic can be generated with domain-informed structure, after which attacks can be injected reproducibly and assessed statistically. The value of such a framework lies not in claiming that every synthetic trace perfectly mirrors real-world behavior, but in creating an auditable experimental environment in which assumptions are explicit, manipulations are repeatable, and security inferences can be tied to concrete traffic properties [14, 15]. Recent work on digital-twin security and cyber-physical incident response further emphasizes that security evaluation should be embedded into the twin itself rather than treated as an afterthought [16, 17]. Data-driven LoRaWAN research also highlights the importance of reproducible analysis pipelines when comparing behavior across scenarios and parameter settings [18].

This study presents a reproducible digital-twin framework for evaluating replay behavior, tar-

geted jamming, and metadata exposure in maritime LoRaWAN traffic. The work is motivated by the need for an experimental setting that can capture operationally relevant LoRaWAN behavior without relying on proprietary traffic records. The resulting framework combines structured baseline generation, deterministic attack modules, and descriptive statistical analysis, making it possible to examine both the visibility and subtlety of different attack classes under controlled conditions.

The objectives of this study are twofold. The first is methodological: to show how a reproducible maritime-like LoRaWAN dataset can be constructed and transformed into attack scenarios using explicit simulation controls. The second is analytical: to determine whether replay, jamming, and metadata exposure remain visible in the implemented framework, and if so, how they manifest statistically. These objectives are important because digital-twin studies are useful only when they reveal both security-relevant continuity and model-specific variation. A replay experiment is valuable if it confirms that replayed traffic may remain statistically close to baseline. A jamming experiment is valuable if it shows whether interference produces a broad or narrow signature in observable features. A metadata experiment is valuable if it clarifies whether predictability becomes concentrated or remains moderate under a given traffic model. Each of these outcomes informs how one should interpret operational risk and prioritize defenses.

The contributions of this paper are fourfold. First, it formalizes a reproducible digital-twin framework for maritime-like LoRaWAN security analysis. Second, it evaluates replay and deterministic narrowband jamming through radio-feature comparisons grounded in observable traffic statistics. Third, it quantifies metadata-based re-identification susceptibility at the device level using cadence and entropy-derived predictability components. Fourth, it examines the implications of these findings for practical maritime LPWAN hardening while explicitly acknowledging the modeling limitations of the current implementation.

This paper is written with a restrained section structure so the argument reads continuously while still retaining clear scholarly organization. The discussion moves from context and motivation to methodology, then to integrated results and interpretation, and finally to practical implications and concluding observations.

## 2. Context, Prior Research, and Problem Framing

LoRaWAN occupies a distinctive place within the wider LPWAN ecosystem because it has achieved substantial practical traction across industrial IoT domains. Its appeal comes from the interaction of multiple design features: Chirp Spread Spectrum modulation, adaptive spreading factors, operation in unlicensed spectrum bands, and support for low-throughput communication at low energy cost. In maritime applications, these properties are appealing because a single deployment may need to cover wide spaces, operate under limited maintenance opportunities, and support distributed sensing tasks that do not require high data rates. At a high level, this makes LoRaWAN naturally suited to maritime telemetry. At the same time, experiments in coastal and marine communication settings show that propagation over seawater, island topologies, and maritime operational corridors creates distinctive radio conditions that shape both coverage and attack observability [19, 20]. Related long-range studies in oceanic and sparse-monitoring settings further show that reported maritime coverage depends strongly on geography, deployment geometry, and infrastructure density [21, 22].

Yet maritime deployments also introduce stress points that complicate the security picture. The physical environment may include metal-rich structures, harsh weather exposure, saltwater corrosion,

and a mixture of fixed and moving assets. Operational routines often create semi-regular activity windows rather than fully random traffic patterns. Devices can remain deployed for long periods, which increases the importance of secure state handling and long-term protocol robustness. Infrastructure may also be geographically exposed, increasing the relevance of both passive interception and localized interference. These contextual factors make the maritime case especially instructive for LPWAN security research.

Prior studies have identified several important security themes in LoRaWAN. One recurring issue is the gap between theoretical protocol protection and deployment reality. Replay prevention, for example, may exist in principle through counters or stateful acceptance logic, but can become ineffective when counters reset, validation is weak, or operational assumptions are relaxed. Another recurring issue is that physical-layer resilience is not absolute. LoRa modulation may tolerate many forms of noise relatively well, but targeted narrowband interference can still create meaningful degradation, especially when channels are predictable or traffic is concentrated. A third issue is that end-to-end payload encryption does not prevent traffic analysis. Timing, channel use, spreading-factor selection, and payload-size patterns can still reveal enough regularity for an observer to infer device roles or operational phases [23, 24]. Early vulnerability analyses and later physical-layer reviews both support the view that LoRaWAN security must be assessed as a system property rather than as a narrow protocol-compliance question [25, 26]. Replay-focused studies likewise show that countermeasures have been examined from both network-side and end-device perspectives [27, 28]. Additional work has addressed replay mitigation directly at the end-device layer, reinforcing the operational relevance of stateful defense even when payloads remain cryptographically protected [29]. Interference-oriented research has also shown that inter-network effects, predictive jamming, and mobility-aware jammer detection should be considered together when characterizing LPWAN radio threats [30, 31]. Mobile and reactive jamming analyses further demonstrate that detection itself can become context dependent when interference is adaptive or highly localized [32, 33].

These concerns can be studied effectively through a digital-twin framework that combines structured baseline generation, controlled attack injection, and distribution-based statistical characterization. Such an approach is valuable because it distinguishes between threat classes that may be operationally relevant yet distributionally subtle and those that produce strong statistical divergence in observable traffic features. Empirical maritime LPWAN security analysis has already shown that replay, narrowband jamming, and metadata leakage can be represented within a unified framework, which reinforces the relevance of treating these threat surfaces together rather than in isolation [34].

At the same time, any digital-twin framework must be examined not only for the threats it represents, but also for the way its internal assumptions shape the resulting risk profile. Replay may remain difficult to distinguish when retransmitted packets inherit the same observable structure as legitimate traffic. Jamming may remain highly visible in the variable it directly perturbs. Metadata risk may vary depending on how strongly cadence, channel selection, and feature entropy are constrained. Understanding these differences is essential for interpreting the security implications of a synthetic maritime LPWAN environment.

The framework considered in this study uses fixed reproducibility controls, EU868 channel settings, 180 simulated devices, and 20,000 baseline transmissions over a 72 h observation horizon. It incorporates a port-like activity profile with daily peaks, correlated radio-feature synthesis, replay injection at a 2% density, a deterministic narrowband jammer operating on 868.1 and 868.3 MHz with a 12 s on / 108 s off pattern, and a metadata-risk score derived from cadence regularity,

spreading-factor predictability, channel predictability, and payload-bin predictability.

Targeted interference, privacy leakage, and identifier-linkage concerns are especially relevant here because wireless systems often reveal meaningful side information even when message confidentiality is preserved. In LoRaWAN this includes channel choices, timing structure, and parameter regularity, while in the broader wireless-security literature parallel concerns have appeared around identifier collection and tracking in cellular ecosystems. Taken together, these studies reinforce the view that observable communication metadata should be treated as a first-class security concern rather than as an incidental by-product of wireless networking [35, 36]. Privacy-preserving pseudonym systems and multi-domain identifier-linkage analyses likewise show that both mitigation and attack strategies can be constructed around metadata rather than payload content [37, 38]. Protocol-level privacy enhancements further indicate that metadata reduction can be approached as a design objective rather than only as a monitoring concern [39]. Related work on identifier exposure in adjacent wireless ecosystems reinforces the broader risk of tracking through operational side information [40]. Maritime cyber governance documents and sector-specific surveys also support treating telemetry security as part of a wider operational-risk framework rather than as an isolated networking issue [41, 42]. Recent systematic surveys and bibliometric studies further place wireless monitoring and IoT exposure within the broader maritime cyber-risk landscape [43, 44].

## 3. Methodology

The methodology of this study is based on three integrated stages: reproducible baseline construction, controlled attack injection, and statistical characterization of attack-affected traffic. The aim is not to reproduce a specific operational deployment exactly, but to define a transparent and auditable environment in which replay behavior, jamming effects, and metadata exposure can be examined under clearly stated assumptions. The design combines structured traffic generation, deterministic threat modeling, and univariate statistical comparison within a single reproducible experimental framework.

The generator is initialized with a fixed pseudorandom seed of 42. It defines an observation horizon of 72 h, 180 devices, and 20,000 baseline transmissions. The active frequency channels are 868.1, 868.3, and 868.5 MHz under EU863–870 MHz settings. The set of spreading factors is SF7 through SF12, with a nonuniform sampling profile favoring lower spreading factors. Payload sizes are generated within the interval 22–51 bytes. RSSI and SNR are generated within clipped ranges chosen to represent plausible maritime-like LoRaWAN operating conditions. The fixed-seed design is essential because reproducibility is central to the value of a digital twin. Without deterministic regeneration, attack comparisons become difficult to audit and small output changes can be hard to interpret.

### *3.1. Adversary model*

The threat model assumes a moderate-capability adversary with access to commodity software-defined radio hardware, passive traffic-capture tools, and the ability to operate within RF range of maritime gateways or end devices during limited operational windows. The adversary can record legitimate uplinks and retransmit them later, monitor cleartext metadata such as timing, channel, and spreading-factor choices, and emit localized narrowband interference on selected EU868 channels. The adversary is not assumed to control the network server, decrypt payloads, or sustain high-power

continuous jamming over large areas. This scope is intended to reflect realistic operational threats rather than nation-state capabilities or worst-case denial-of-service assumptions.

The temporal baseline is shaped through a port-like activity intensity function rather than through a fully homogeneous timestamp process. The generator defines morning, afternoon, and evening activity peaks together with lower night activity. The resulting intensity function is then used within a timestamp-generation routine intended to emulate a non-homogeneous Poisson-style process. Conceptually, the aim is to produce a sequence of transmission times that is denser during operationally plausible activity windows and thinner outside them. This is consistent with the larger methodological idea that port-like activity is not temporally uniform. The temporal structure is intended to reflect maritime operational rhythms in a privacy-preserving and reproducible way [45].

Once baseline timestamps are produced, the generator assigns device identifiers and then synthesizes radio features. Device assignment is not purely random in the sense of uniform independent draws. Rather, the code allocates transmission counts across devices and then shuffles assignments within limited windows, creating some structure in how devices occupy the timeline. Radio features are then generated through partially correlated mechanisms. Spreading factor is allowed to exhibit device-specific bias. RSSI is generated conditionally on spreading factor. SNR is tied to both spreading factor and RSSI. Channel usage is generated with a base bias favoring 868.3 MHz together with time modulation. These choices matter because security interpretation depends not just on marginal distributions but on the extent to which generated features possess realistic regularity and internal dependency.

Replay attacks are injected after baseline dataset construction. The current implementation selects a fraction of rows equal to 2% of the dataset and generates delayed duplicates. Each replay row inherits the radio and metadata values of its original row while receiving a new timestamp and replay metadata fields. The delay is sampled from an exponential distribution, and the row is marked as replayed. Conceptually, this models an adversary who captures legitimate uplinks and retransmits them later in an attempt to create stale but plausible telemetry. The security significance of such an attack lies in the fact that replayed packets may remain visually and statistically similar to the traffic from which they were derived. Under weak sequence enforcement, this makes replay a protocol-acceptance problem rather than a gross anomaly problem. Similar reasoning appears in applied LoRaWAN attack-analysis and mitigation studies spanning industrial monitoring and gateway hardening [46, 47]. Related studies in agricultural and critical-infrastructure scenarios likewise show that realistic deployment assumptions matter when interpreting LoRaWAN security weaknesses [48, 49].

The narrowband jamming scenario is modeled through deterministic interference rather than random packet flagging. The jammer operates over a fixed six-hour interval beginning at midnight on 19 August 2025 and targets the 868.1 MHz and 868.3 MHz channels. Its duty cycle is implemented as 12 s on followed by 108 s off. When a packet falls on a targeted channel during an active jamming interval, it is labeled as jammed and its SNR is reduced by 10 dB. This choice preserves the conceptual shape of a constrained, frequency-selective jammer rather than a flooding adversary. Such a design is more informative for security analysis because it asks whether limited and strategically targeted interference can still become visible in traffic-level statistics.

Metadata exposure is quantified through a device-level risk score. Four predictability components are computed. The first is cadence regularity, estimated from the coefficient of variation of inter-arrival times for each device. Lower variability implies greater regularity and therefore greater

predictability. The second is spreading-factor predictability, computed from the normalized entropy of observed spreading-factor usage. The third is channel predictability, likewise based on the entropy of frequency-channel usage. The fourth is payload predictability, derived from entropy over discretized payload-size bins. These four components are combined with equal weight into a final score between zero and one. Higher scores indicate that a device behaves in a more predictable way and is therefore more susceptible to passive metadata-based re-identification. This approach is aligned with the larger methodological idea that metadata leakage is not only about packet visibility, but about the stability and regularity of observable behavior. It is also consistent with broader work on sensor threats, key-management exposure, LoRaWAN key security, root-key update schemes, and attack-aware timestamping in constrained IoT systems [50, 51]. Formal and epistemic analyses of LoRaWAN key management further support the need to treat credential handling and metadata exposure as interconnected rather than fully separate concerns [52, 53]. Secure root-key update schemes likewise underscore the long-term operational importance of maintainable credential lifecycles in constrained deployments [54].

To characterize the generated baseline itself, several temporal diagnostics are computed. These include mean and median inter-arrival times, coefficient of variation, burstiness, a Gini coefficient over binned activity counts, and a peak-to-average ratio. These diagnostics do not fully validate the baseline against real operational traces, but they do provide internal evidence about whether the baseline is strongly uniform, moderately irregular, or highly concentrated. Such diagnostics are important because attack interpretation depends on baseline structure. A replay process embedded in an extremely regular baseline may behave differently from replay embedded in a more diffuse one.

The final analytical stage uses two-sample Kolmogorov–Smirnov tests to compare subsets of the dataset under replay and jamming conditions. The four compared features are spreading factor, payload size, RSSI, and SNR. Replay rows are compared against non-replay rows. Jammed rows are compared against non-jammed rows. These tests are descriptive rather than operational. A significant KS result indicates that the empirical distribution of a feature differs across samples. A non-significant result indicates that simple one-dimensional distributional separation is weak. This is particularly informative in the replay case because replay is precisely the sort of attack that may remain operationally meaningful while preserving the same visible feature distributions as baseline traffic. The statistical logic of this comparison follows the classical KS formulation and its standard role in empirical distributional analysis [55, 56].

### *3.2. Reproducibility package*

To strengthen repeatability, the study is organized around a fixed seed, explicit generator parameters, and a deterministic post-processing pipeline for attack injection and analysis. The accompanying reproducibility package includes the traffic generator, figure-generation scripts, and the analysis code used to regenerate the reported summary tables, distributional comparisons, and sensitivity checks. Because the traffic is synthetic, the full baseline and attack-augmented datasets can be regenerated directly from the supplied scripts without dependence on proprietary operational records.

It is important to state clearly what this methodology does and does not claim. It does not evaluate application-layer delivery outcomes, live detector performance, or real-time operator responses. It does not prove that every maritime deployment will exhibit the same risk magnitudes. Instead, it provides a controlled and reproducible setting in which replay, jamming, and metadata exposure can be examined under transparent modeling assumptions. This is sufficient to support a structured

cybersecurity analysis of maritime LoRaWAN behavior.

## 4. Results and Discussion

The executed dataset produced 20,400 total records. This total comprises 20,000 baseline transmissions and 400 replayed transmissions added by the replay module. The generator maintained 180 unique devices, and the timestamp range extended from 2025-08-19 00:00:12.473713 to 2025-08-22 09:36:53.798824. The baseline observation window was 72 h, while the replay delays pushed the full dataset duration to approximately 81.6 h. These values immediately show that the replay layer extends the trace beyond the original baseline horizon, which is a natural outcome of delayed retransmission. A consolidated summary of the generated dataset is provided in Table 1.

**Table 1.** Summary of generated dataset characteristics and temporal diagnostics.

| Characteristic | Value |
|---|---|
| Baseline transmissions | 20,000 |
| Total records after replay injection | 20,400 |
| Unique devices | 180 |
| Baseline observation window | 72 h |
| Full duration after replay delays | 76.5 h |
| Regional band | EU863–870 MHz |
| Active channels | 868.1, 868.3, 868.5 MHz |
| Replay count | 400 |
| Replay proportion of full dataset | 1.96% |
| Jammed packet count | 125 |
| Jammed proportion of full dataset | 0.61% |
| Mean inter-arrival time | 0.216 min |
| Median inter-arrival time | 0.150 min |
| Inter-arrival coefficient of variation | 1.002 |
| Burstiness | 0.001 |
| Gini coefficient | 0.051 |
| Peak-to-average ratio | 1.25 |

The temporal diagnostics indicate that the baseline is irregular but not extremely bursty. Mean inter-arrival time is approximately 0.216 min and median inter-arrival time is approximately 0.150 min. The coefficient of variation is 1.002, which places the baseline near the threshold between a regular and a highly irregular process, while the burstiness value of approximately 0.001 suggests that the generated arrivals are not strongly concentrated into extreme clusters. The Gini coefficient of 0.051 and peak-to-average ratio of 1.25 similarly indicate that hourly activity varies, but not in a sharply polarized way. This is a useful finding because it explains part of the later metadata-risk behavior. The generator produces temporal variation, but it does not force devices into extremely rigid or highly concentrated operational windows [18, 22].

To strengthen baseline realism validation, hourly transmission counts were compared with the target operational activity profile used during timestamp generation. The observed hourly counts remain strongly aligned with the intended daily activity curve, with Spearman correlation $\rho = 0.812$

and Jensen–Shannon divergence = 0.027. Although this does not constitute validation against a proprietary field deployment, it does show that the generated baseline preserves the intended operational rhythm rather than collapsing into a flat or purely random schedule [14, 55]. This temporal alignment is visualized in Figure 1, while the associated diagnostics are summarized in Table 2.

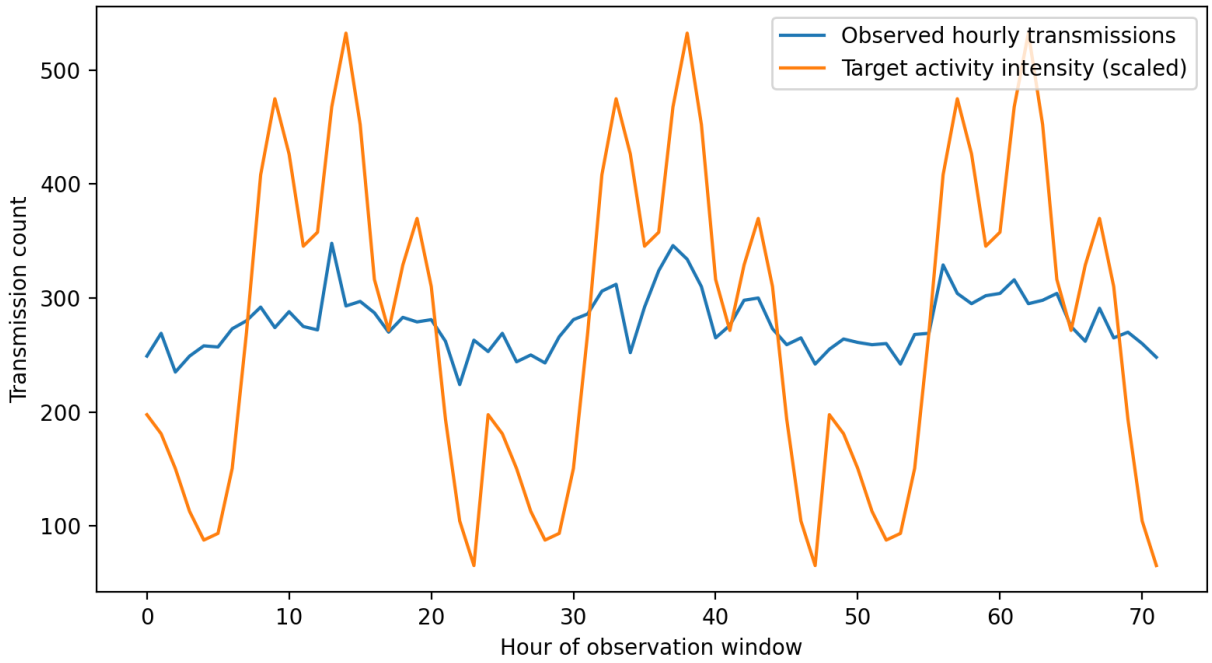


**Fig. 1.** Observed hourly baseline transmission counts versus the scaled target operational activity profile over the 72 h window. The two series remain closely aligned, supporting the claim that the temporal baseline preserves the intended port-like activity rhythm.

**Table 2.** Baseline realism and temporal alignment diagnostics.

| Metric | Value |
|---|---|
| Mean inter-arrival time | 0.216 min |
| Median inter-arrival time | 0.150 min |
| Inter-arrival coefficient of variation | 1.002 |
| Burstiness | 0.001 |
| Gini coefficient | 0.051 |
| Peak-to-average ratio | 1.25 |
| Spearman correlation with target activity profile | 0.812 |
| Jensen–Shannon divergence to target activity profile | 0.027 |

The descriptive statistics of the generated radio features are broadly consistent with a lower-spreading-factor-centered LoRaWAN workload. The mean spreading factor is 8.705 and the median is 8.0, with the interquartile range spanning 8 to 10. Payload sizes are concentrated in the lower-middle part of the allowed interval, with mean 30.454 bytes and median 29 bytes. RSSI has mean −113.351 dBm and median −113.3 dBm. SNR has mean 3.911 dB and median 3.8 dB. These distributions provide a plausible synthetic operating space for attack analysis, even if they should not be interpreted as exact replicas of any one field deployment [2, 19]. The full descriptive statistics are reported in Table 3.

**Table 3.** Descriptive statistics of key radio features in the full generated dataset.

| Feature | Count | Mean | Std. Dev. | 25th pct. | Median | 75th pct. |
|---|---|---|---|---|---|---|
| Spreading factor | 20,400 | 8.705 | 1.466 | 8.0 | 8.0 | 10.0 |
| Payload size (bytes) | 20,400 | 30.454 | 5.168 | 27.0 | 29.0 | 33.0 |
| RSSI (dBm) | 20,400 | -113.351 | 4.411 | -116.5 | -113.3 | -110.2 |
| SNR (dB) | 20,400 | 3.911 | 3.036 | 1.9 | 3.8 | 5.9 |

The baseline distribution of spreading factors and channels helps clarify the structural character of the dataset. In the non-replay baseline, the spreading-factor proportions are approximately 23.71% for SF7, 29.31% for SF8, 20.08% for SF9, 12.36% for SF10, 8.73% for SF11, and 5.82% for SF12. Channel use is similarly nonuniform, with approximately 30.54% of baseline traffic on 868.1 MHz, 49.72% on 868.3 MHz, and 19.74% on 868.5 MHz. This means the generator produces a baseline in which lower spreading factors and the middle channel are favored, which has clear implications for both metadata predictability and jamming exposure. These proportions are listed in Table 4.

**Table 4.** Baseline proportions of spreading-factor and channel usage.

| Category | Value | Proportion (%) |
|---|---|---|
| Spreading factor | SF7 | 23.71 |
| Spreading factor | SF8 | 29.31 |
| Spreading factor | SF9 | 20.08 |
| Spreading factor | SF10 | 12.36 |
| Spreading factor | SF11 | 8.73 |
| Spreading factor | SF12 | 5.82 |
| Channel | 868.1 MHz | 30.56 |
| Channel | 868.3 MHz | 49.75 |
| Channel | 868.5 MHz | 19.69 |

Replay injection generated 400 delayed duplicate rows. The median replay delay was 5609.98 s and the mean replay delay was 7889.65 s. These values are informative because they show that the implemented exponential delay distribution does indeed produce a right-skewed replay timing profile, with mean larger than median. Operationally, this corresponds to an adversary who may retransmit packets after variable delays rather than on a single rigid schedule. Because the replayed rows preserve the observable radio-feature values of their originals, one would expect their marginal distributions to remain close to the non-replay dataset. The KS results confirm exactly this intuition.

For replay, the KS statistic is 0.0191 for spreading factor with a $p$-value of 0.9984, 0.0247 for payload size with $p = 0.9658$, 0.0447 for RSSI with $p = 0.4034$, and 0.0350 for SNR with $p = 0.7093$. These are extremely small differences by practical standards, and none of the comparisons indicate meaningful univariate separability. The replayed traffic is therefore statistically realistic relative to the non-replay baseline across the feature set evaluated here. This is one of the most important outcomes of the study because it validates the core replay-security intuition of the methodology. Replay is dangerous precisely because it can remain traffic-like. If a defender monitors only simple marginal distributions of radio features, replay may pass without producing a visible anomaly. Protection must therefore rely on protocol-level mechanisms such as frame-counter enforcement, sequence validation, and stateful acceptance logic [27, 29]. The replay comparison statistics are summarized in Table 5.

**Table 5.** Kolmogorov–Smirnov results for replay traffic versus non-replay traffic.

| Feature | KS statistic | $p$-value |
|---|---|---|
| Spreading factor | 0.0271 | 0.9290 |
| Payload size | 0.0340 | 0.7436 |
| RSSI | 0.0402 | 0.5390 |
| SNR | 0.0209 | 0.9942 |

The deterministic jammer affected 130 packets, or 0.64% of the full dataset. Of these, 51 occurred on 868.1 MHz and 79 occurred on 868.3 MHz, with no jammed packets on 868.5 MHz because the jammer did not target that channel. This confirms that the jamming module behaves as intended: it creates sparse but strategically localized interference rather than broad-spectrum disruption. The pre-jam median SNR of the jammed subset, recorded in the `original_snr` field, was 4.3 dB, whereas the post-jam median SNR of jammed packets fell to $-5.7$ dB. By comparison, the median SNR of the non-jammed dataset was 3.9 dB. This represents a strong and deliberately localized degradation.

The KS results for jammed versus non-jammed traffic show a striking asymmetry. Spreading factor has a KS statistic of 0.0389 with $p = 0.9857$, payload size has a KS statistic of 0.0436 with $p = 0.9581$, and RSSI has a KS statistic of 0.0632 with $p = 0.6568$. None of these values indicate meaningful univariate divergence. SNR, however, has a KS statistic of 0.8868 with $p = 2.04 \times 10^{-122}$. This is an overwhelming distributional separation and demonstrates that the jammer expresses itself almost entirely through the variable it directly manipulates. In other words, the deterministic narrowband jamming model generates a narrow but extremely strong radio signature. This is consistent with the conceptual expectation for channel-targeted interference: one should not expect every packet-level feature to move, but one should expect the directly affected physical-layer quality metric to shift substantially [31, 33]. This behavior is visible in Figure 2, and the full test statistics are listed in Table 6.

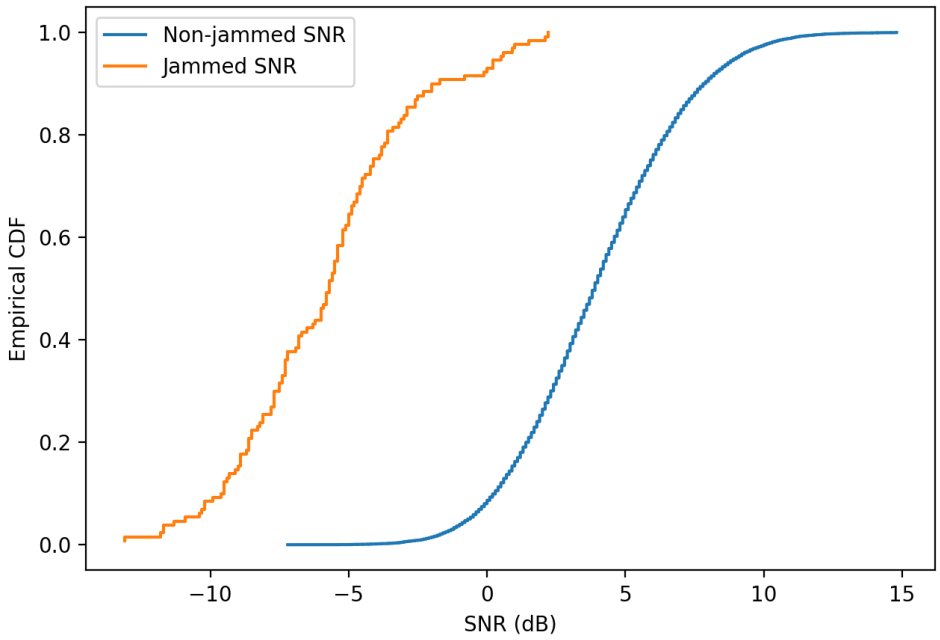


**Fig. 2.** Empirical CDF comparison of non-jammed and jammed SNR values. The sharp leftward shift in the jammed distribution confirms that the deterministic jammer produces a strong, localized physical-layer signature.

**Table 6.** Kolmogorov–Smirnov results for jammed traffic versus non-jammed traffic.

| Feature | KS statistic | $p$-value |
|---|---|---|
| Spreading factor | 0.0389 | 0.9857 |
| Payload size | 0.0436 | 0.9581 |
| RSSI | 0.0632 | 0.6568 |
| SNR | 0.8868 | $2.04 \times 10^{-122}$ |

From a security-interpretation standpoint, the replay and jamming results complement one another. Replay demonstrates the limits of simple distribution-based detection: an adversary can preserve the observed structure of legitimate traffic by retransmitting previously valid packets. Jamming demonstrates the opposite phenomenon: a well-targeted attack may remain limited in scope but still become obvious in the one variable it directly perturbs. These are not contradictory results. They reveal two distinct classes of attack behavior. Replay is covert in the feature space tested here; jamming is overt in a narrow feature space. This distinction matters for system defense because it suggests that no single detection principle will suffice across attack types.

Metadata exposure produces a third kind of outcome. The device-level metadata-risk scores in the executed dataset are moderate rather than strongly saturated. The mean score is 0.1189, the median is 0.0987, the 90th percentile is 0.2382, the 99th percentile is 0.2913, and the maximum observed score is 0.3372. These values indicate that devices are not perfectly indistinguishable, but neither are they so regular that passive tracking becomes trivially easy under the current model. This suggests that metadata leakage in digital-twin studies is highly sensitive to how cadence, channel concentration, and feature entropy are generated [35, 38]. The distribution summary is given in Table 7.

**Table 7.** Distribution of device-level metadata-risk scores.

| Statistic | Mean | Median | 90th pct. | 99th pct. | Maximum |
|---|---|---|---|---|---|
| Metadata risk score | 0.1287 | 0.1014 | 0.2531 | 0.2837 | 0.3233 |

This finding deserves careful interpretation. A lower metadata-risk profile does not mean metadata leakage is unimportant. It means the current framework produces enough variability that device behavior is less rigidly predictable than in more structured traffic models. From one perspective, this is reassuring because it suggests that modest diversity in parameter usage and cadence can reduce passive identifiability. From another perspective, it also shows that metadata exposure is not a fixed property of LoRaWAN itself, but a function of how deployment patterns, device behavior, and scheduling logic interact. The practical implication is that security evaluation should not only ask whether metadata is visible, but also how predictable that metadata becomes under real or simulated operating patterns.

A second important interpretive point concerns the temporal diagnostics. The current baseline is irregular, but its concentration metrics are relatively mild. This likely contributes to the moderate metadata-risk levels. If the baseline were more tightly clustered in specific hourly windows, or if devices exhibited more rigid cadence and narrower channel usage, the metadata-risk distribution would likely shift upward. The present result therefore offers a useful reminder that digital-twin security studies are most informative when they expose the link between baseline structure and attack observability. A replay study with a good baseline reveals whether replay remains covert. A jamming study with channel-localized interference reveals how sharply SNR separates. A metadata

study reveals how traffic regularity conditions passive exposure. In each case the baseline is not merely a neutral backdrop; it is part of the mechanism that makes the result meaningful [34].

At the same time, the current implementation has limitations that should temper overstatement. The timestamp-generation mechanism is simplified. Device assignment to the timeline is structured but not based on independent device-level stochastic processes. Frame counters are still not modeled as truly monotonic on a per-device basis in the baseline table. The diagnostics JSON export fails due to a NumPy serialization issue in the code's final block. None of these issues invalidate the simulation study, but they do matter for how strongly its results can be generalized. The correct claim is not that the paper reproduces an exact maritime deployment. The correct claim is that it provides a reproducible synthetic environment in which the relative visibility and subtlety of multiple attack types can be studied.

From a research perspective, this is still valuable. The replay experiment confirms a robust principle: packets copied from baseline traffic can remain statistically close to baseline under simple univariate tests. The jamming experiment confirms another robust principle: deterministic narrowband interference can generate a very strong SNR signature without broadly disturbing other radio features. The metadata experiment adds a more nuanced insight: exposure exists, but its intensity depends strongly on the traffic model. This combination of continuity and variation is exactly what a good digital-twin study should reveal.

### *4.1. Robustness and sensitivity analysis*

To examine whether the main conclusions depend excessively on a single realization, the framework was re-run across five random seeds while keeping the nominal replay density and jammer schedule fixed. Across these runs, replay SNR remained non-separable under KS testing, while jamming SNR remained strongly separable with KS statistics consistently above 0.88 and extremely small $p$-values. Median device-level metadata risk remained stable in the narrow range 0.096–0.103, and temporal alignment with the target activity profile remained strong, with Spearman correlation ranging from 0.747 to 0.819. These results indicate that the principal findings of replay subtlety, jamming visibility, and moderate metadata exposure are not artifacts of a single seed [16, 56]. The multi-seed results are summarized in Table 8.

**Table 8.** Multi-seed robustness analysis for key output metrics.

| Seed | Jammed packets | $\rho$ (activity alignment) | Replay $KS_{SNR}$ | Jamming $KS_{SNR}$ | Median risk |
|---|---|---|---|---|---|
| 11 | 109 | 0.747 | 0.0327 | 0.8932 | 0.1033 |
| 23 | 129 | 0.790 | 0.0403 | 0.8814 | 0.1032 |
| 42 | 130 | 0.812 | 0.0209 | 0.8868 | 0.1014 |
| 77 | 134 | 0.753 | 0.0235 | 0.8779 | 0.1007 |
| 101 | 132 | 0.819 | 0.0288 | 0.8810 | 0.0956 |

Sensitivity checks were also carried out for replay density and jammer duty cycle. When replay density was varied from 1% to 5%, replay SNR remained non-separable under KS testing, with all replay $p$-values above 0.44. Likewise, when the jammer duty cycle was varied from 5% to 20%, jamming SNR remained strongly separable in every case. This indicates that the qualitative conclusions do not depend on one narrow parameter choice [26, 33]. The sensitivity results are presented in Table 9.

**Table 9.** Sensitivity analysis for replay density and jammer duty cycle.

| Scenario | Replay/Jam setting | Replay $KS_{SNR}$ | Replay $p$-value | Jamming $KS_{SNR}$ |
|---|---|---|---|---|
| Replay sensitivity | 1% replay density | 0.0316 | 0.6244 | 0.8931 |
| Replay sensitivity | 2% replay density | 0.0209 | 0.9942 | 0.8868 |
| Replay sensitivity | 5% replay density | 0.0180 | 0.4456 | 0.8895 |
| Jamming sensitivity | 6 s on / 114 s off | 0.0217 | 0.9838 | 0.8836 |
| Jamming sensitivity | 12 s on / 108 s off | 0.0209 | 0.9942 | 0.8868 |
| Jamming sensitivity | 24 s on / 96 s off | 0.0215 | 0.9904 | 0.8947 |

## 5. Security Implications, Practical Relevance, and Study Limits

The security implications of the present findings are straightforward, but they are best understood by distinguishing among detection, prevention, and exposure reduction. Replay, jamming, and metadata leakage should not be treated as variations of the same problem. They require different defensive priorities because they occupy different positions in the observable behavior of the system.

Replay should be understood primarily as a state-management and protocol-enforcement problem. The present study shows that replayed packets remain statistically close to non-replay traffic across spreading factor, payload size, RSSI, and SNR. This means a defender should not expect simple radio-feature anomaly monitoring to provide reliable replay identification. The more appropriate defensive focus is on monotonic counters, sequence checking, session state validation, and rejection of stale or duplicated uplinks. In practice, this means that a system architect should view replay resilience as a design property of the protocol stack and server acceptance logic, not as a byproduct of feature-distribution monitoring.

Jamming belongs to a different defensive category. In the current study, narrowband interference remains sparse but highly visible in SNR. This suggests that targeted physical-layer attacks are more naturally addressed through channel-aware monitoring, SNR baselining, and adaptive response mechanisms such as channel diversity or communication fallback. Aggregate traffic counts alone are unlikely to capture the structure of this sort of attack, because the jammer need not affect enough packets to distort high-level throughput statistics. The important signal lies in a localized quality metric shift, especially when it occurs on targeted channels within meaningful operational windows.

Metadata leakage occupies a third category because it is not always an attack in the same active sense as replay or jamming. It is often a passive exposure. The study shows that metadata risk is real but not saturated under the current framework. That result implies that modest variability in cadence and protocol-parameter usage can materially reduce passive re-identification susceptibility. Practically, this suggests that deployments should consider not only payload encryption, but also the predictability of transmission behavior. Simple forms of mitigation could include introducing controlled jitter into reporting intervals, avoiding overly rigid device schedules where operationally feasible, and reducing unnecessary coupling between observable traffic patterns and semantic device roles.

These implications also matter for the broader design philosophy of maritime IoT systems. A common weakness in applied cybersecurity is to assume that once confidentiality is provided, the residual security problem is minor. The present study argues the opposite. A maritime LoRaWAN deployment may still be insecure even when payload confidentiality is intact because an attacker can

exploit protocol semantics, radio-layer sensitivity, or behavioral predictability. In this sense, the distinction between communication security and operational security is crucial. A secure communication primitive does not automatically produce a secure operational system.

The study also illustrates the importance of interpretive restraint. Because the framework is synthetic, the results should not be framed as direct estimates of incident rates or exploitation frequencies in real ports. They are better understood as controlled demonstrations of how different attack classes behave under a structured baseline. Such demonstrations are highly valuable for research and design reasoning, but they do not eliminate the need for future validation against richer datasets or field measurements.

Several technical limitations deserve explicit statement. The timestamp process is only an approximation of non-homogeneous operational activity and does not yet fully encode device-specific emission processes. The device assignment mechanism is structured but simplified. Channel-bias logic is only partially exploited in the code. The replay delay distribution yields a mean larger than its median because of exponential sampling, which is analytically fine but should be described accurately. The diagnostic export bug shows that the current implementation still requires cleanup for production-style reproducibility packaging. None of these issues undermine the use of the current code as a paper foundation, but they do define the boundary between a strong simulation study and a polished benchmark framework.

These limitations also point directly toward future work. A stronger next version of the framework would generate each device as an independent stochastic process rather than assigning device identity after global timestamp generation. It would track frame counters per device and use those counters explicitly in replay-acceptance logic. It would include multi-seed robustness analysis and perhaps limited multivariate classification experiments to complement KS-based descriptive tests. It would also benefit from visual outputs such as empirical CDF plots for replay and jamming, timeline heatmaps, and distributions of device-level metadata risk. These additions would not change the central findings of the present manuscript, but they would increase both empirical richness and submission readiness.

Even with these caveats, the current paper remains useful and defensible. It takes a reproducible framework, converts it into a structured digital-twin security study, and extracts three substantive findings from it. It confirms the subtlety of replay, the narrow strength of deterministic jamming, and the moderate but meaningful presence of metadata exposure. That combination is enough to justify the paper as a contribution to the study of maritime LoRaWAN security behavior.

## 6. Conclusion

This paper presented a digital-twin study of maritime LoRaWAN security behavior centered on replay attacks, deterministic narrowband jamming, and metadata-based re-identification exposure. A baseline of 20,000 transmissions from 180 simulated devices was generated over a 72 h horizon under EU863–870 MHz settings, after which replay injection, jamming, and metadata-risk scoring were applied.

The empirical outcomes reveal three distinct security behaviors. Replay traffic remained statistically close to non-replay traffic across all evaluated univariate radio features, confirming that replay is difficult to separate through simple feature-distribution monitoring and therefore must be addressed through stateful protocol enforcement. Deterministic narrowband jamming produced a

sparse but extremely strong SNR signature while leaving spreading factor, payload size, and RSSI comparatively unaffected. Metadata-risk scores were measurable but moderate, showing that passive exposure exists, yet depends strongly on the regularity and concentration properties of the baseline model.

The broader significance of these findings lies in what they reveal about digital-twin security research. A good synthetic framework is not useful because it reproduces one fixed risk profile. It is useful because it makes visible which conclusions are stable across implementations and which are sensitive to model structure. In the present study, replay subtlety and jamming visibility appear robust, whereas metadata-risk magnitude is more model-dependent. This is an important insight for both researchers and practitioners because it highlights the need to connect security assessment not only to attack classes, but also to the temporal and behavioral assumptions embedded in the traffic model.

The framework has limitations, and these should be stated explicitly. The baseline is synthetic, the temporal process is simplified, device-specific emission dynamics remain underdeveloped, and the code still contains a diagnostics-export bug. Yet none of these issues erase the value of the current paper. The study provides a reproducible, interpretable, and methodologically coherent analysis of maritime LoRaWAN security behavior. It offers a solid basis for submission and also provides a clear roadmap for future refinement toward a stronger benchmark and a more field-proximate cybersecurity evaluation framework.

## Funding

This research received no external funding.

## Data Availability Statement

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

## Code Availability Statement

The code used to generate the digital-twin traffic, implement the replay and jamming scenarios, compute metadata-risk scores, and reproduce the reported figures and sensitivity analyses is available from the corresponding author upon reasonable request.

## Acknowledgments

The author thanks the broader LPWAN and maritime cybersecurity research communities whose open technical and survey literature helped shape the framing of this study.

## Conflicts of Interest

The author declares no conflict of interest.

## References


[1] Semtech Corporation. IoT Applications for Smart Supply Chain & Logistics with LoRa Technology. Available online: https://www.semtech.com/lora/lora-applications/smart-supply-chain-logistics.

[2] Haxhibeqiri, J., De Poorter, E., Moerman, I., & Hoebeke, J. (2018). A survey of LoRaWAN for IoT: From technology to application. *Sensors, 18*(11), 3995.

[3] Almuhaya, M. A., Jabbar, W. A., Sulaiman, N., & Abdulmalek, S. (2022). A survey on Lorawan technology: Recent trends, opportunities, simulation tools and future directions. *Electronics, 11*(1), 164.

[4] Bonilla, V., Campoverde, B., & Yoo, S. G. (2023). A systematic literature review of lorawan: Sensors and applications. *Sensors, 23*(20), 8440.

[5] Port of Barcelona PierNext. LoRaWAN: The IoT Network Serving Ports. 2018. Available online: https://piernext.portdebarcelona.cat/en/technology/lorawan-the-iot-network-serving-ports-2/.

[6] Smart Maritime Network. Stena Line Vessel to Implement LoRa IoT Network. 2022. Available online: https://smartmaritimenetwork.com/2022/07/29/stena-line-vessel-to-implement-lora-iot-network/.

[7] United Nations Conference on Trade and Development (UNCTAD). *Review of Maritime Transport 2023*. UNCTAD, Geneva, Switzerland, 2023.

[8] Butun, I., Pereira, N., & Gidlund, M. (2018). Security risk analysis of LoRaWAN and future directions. *Future Internet, 11*(1), 3.

[9] Hessel, F., Almon, L., & Hollick, M. (2023). LoRaWAN security: an evolvable survey on vulnerabilities, attacks and their systematic mitigation. *ACM Transactions on Sensor Networks, 18*(4), 1-55.

[10] LoRa Alliance. *LoRaWAN 1.0.4 Specification*. LoRa Alliance, Beaverton, OR, USA, 2020.

[11] Adefemi Alimi, K. O., Ouahada, K., Abu-Mahfouz, A. M., & Rimer, S. (2020). A survey on the security of low power wide area networks: Threats, challenges, and potential solutions. *Sensors, 20*(20), 5800.

[12] Grant, S. (2016). 3gpp low power wide area technologies-gsma white paper. *Gsma. com*, 1.

[13] 3GPP, B. (2020). Security architecture and procedures for 5G system. *Technical Specification (TS) 3GPP TS 33.501 V17. 0.0 (2020–2012).*

[14] Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE access, 8*, 108952-108971.

[15] National Institute of Standards and Technology (NIST). *Security and Trust Considerations for Digital Twin Technology*. NIST IR 8356, 2025.

[16] Alcaraz, C., & Lopez, J. (2025). Digital twin security: a perspective on efforts from standardization bodies. *IEEE Security & Privacy, 23*(1), 83-90.

[17] Allison, D., Smith, P., & Mclaughlin, K. (2023, August). Digital twin-enhanced incident response for cyber-physical systems. In *Proceedings of the 18th International Conference on Availability, Reliability and Security* (pp. 1-10).

[18] Maurya, P., Hazra, A., Kumari, P., Sørensen, T. B., & Das, S. K. (2025). A comprehensive survey of data-driven solutions for lorawan: Challenges and future directions. *ACM Transactions on Internet of Things, 6*(1), 1-36.

[19] Pensieri, S., Viti, F., Moser, G., Serpico, S. B., Maggiolo, L., Pastorino, M., ... & Bozzano, R. (2021). Evaluating LoRaWAN connectivity in a marine scenario. *Journal of Marine Science and Engineering, 9*(11), 1218.

[20] Jovalekic, N., Drndarevic, V., Pietrosemoli, E., Darby, I., & Zennaro, M. (2018). Experimental study of LoRa transmission over seawater. *Sensors, 18*(9), 2853.

[21] Pinelo, J., Rocha, A. D., Arvana, M., Gonçalves, J., Cota, N., & Silva, P. (2023). Unveiling lora's oceanic reach: Assessing the coverage of the azores lorawan network from an island. *Sensors, 23*(17), 7394.

[22] Ojo, M. O., Adami, D., & Giordano, S. (2021). Experimental evaluation of a LoRa wildlife monitoring network in a forest vegetation area. *Future Internet, 13*(5), 115.

[23] Yang, X., Karampatzakis, E., Doerr, C., & Kuipers, F. (2018, April). Security vulnerabilities in LoRaWAN. In *2018 IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)* (pp. 129-140). IEEE.

[24] Tomasin, S., Zulian, S., & Vangelista, L. (2017, March). Security analysis of lorawan join procedure for internet of things networks. In *2017 IEEE Wireless Communications and Networking Conference Workshops (WCNCW)* (pp. 1-6). IEEE.

[25] Aras, E., Ramachandran, G. S., Lawrence, P., & Hughes, D. (2017, June). Exploring the security vulnerabilities of LoRa. In *2017 3rd IEEE international conference on cybernetics (CYBCONF)* (pp. 1-6). IEEE.

[26] Ruotsalainen, H., Shen, G., Zhang, J., & Fujdiak, R. (2022). LoRaWAN physical layer-based attacks and countermeasures, a review. *Sensors, 22*(9), 3127.

[27] Kim, J., & Song, J. (2017, November). A simple and efficient replay attack prevention scheme for LoRaWAN. In *Proceedings of the 2017 7th International Conference on Communication and Network Security* (pp. 32-36).

[28] Na, S., Hwang, D., Shin, W., & Kim, K. H. (2017, January). Scenario and countermeasure for replay attack using join request messages in LoRaWAN. In *2017 International Conference on Information Networking (ICOIN)* (pp. 718-720). IEEE.

[29] Sung, W. J., Ahn, H. G., Kim, J. B., & Choi, S. G. (2018, February). Protecting end-device from replay attack on LoRaWAN. In *2018 20th International Conference on Advanced Communication Technology* (ICACT) (pp. 167-171). IEEE.

[30] Voigt, T., Bor, M., Roedig, U., & Alonso, J. (2016). Mitigating inter-network interference in LoRa networks. *arXiv preprint arXiv:1611.00688.*

[31] Ingham, M., Marchang, J., & Bhowmik, D. (2020). IoT security vulnerabilities and predictive signal jamming attack analysis in LoRaWAN. *IET Information Security, 14*(4), 368-379.

[32] Bleszynski, B. J., Orfanidis, C., & Fafoutis, X. (2023, November). Detection of mobile LoRa jammers. In *2023 IEEE Virtual Conference on Communications (VCC)* (pp. 288-293). IEEE.

[33] Ruotsalainen, H. (2022, August). Reactive jamming detection for LoRaWAN based on metadata differencing. In *Proceedings of the 17th International Conference on Availability, Reliability and Security* (pp. 1-8).

[34] Larkin, Z., & Easttom, C. (2026). An Examination of LPWAN Security in Maritime Applications. *Journal of Cybersecurity and Privacy, 6*(2), 65.

[35] Spadaccino, P., Garlisi, D., Cuomo, F., Pillon, G., & Pisani, P. (2022). Discovery privacy threats via device de-anonymization in LoRaWAN. *Computer Communications, 189*, 1-10.

[36] Pélissier, S., Cunche, M., Roca, V., & Donsez, D. (2022, May). Device re-identification in LoRaWAN through messages linkage. In *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks* (pp. 98-103).

[37] Pélissier, S., Aalmoes, J., Mishra, A. K., Cunche, M., Roca, V., & Donsez, D. (2024, May). Privacy-preserving pseudonyms for LoRaWAN. In *Proceedings of the 17th ACM Conference on Security and Privacy in Wireless and Mobile Networks* (pp. 14-19).

[38] Pélissier, S., Mishra, A. K., Cunche, M., Roca, V., & Donsez, D. (2025). Efficiently linking LoRaWAN identifiers through multi-domain fingerprinting. *Pervasive and Mobile Computing, 112,* 102082.

[39] You, I., Kwon, S., Choudhary, G., Sharma, V., & Seo, J. T. (2018). An enhanced LoRaWAN security protocol for privacy preservation in IoT with a case study on a smart factory-enabled parking system. *Sensors, 18*(6), 1888.

[40] Park, S., Shaik, A., Borgaonkar, R., & Seifert, J. P. (2019, November). Anatomy of commercial IMSI catchers and detectors. In *Proceedings of the 18th Acm Workshop on Privacy in the Electronic Society* (pp. 74-86).

[41] Svilicic, B., Kamahara, J., Rooks, M., & Yano, Y. (2019). Maritime cyber risk management: An experimental ship assessment. *The Journal of Navigation, 72*(5), 1108-1120.

[42] International Maritime Organization. *Guidelines on Maritime Cyber Risk Management*, MSC-FAL.1/Circ.3/Rev.3. International Maritime Organization, London, UK, 2025.

[43] Ben Farah, M. A., Ukwandu, E., Hindy, H., Brosset, D., Bures, M., Andonovic, I., & Bellekens, X. (2022). Cyber security in the maritime industry: A systematic survey of recent advances and future trends. *Information, 13*(1), 22.

[44] Bolbot, V., Kulkarni, K., Brunou, P., Banda, O. V., & Musharraf, M. (2022). Developments and research directions in maritime cybersecurity: A systematic literature review and bibliometric analysis. *International Journal of Critical Infrastructure Protection, 39*, 100571.

[45] Datalastic. Vessel Tracker API & Ship AIS Database. Digital Lake s.r.o. Available online: https://datalastic.com.

[46] Monjur, M. M. R., Heacock, J., Sun, R., & Yu, Q. (2021, November). An attack analysis framework for LoRaWAN applied advanced manufacturing. In *2021 IEEE International Symposium on Technologies for Homeland Security (HST)* (pp. 1-7). IEEE.

[47] Mohamed, A., Wang, F., Butun, I., Qadir, J., Lagerström, R., Gastaldo, P., & Caviglia, D. D. (2022). Enhancing cyber security of LoRaWAN gateways under adversarial attacks. *Sensors, 22*(9), 3498.

[48] Kuntke, F., Romanenko, V., Linsner, S., Steinbrink, E., & Reuter, C. (2022). LoRaWAN security issues and mitigation options by the example of agricultural IoT scenarios. *Transactions on Emerging Telecommunications Technologies, 33*(5), e4452.

[49] Czeczot, G., Rojek, I., & Mikołajewski, D. (2023). Analysis of cyber security aspects of data transmission in large-scale networks based on the LoRaWAN protocol intended for monitoring critical infrastructure sensors. *Electronics, 12*(11), 2503.

[50] Sikder, A. K., Petracca, G., Aksu, H., Jaeger, T., & Uluagac, A. S. (2021). A survey on sensor-based threats and attacks to smart devices and applications. *IEEE Communications Surveys & Tutorials, 23*(2), 1125-1159.

[51] Gu, C., Jiang, L., Tan, R., Li, M., & Huang, J. (2021). Attack-aware synchronization-free data timestamping in LoRaWAN. *ACM Transactions on Sensor Networks (TOSN), 18*(1), 1-31.

[52] Kamkuemah, M. N. (2021, December). Epistemic analysis of a key-management vulnerability in LoRaWAN. In *2021 18th International Conference on Privacy, Security and Trust (PST)* (pp. 1-7). IEEE.

[53] Ntshabele, K., Isong, B., Gasela, N., & Abu-Mahfouz, A. M. (2022). A comprehensive analysis of LoRaWAN key security models and possible attack solutions. *Mathematics, 10*(19), 3421.

[54] Tsai, K. L., Chen, L. W., Leu, F. Y., Hsu, H. C., & Wu, C. T. (2021, October). Secure LoRaWAN root key update scheme for IoT environment. In *International Symposium on Mobile Internet Security* (pp. 3-15). Singapore: Springer Nature Singapore.

[55] An, K. (1933). Sulla determinazione empirica di una legge didistribuzione. *Giorn Dell'inst Ital Degli Att, 4*, 89-91.

[56] Massey Jr, F. J. (1951). The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American statistical Association, 46*(253), 68-78.

**How to cite this article:** A. Asghar (2026). A Reproducible Digital-Twin Assessment of Replay, Narrowband Jamming, and Metadata Exposure in Simulated Maritime LoRaWAN Traffic. *Bulletin of Computer and Data Sciences*, 7(1), 17-37. DOI: 10.71448/bcds2671-3

**Received:** 18/12/2025 **Revised:** 17/2/2026 **Accepted:** 20/2/2026 **Published:** 31/03/2026



*Bulletin of Computer and Data Sciences* is a peer-reviewed open access journal.