electronics

MDPI

Article

Secure QR Code Attendance over RPL-Based Low-Power IoT Networks: Security Analysis and Scalability Evaluation in Contiki-NG/Cooja

Rashid Mustafa and Nurul I. Sarkar *

Department of Computer and Information Sciences, Auckland University of Technology, Auckland 1010, New Zealand; rashid.mustafa@autuni.ac.nz
* Correspondence: nurul.sarkar@aut.ac.nz

Abstract

QR code attendance offers a low-cost way to digitise presence records, but a visible code by itself does not show that an attendance transaction is fresh, belongs to an authorised participant, or has been delivered reliably over a constrained network. This paper examines a post-scan QR attendance transaction in a Routing Protocol for Low-Power and Lossy Networks (RPL)-based Internet of Things environment using Contiki-NG/Cooja. Unlike attendance studies that mainly strengthen QR freshness or identity checks and RPL studies that mainly examine routing behaviour, the present work evaluates network delivery, application acceptance, and exact replay rejection along the same transaction path. The baseline combines token validation, attendance-rule checking, and server-side duplicate state with IEEE 802.15.4, IPv6 over Low-Power Wireless Personal Area Networks (6LoWPAN), RPL Lite, Carrier-Sense Multiple Access (CSMA), and the User Datagram Protocol (UDP). The security claim is deliberately limited to the evidence available from the recorded experiments: exact replayed attendance transactions are evaluated, whereas confidentiality, cryptographic origin authentication, routing-control protection, and physical compromise are not claimed as implemented defences. Three recorded network sizes (5, 20, and 50 motes) are analysed using packet delivery ratio (PDR), packet loss, average end-to-end delay, attendance success rate (ASR), and replay rejection. At 50 motes, PDR is 91.7%, average delay is 37 ms, and ASR is 89.4%. All 38 recorded exact replay attempts are rejected; the corresponding two-sided exact 95% binomial confidence interval is 90.75–100%, so the observed result is not presented as universal replay immunity. The PDR–ASR gap remains small across the recorded scenarios, but the available aggregate logs do not permit the gap to be assigned uniquely to application rejection, queueing, retransmission, or radio effects. The study therefore provides a bounded application-security and scalability result, identifies the limits of the present evidence, and defines the additional repeated-run, multi-hop, energy, cryptographic, and hardware measurements required for deployment-oriented validation.

Keywords: wireless sensor networks; low-power and lossy networks; Internet of Things; QR code attendance; RPL; 6LoWPAN; Contiki-NG; Cooja; replay attack; application security; scalability

Check for updates

Academic Editor: Hung-Yu Chien

Received: 21 September 2026
Revised: 17 October 2026
Accepted: 23 October 2026
Published: 25 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Low-power Internet of Things (IoT) networks increasingly carry transactions whose correctness matters at the application layer, not only packets whose delivery matters at

the network layer. This distinction is important for wireless sensor networks (WSNs) and low-power and lossy networks (LLNs), where modest processors, limited memory, shared radio channels, and restricted energy budgets constrain both communication and security design [1–3]. A protection mechanism that adds state, cryptographic operations, packet bytes, or retransmissions can strengthen assurance while also consuming scarce resources. Conversely, a network that is economical but unreliable can undermine an otherwise correct application because a legitimate transaction that never reaches the validation service cannot be accepted.

The Routing Protocol for Low-Power and Lossy Networks (RPL) was standardised for IPv6 routing in LLNs and remains a central protocol for constrained multi-hop IoT deployments [4]. Contiki-NG and Cooja provide a practical environment for studying such systems across routing, radio, and application behaviour [5,6]. Comparative benchmarking of low-end IoT operating systems also places Contiki-NG within the broader resource-constrained software landscape [7]. Recent work illustrates the security–performance coupling of RPL networks: replayed control traffic can disturb routing behaviour [8], Hello Flood attacks can degrade delivery and power-related metrics [9], and energy-aware parent selection introduces additional routing logic in exchange for resource and quality-of-service benefits [10]. These observations motivate application studies in which network behaviour and security outcomes are evaluated together.

QR code attendance is a useful example because the visible interaction is simple while the trust problem is not. In a lecture, laboratory, workplace training session, or smart building event, many participants may submit attendance within a short interval. A copied image can be forwarded to an absent participant, a previously accepted transaction can be retransmitted, an identifier can be altered, and legitimate traffic can be lost or delayed during a burst. The system therefore has two simultaneous obligations: the communication path must carry legitimate messages with acceptable delay, and the application must avoid turning repeated or invalid messages into durable attendance records.

Recent attendance systems have used hashing and public key cryptography [11], dynamic QR values and UUIDs [12], additional identity checks [13], and QR/GPS combinations [14]. These approaches primarily strengthen the user-facing attendance workflow. By contrast, RPL security research generally concentrates on control plane or routing attacks rather than the semantic success of an attendance transaction [15,16]. The present study connects these two research directions by evaluating a post-scan QR attendance transaction over an RPL-based LLN and by keeping network delivery, application acceptance, and replay rejection as distinct but related observations.

The implemented replay resistance mechanism is intentionally transparent. When a transaction passes the token and attendance rules, the server retains sufficient transaction state for the active attendance context. A later byte-equivalent or logically duplicate submission that maps to an already accepted event fails the duplicate/freshness predicate and is rejected. This mechanism is not presented as a new cryptographic primitive, nor is it equivalent to proving device possession of a nonce, one-time password (OTP), or time-based secret. Time-based tokens, nonces, and OTPs generate or validate freshness evidence; the mechanism studied here focuses on server-side prevention of a second durable attendance record for the same accepted logical event. The research contribution lies in analysing that application property together with RPL network scalability and in explicitly separating what the baseline proves from properties that require cryptographic protection.

1.1. Research Challenges, Gaps, and Novelty

The first challenge is the difference between a visible QR value and an authenticated transaction. A QR image can be public or copied, so trust should ultimately be associated

with the response sent to the server rather than secrecy of the visual code itself. The second challenge is duplicate state: replay resistance requires enough short-lived server state to recognise an already accepted logical event without allowing the cache to grow indefinitely. The third challenge is burst scalability. Attendance requests often arrive within a narrow window; increasing the number of active motes may raise contention, packet loss, and delay before validation is even reached. The fourth challenge is resource accountability. Terms such as "lightweight" and "energy efficient" should be supported by memory, CPU, radio, or energy measurements rather than inferred from a small algorithm alone.

The related literature leaves a gap at the boundary between these concerns. Secure attendance studies generally evaluate the application, authentication method, or mobile interface, while RPL security studies evaluate routing behaviour or attack detection. Few studies report, within one attendance-oriented RPL experiment, both network-level delivery and application-level acceptance while also quantifying application replay rejection. A second gap concerns claim scope: a replay check is sometimes discussed as if it implied general authentication, even though QR cloning, message modification, node impersonation, gateway compromise, and routing attacks require different controls. A third gap is experimental traceability. Scalability claims are stronger when the measured range, unavailable metrics, and statistical limits are made explicit instead of being filled with assumed simulator defaults.

The novelty of this study is therefore methodological and architectural rather than cryptographic. First, it treats the QR scan as the start of a networked attendance transaction and evaluates PDR, loss, delay, ASR, and replay rejection along the same end-to-end path. Second, it formalises the baseline acceptance rule and shows the precise condition under which an exact replay is rejected. Third, it distinguishes the validated replay property from unimplemented properties such as authenticated payload integrity, secure enrolment, routing-layer defence, and measured energy efficiency. Fourth, it reports the demonstrated scaling envelope up to 50 motes without extrapolating a maximum deployment size that was not experimentally reached. Finally, it defines a concrete extension using NIST-standardised Ascon authenticated encryption, secure device enrolment, repeated runs, energy instrumentation, and hardware validation [17,18].

1.2. Research Questions and Key Contributions

The research questions were formulated from two issues that emerge from the problem addressed in this study. First, an attendance transaction carried over a low-power and lossy network depends on more than the correctness of the application logic; changes in network scale can affect whether the transaction reaches the receiving application within an acceptable time and is ultimately recorded. Examining network performance without considering the resulting attendance outcome would therefore provide only a partial view of system behaviour. Second, rejecting a repeated transaction is an important security property, but it should not be confused with complete authentication or end-to-end protection. The evaluation must consequently establish what the implemented validation mechanism actually demonstrates while identifying the security properties that remain outside the present implementation. These considerations motivate the following research questions:

- RQ1: What relationship exists between network scale, packet delivery, communication delay, and successful attendance recording in the evaluated RPL-based low-power IoT network?
- RQ2: What evidence does the implemented token and duplicate validation mechanism provide for replay resistance, and what additional safeguards are required before stronger authentication and energy efficiency claims can be supported?

Justification for RQ1. The first research question connects communication behaviour with the outcome that is meaningful to the attendance application. Increasing the number of participating motes can influence contention, packet delivery, packet loss, and delay, but these network effects become practically important only when they affect the completion of attendance transactions. RQ1 therefore examines packet delivery ratio (PDR), packet loss, end-to-end delay, and attendance success rate (ASR) together. This enables the evaluation to determine whether application-level performance follows changes observed in the underlying RPL network rather than treating network and application metrics as unrelated measurements.

Justification for RQ2. The second research question defines the security boundary of the present work. The implemented mechanism uses token checking and duplicate/freshness validation to determine whether a received attendance transaction should be accepted. The recorded replay experiments provide evidence about repeated submissions under the evaluated conditions, but they do not by themselves demonstrate cryptographic message integrity, secure device identity, resistance to routing manipulation, or measured energy efficiency. RQ2 was therefore formulated to assess the replay property that can be supported by the available evidence while making clear which additional mechanisms and experiments would be necessary for a stronger end-to-end security claim.

The study makes four principal contributions:

1. An integrated low-power attendance architecture. The study models QR-based attendance as an end-to-end network transaction rather than as an isolated QR-scanning operation. The architecture separates attendance initiation, RPL-based low-power communication, gateway delivery, and server-side validation. This separation makes it possible to identify where communication failure and application rejection occur and provides a clearer basis for evaluating the interaction between network and attendance behaviour.
2. A clearly bounded replay resistance analysis. The work defines the conditions under which an attendance transaction is accepted and examines how token and duplicate/freshness checks respond when an already accepted transaction is submitted again. The security contribution is deliberately limited to the property demonstrated by the implementation and recorded replay experiments. The accompanying threat analysis also distinguishes this result from stronger properties such as cryptographic authentication, message-integrity protection, node impersonation resistance, and routing-layer security.
3. An application-aware scalability evaluation. The experimental study evaluates the architecture with 5, 20, and 50 motes in Contiki-NG/Cooja and considers both network and application outcomes. Packet delivery ratio, packet loss, end-to-end delay, and attendance success are analysed together, allowing changes in low-power network performance to be related directly to the success of the attendance service. The evaluation therefore provides a measured scalability range without extrapolating an unsupported maximum network capacity.
4. A deployment-oriented security and validation pathway. Beyond reporting the current simulation results, the study identifies the controls and measurements needed to progress toward a more complete operational system. These include authenticated transaction protection, secure device enrolment, routing-aware defence, repeated experiments across independent simulation runs, direct energy and memory measurements, and hardware-based validation under realistic wireless conditions. This contribution establishes a practical baseline against which the performance and resource cost of future security extensions can be assessed.

Together, these contributions position the work as an evaluation of how the baseline attendance-validation logic interacts with RPL-based communication, rather than as a claim of a new cryptographic primitive. The resulting framework links application success, network scalability, and a precisely defined replay property while keeping stronger security and energy efficiency claims within the limits of experimentally demonstrated evidence. The remainder of the paper is organised as follows. Section 2 provides the technical background and motivates the choice of QR interaction for the target setting. Section 3 reviews recent RPL security, lightweight authentication, and secure attendance studies and positions the contribution. Section 4 presents the architecture, protocol semantics, and threat model. Section 5 defines the simulation methodology, metrics, and statistical treatment. Section 6 reports the measured results. Section 7 interprets security, scalability, and deployment implications; Section 8 states the validity limits; Section 8.2 defines the next validation stage; and Section 9 concludes the paper.

2. Technical Background

The attendance architecture operates over a constrained IPv6 communication stack in which short application messages are transported through a low-power wireless network. IEEE 802.15.4 provides the low-rate link, 6LoWPAN adapts IPv6 to constrained frames, RPL Lite supplies routing toward the root or gateway, CSMA coordinates medium access, and UDP carries the attendance payload. RPL is especially relevant because a larger population can increase contention and forwarding activity even when the application message itself is small [4,5]. The study therefore evaluates the attendance service as a networked application rather than treating the QR interaction as an isolated web or mobile function.

2.1. QR Attendance as a Network Transaction

The simulation begins after visual recognition. A successful scan is abstracted into an attendance message containing the information needed by the validation service, such as participant identity, session context, timing/freshness information, and a token or transaction identifier. This boundary avoids attributing camera or QR-decoding accuracy to the network experiment. It also exposes the security-relevant event: the server must decide whether a delivered message is acceptable and whether the same logical attendance event has already been recorded.

Network delivery and attendance acceptance are deliberately separated. A legitimate message may fail before reaching the application because of packet loss, while a successfully delivered message may still be rejected by the attendance rules. This motivates the joint use of PDR and ASR later in the paper. A small PDR–ASR gap suggests that application success is tracking network delivery; a widening gap would indicate additional rejection or application-side failure after delivery.

2.2. Why QR Codes for the Target Setting?

QR codes are attractive when the deployment already has displays and camera-equipped phones or terminals because no dedicated contactless reader is required. A session value can also be rotated in software, making QR suitable for short attendance windows. These advantages do not make QR intrinsically secure: the image is observable and can be photographed or forwarded. The appropriate design assumption is therefore that the visual code is a challenge or session reference, while authenticity and freshness are enforced in the response transaction.

Table 1 contrasts QR interaction with common alternatives at an engineering level. NFC and RFID provide convenient proximity-based interactions but normally introduce

dedicated tags or reader infrastructure; BLE supports passive proximity but adds beacon, scanning, and radio management complexity; a mobile identity-only approach can provide strong account integration but depends more heavily on enrolment and backend identity services. QR remains useful where low deployment cost and easy session rotation are priorities, provided that visual copying is treated as an expected attack surface rather than ignored [11,13,14].

Table 1. Engineering comparison of attendance interaction options for the target low-power deployment.

Approach	Typical Infrastructure	Main Operational Advantage	Security/Deployment Consideration
QR code	Display/printed code plus camera-equipped client	Low additional hardware cost; session values can be changed in software	Visual content can be copied or forwarded; requires short-lived context and authenticated response for stronger assurance.
NFC	NFC-capable device and reader/tag	Deliberate close-range interaction	Requires compatible hardware and enrolment; physical proximity alone does not prevent credential sharing or device compromise.
RFID	Tags/cards and RFID readers	Fast identification and established access-control ecosystem	Adds reader/tag infrastructure; card sharing or cloning must be addressed separately.
BLE beacon	Beacons and BLE-capable clients/gateways	Passive or low-friction proximity detection	RSSI is an imperfect proxy for presence; scanning, privacy, and beacon maintenance add design complexity.
Mobile identity	Enrolled application/account and backend identity service	Can bind attendance directly to an account or device credential	Strongly depends on secure enrolment, credential lifecycle, stolen-device handling, and backend availability.

2.3. Replay Semantics and Security Scope

Replay can occur at several layers. RPL control message replay manipulates routing state and can affect availability or topology [8]; application replay resubmits an attendance event that was previously valid. The two attacks share the notion of stale information being reused but differ in object, consequence, and mitigation point. The present quantitative experiment addresses the application case. RPL attacks remain part of the broader threat model and require routing-aware controls.

The baseline uses deterministic token, authorisation, and duplicate checks rather than authenticated encryption. This design is small in logical complexity, but "lightweight" should not be interpreted as measured energy efficiency. NIST SP 800-232 makes Ascon-AEAD128 a relevant future protection mechanism for constrained devices because it provides authenticated encryption in a standardised lightweight family [17]. Secure enrolment is equally important because a cryptographic key must be associated with a trusted device or user before origin authentication can be claimed [18,19]. These distinctions are carried into the security model in Section 4.

3. Related Work and Research Positioning

Research relevant to this study falls into three connected areas: RPL security, authentication for constrained devices, and secure digital attendance. These areas increasingly overlap, but they are usually evaluated with different objectives. RPL studies tend to focus on routing availability and attack detection, lightweight authentication studies concentrate on cryptographic assurance and resource cost, and attendance studies usually emphasise identity, freshness, or resistance to proxy attendance. The comparison below is therefore organised around the specific question addressed here: how network delivery, application acceptance, and replay handling are related when an attendance transaction is carried over a constrained RPL path.

RPL security studies show that an application can fail even when its own validation logic is correct. Albinali and Azzedin studied replay attacks against RPL and demonstrated the effect of stale routing information on network operation [8]. Hkiri et al. evaluate Hello Flood behaviour in Cooja using delivery, delay, throughput, ETX, and power-related measures [9]. Arshad et al. proposed THC-RPL, a trust-enabled RPL scheme for Sybil-attack mitigation and reported packet-loss ratios of 15–25%, average node energy consumption of 60–80 mJ, approximately 40% greater node-level energy conservation, a 50% increase in network lifetime, and 10% fewer message exchanges in their evaluation [20]. The IoT-RPL 2021 dataset of Dhifallah et al. provides Cooja-generated traces for four RPL attack classes—Blackhole, Flooding, DODAG Version Number, and Decreased Rank—for intrusion detection research [21]. Prajapati et al. surveyed RPL attacks and countermeasures [15], while learning-based work examines black-hole and multiclass routing attack detection [16,22,23]. These studies establish that an application-level duplicate check cannot be used as evidence of routing security. Energy-aware RPL work likewise shows that routing choices and resource use are coupled [10].

A second body of work addresses authentication, cross-layer protection, and the cost of security on constrained devices. Guşiţa et al. and Haider et al. reviewed resource-conscious cryptographic and communication mechanisms for IoT systems [2,3], while broader multi-layer and cross-layer studies have emphasised the interactions between device, network, and application controls [24–26]. Garcia Ribera et al. evaluated an intrusion detection approach for RPL-based IoT with overhead considerations [27], and Matos et al. provided a useful deployed WSN contrast in which low-power operation is measured rather than inferred [28]. NIST SP 800-232 standardises Ascon-based lightweight cryptography, including Ascon-AEAD128 [17]. Fedrecheski et al. addressed secure zero-touch enrolment [18], and PUF-based protocols illustrated device-bound authentication with formal and practical evaluation [19,29]. These mechanisms address properties that the present duplicate-state rule does not provide, including cryptographic message integrity and proof of possession of a device-bound credential.

Secure attendance research has meanwhile moved beyond static QR registration. Irawan et al. combined RSA and hashing and examined a security–resource trade-off [11]. Annurrahma et al. used dynamic QR values with UUIDv4 and contextual validation [12]. Oyerinde et al. developed a secure higher-education QR attendance workflow [13], while Siregar et al. combined dynamic QR codes with GPS geofencing [14]. Panimozhi and Murthy presented a web-based QR attendance prototype with proxy mitigation measures and reported successful load handling in the application prototype [30]. Alabi et al. and Abirami and Rohann provide further recent examples of enhanced QR-based attendance designs [31,32]. These approaches strengthen the application-facing attendance process. The present study addresses a different measurement question by pairing attendance acceptance with RPL network delivery and delay at 5, 20, and 50 motes.

The present baseline is intentionally different from a nonce, OTP, or cryptographic authentication protocol. Once a logical attendance event has been accepted, the server retains duplicate state for the active session and prevents the same event from creating a second durable record. This provides a directly testable exact-replay property, but it is weaker than cryptographically binding fresh material to an enrolled device. The distinction is central to the scope of this paper and to the comparison in Table 2.

The comparison makes the novelty boundary explicit. The paper does not claim a new cryptographic primitive and does not treat duplicate checking as a substitute for dynamic QR generation, authenticated encryption, secure enrolment, or RPL attack detection. Its contribution is the joint interpretation of network-level delivery, application-level accep-

tance, and a precisely bounded replay property, together with a transparent statement of the measurements still needed to evaluate stronger security controls.

Table 2. Structured comparison of representative RPL security, constrained authentication, and QR attendance studies.

Reference	Primary Focus	Security/Validation Approach	RPL/LLN	Evaluation Evidence	Position Relative to This Study
Albinali and Azzedin [8]	RPL replay attacks	Routing-control replay analysis	Yes	Routing-level attack behaviour	Establishes replay risk at the routing layer rather than attendance semantics.
Hkiri et al. [9]	Hello Flood impact	Multi-metric attack evaluation	Yes	Delivery, delay, throughput, ETX and power-related metrics	Demonstrates network-layer security/performance coupling in Cooja.
Mokrani et al. [10]	Energy-aware RPL	Energy/QoS-aware parent selection	Yes	Routing and energy/QoS evaluation	Provides a routing and energy direction for a secured follow-on design.
Panda and Muthuraman [16]	RPL attack detection	Multiclass learning-based detection	Yes	Attack-classification evaluation	Covers broader routing attacks than the deterministic application check used here.
Irawan et al. [11]	QR attendance security	RSA and hashing	No	Attendance/security evaluation	Examines cryptographic attendance protection without RPL scaling.
Annurrahma et al. [12]	Dynamic QR attendance	UUIDv4 and contextual freshness	No	Application-level attendance validation	Strengthens code/session freshness at the application layer.
Oyerinde et al. [13]	Secure higher-education attendance	QR workflow with identity/security controls	No	Application/workflow evaluation	Strengthens attendance workflow but does not evaluate an RPL transport path.
Siregar et al. [14]	Dynamic QR + GPS	Dual-factor context/geofencing	No	Context-aware attendance validation	Adds contextual proxy-resistance evidence rather than LLN transport analysis.
Lim et al. [19]	Constrained-device authentication	PUF-based authenticated protocol	No	Formal and practical authentication evidence	Illustrates device-bound authentication beyond the present duplicate-state rule.
Arshad et al. [20]	RPL Sybil defence	Trust-enabled routing/malicious-node isolation	Yes	Packet loss, energy, lifetime and message overhead	Provides a routing-security baseline; it does not evaluate attendance semantics or ASR.
Dhifallah et al. [21]	RPL attack dataset	Cooja traces for four routing attacks	Yes	Four labelled attack classes	Supplies routing attack data rather than an attendance acceptance mechanism.
Panimozhi and Murthy [30]	QR attendance/proxy mitigation	Web-based QR workflow with proxy controls	No	Prototype and load-handling evaluation	Closest attendance-oriented comparator, but not an RPL/LLN transport evaluation.
This study	Attendance integrity + network scale	Token/rule checks + duplicate state; exact-replay test	Yes	PDR, loss, mean delay, ASR and replay counts	Relates network delivery, application acceptance and exact replay rejection along the same transaction path.

Quantitative Positioning Against Representative Comparison Studies

A numerical comparison is useful only when the measurement scope is kept visible. THC-RPL, IoT-RPL 2021, and the QR attendance prototype address different layers of the problem, so their numbers should not be treated as a league table. Table 3 therefore reports the quantitative evidence that each source contributes and places it next to the measurements from the present study. The present-study values used for comparison are drawn directly from the recorded measurements reported in the Results section.

The numerical evidence supports the paper's contribution without implying superiority over studies that measure different phenomena. The largest recorded scenario still delivers 91.7% of transmitted packets with a 37 ms mean delay and records 89.4% of intended attendance events, while all 38 recorded exact replays are rejected. Across the three scenarios, PDR falls by 6.5 percentage points and ASR by 7.1 percentage points, with the PDR–ASR gap remaining between 1.6 and 2.3 percentage points. These results are the publication-relevant evidence for the cross-layer question posed here: network delivery and attendance acceptance are measured together, rather than one being inferred from the

other. Because the scenarios are single aggregate runs, the manuscript continues to treat these changes as descriptive rather than statistically significant.

Table 3. Quantitative positioning against three representative comparison studies. Values are reported in their original evaluation scope and are not interpreted as directly interchangeable performance measures.

Study	Evaluation Target	Reported Numerical Evidence	Relevance to the Present Study
Arshad et al. (THC-RPL) [20]	Sybil-resistant RPL routing	Packet-loss ratio 15–25%; average node energy 60–80 mJ; about 40% greater node-level energy conservation; about 50% longer network lifetime; 10% fewer message exchanges.	Strong routing-security comparator, but it evaluates Sybil mitigation rather than attendance acceptance or exact application replay.
Dhifallah et al. (IoT-RPL 2021) [21]	RPL attack dataset for IDS research	Four Cooja-generated attack classes: Blackhole, Flooding, DODAG Version Number, and Decreased Rank.	Demonstrates broader routing attack coverage; it is a dataset rather than an attendance transaction or replay-mitigation evaluation.
Panimozhi and Murthy [30]	QR attendance with proxy mitigation	Web-based prototype; the publication reports successful load handling without errors in the tested application setup.	Closest application comparator, but its evaluation emphasis is proxy mitigation and web-application operation rather than RPL PDR/ASR coupling.
Present study	Cross-layer attendance transaction over RPL	5 motes: PDR 98.2%, loss 1.8%, delay 12 ms, ASR 96.5%, replay 5/5 rejected; 20 motes: 95.4%, 4.6%, 21 ms, 93.8%, 12/12; 50 motes: 91.7%, 8.3%, 37 ms, 89.4%, 21/21. Aggregate exact replay: 38/38 rejected, exact 95% CI 90.75–100%.	Adds a numerical link between LLN delivery, application acceptance, and a bounded replay result over the same attendance path.

4. System Architecture and Security Model

The system is described as an end-to-end attendance chain rather than as a QR reader connected directly to a database. In a deployment, the attendance service creates or activates a session value that can be represented by a QR code. After the visual interaction, the client forms an application transaction containing the participant identity, session context, freshness-related information, and a token or transaction identifier. The network then carries that payload through the IEEE 802.15.4/6LoWPAN/IPv6/RPL/UDP stack toward the border router or gateway, after which the receiving application applies the attendance rules and duplicate state before committing a record. The simulation begins at the post-scan transaction boundary; it does not model camera recognition or the visual QR encoder itself.

The baseline is designed to answer a narrow security question. It can show whether the same accepted logical attendance transaction is rejected when it is submitted again while the relevant server state is retained. It cannot, without additional mechanisms, establish confidentiality, cryptographic origin authentication, protection against modified replay variants, RPL control plane security, resistance to physical device capture, or backend compromise. Those threats are kept outside the measured claim because the present implementation does not contain the cryptographic keys, routing monitor, trusted hardware, or platform controls needed to test them meaningfully.

4.1. Acceptance Rule and Replay Property

Let an attendance transaction be represented as

$$m_i = \langle u_i, s_i, t_i, x_i \rangle, \quad (1)$$

where u_i is the participant identity, s_i is the attendance-session context, t_i denotes freshness-related information available to the application, and x_i denotes the token or transaction identifier. The server accepts a transaction only when the token is valid, the identity/session combination is authorised, and the logical transaction has not already been accepted for that session:

$$\mathcal{A}(m_i) = V_{tok}(m_i) \wedge V_{auth}(m_i) \wedge [h(m_i) \notin \mathcal{S}_{s_i}], \tag{2}$$

where $\mathcal{S}_{s_i}$ is the retained set of accepted logical identifiers for session s_i. When Equation (2) is true, the attendance record is committed, and the identifier is added to $\mathcal{S}_{s_i}$.

The replay property is therefore bounded. If m_i has already been accepted and the same logical identifier remains in $\mathcal{S}_{s_i}$, an exact replay fails the final predicate and cannot create a second durable attendance record. A crafted transaction that changes fields and produces a different identifier is a different case and was not quantified in the recorded replay experiment. A benign retransmission after an acknowledgement loss is also operationally different from a malicious duplicate: an implementation should treat it idempotently (returning the prior outcome without creating another record), but false positive behaviour for such retries was not measured in the available data. These two cases are now stated explicitly so that exact replay rejection is not confused with general replay-safe protocol design.

4.2. Layered Architecture and Transaction Path

Figure 1 separates the attendance source, low-power network, gateway, and server-side attendance application. The intended processing sequence is (1) a session value is made available to the client through the QR interaction; (2) the post-scan client forms the tokenised attendance payload; (3) UDP carries the payload over IPv6/6LoWPAN, with RPL supplying a route toward the root; (4) the border router or gateway delivers the application payload to the receiving service; and (5) the server evaluates token validity, authorisation, and duplicate state before accepting or rejecting the event. This separation makes it possible to distinguish a packet that never reaches the application from a delivered transaction that fails an attendance rule.

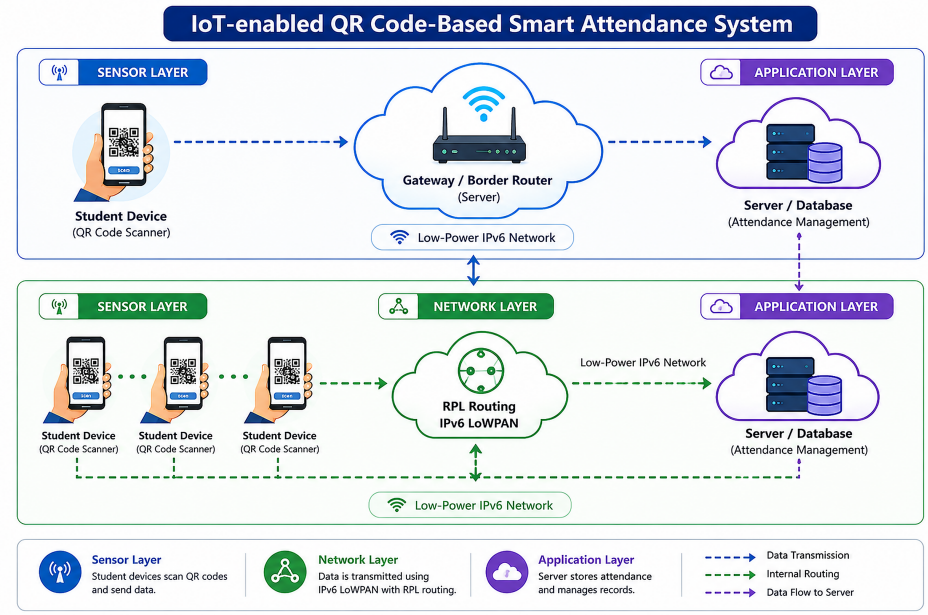


Figure 1. Layered architecture of the QR code attendance system. The simulation begins after the QR interaction. The post-scan transaction is carried through the low-power IPv6/RPL path to the gateway and then evaluated by the attendance application before a durable record is created.

Figure 2 makes the complete chain explicit by including optional intermediate RPL forwarders between the source and the sink/gateway. A node that can reach the root directly may use a one-hop path, whereas a node outside direct reach requires one or more RPL forwarding hops. The architecture therefore permits multi-hop operation. However, the recorded quantitative dataset does not contain per-packet route depth, parent-change logs, or hop-count distributions, so this paper does not claim to measure routing-performance differences between one-hop and multi-hop paths. That limitation is carried into the methodology and future validation plan rather than being filled with assumed routes.

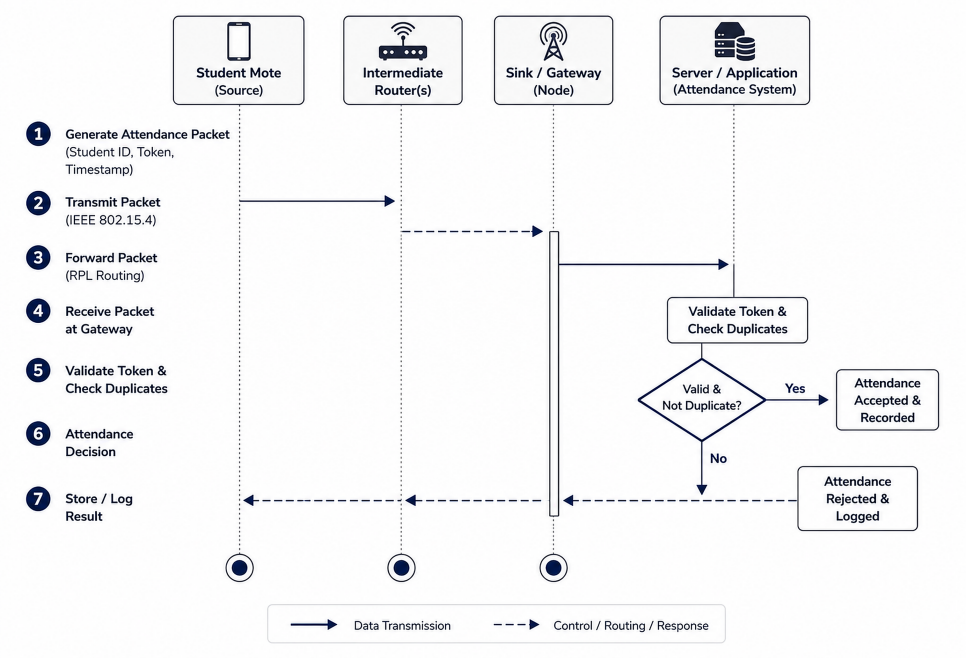


Figure 2. End-to-end transaction sequence from a student mote through optional intermediate RPL router(s), the sink/gateway, and the attendance application.

Figure 3 shows the application-level decision process that takes place after a packet has been delivered successfully. Receipt of the packet alone does not create an attendance record. The server first checks the submitted token, confirms that the user is authorised for the relevant session, and applies duplicate and freshness checks. Only requests that satisfy these conditions are accepted. This distinction also helps explain why packet delivery ratio (PDR) and attendance success rate (ASR) may differ slightly, since successful network delivery does not always result in successful application-level acceptance.

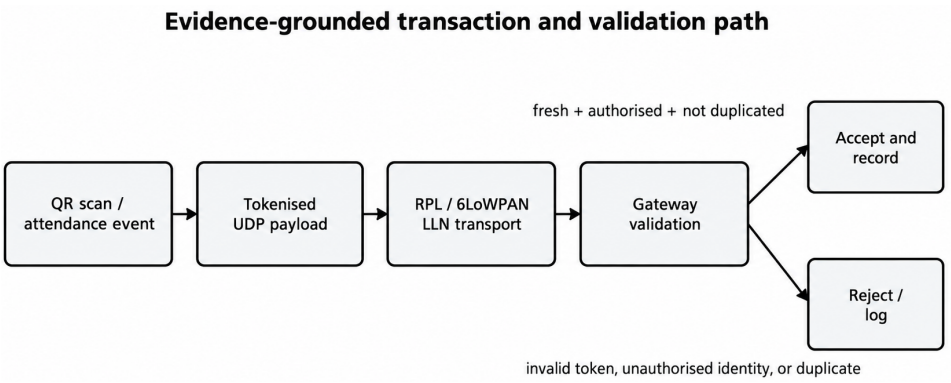


Figure 3. Validation path used to distinguish network delivery from application acceptance and duplicate rejection.

4.3. Threat Model and Attack Surfaces

The threat model distinguishes outsider, insider, and physical-compromise cases. An outsider is assumed able to observe a displayed QR code, capture or reconstruct application traffic within radio reach, submit unauthorised identifiers, and replay previously observed transactions. An enrolled insider possesses a valid identity or device but may attempt to reuse a legitimate attendance event, submit attendance on behalf of another user, or exploit a copied QR value. A physical-compromise adversary may obtain a device and extract locally stored material; this case is considered a deployment risk but is not experimentally defended by the baseline. Gateway compromise and backend administrative compromise are also outside the measured defence boundary because a compromised validation service can alter the acceptance state directly.

The main attack surfaces are the visual QR channel, the client/mote, the IEEE 802.15.4/6LoWPAN/RPL communication path, the gateway, and the attendance backend. Table 4 maps representative attacks to the baseline response and the control still required.

Table 4. Threat model, baseline response, and residual security requirement.

Attack Surface	Adversary Action	Baseline Response	Residual Requirement
Application replay	Re-submit an already accepted attendance transaction	Duplicate/freshness state rejects the repeated logical event in the evaluated exact-replay cases	Authenticated freshness is still needed against crafted variants or state-reset attacks.
QR cloning/forwarding	Photograph or share the visible session code	Visual copying itself is not prevented; the resulting transaction must still pass server-side checks	Short-lived session challenge plus identity/device-bound authenticated response.
Message modification/MITM	Alter identity, token, or session fields in transit	Rule violations may be rejected, but undetected modification cannot be excluded without a cryptographic tag	Authenticated encryption or message authentication with replay-safe nonce/counter handling.
Node impersonation	Submit traffic claiming another participant/device	Invalid identity/session combinations can be rejected; cryptographic proof of possession is not established	Secure enrolment and a device/user-bound key or PUF credential.
RPL manipulation	Replay control messages, alter rank, or induce sinkhole/Hello Flood behaviour	Outside application duplicate logic	RPL-aware monitoring, secure routing extensions, or gateway-side detection.
Physical device capture	Steal an enrolled device or extract local secrets	Not addressed by the simulation baseline	Protected key storage, revocation, device binding, and compromise recovery.
Gateway/backend compromise	Modify validation state or attendance records	Outside the measured defence boundary	Platform hardening, access control, audit logging, and protected backend communication.

For deployment planning, Table 5 adds a qualitative prioritisation. These labels are not measured probabilities; they are engineering judgements based on the ease of attempting the attack in a classroom/workplace setting and the potential consequence if it succeeds. They are included to make the residual risk discussion operational rather than to present an empirical risk model.

Table 5. Qualitative prioritisation of residual deployment risks.

Residual risk	Likelihood	Impact	Deployment interpretation
QR cloning/forwarding	High	Medium–High	The visual code is easy to copy; impact rises when the response is not bound to an enrolled identity/device.
Node impersonation	Medium	High	Requires a usable identity/credential path, but successful impersonation directly affects attendance integrity.
Message modification/MITM	Medium	High	Requires suitable network access; authenticated payload protection is needed to detect undetected field modification.
RPL manipulation/flooding	Medium	High	A nearby attacker can target availability even without defeating application validation.
Physical device capture	Low–Medium	High	Less frequent than QR sharing, but compromise of stored credentials can defeat identity assumptions.
Gateway/backend compromise	Low	Very high	Harder to achieve in normal classroom use, but compromise can alter validation state or records at scale.

A formal protocol tool such as ProVerif, Scyther, AVISPA, or Tamarin would be appropriate once the authenticated transaction format and key-establishment procedure are fully specified. Applying such a tool to the present non-cryptographic duplicate rule would overstate what is being verified. The current paper therefore provides a formal acceptance predicate and a bounded replay argument, while identifying machine-checked authentication/freshness proofs as a requirement for the cryptographic extension [19].

Figure 4 summarises the two application-level security cases examined in this work. The replay workflow is supported by the recorded experimental results, while the unauthorised-identity workflow is included to illustrate the broader threat scenario considered in the security analysis.

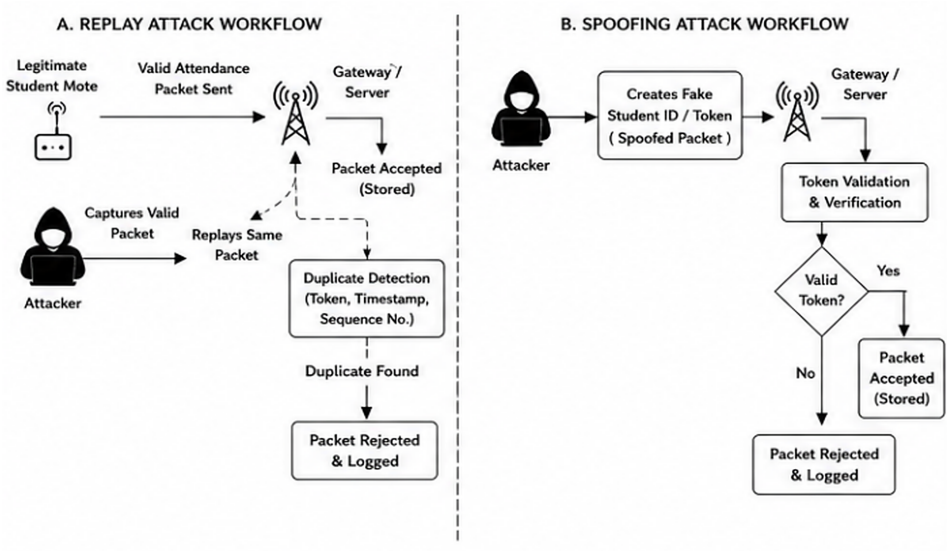


Figure 4. Application-level replay and unauthorised identity workflows considered by the security model. Quantitative attack counts are available for exact replay; unauthorised identity behaviour is interpreted qualitatively.

5. Simulation Methodology and Experimental Evidence

The experimental evaluation examines the recorded behaviour of the attendance mechanism in Contiki-NG/Cooja as the number of participating motes increases. The available

evidence includes scenario-level aggregates for packet delivery, packet loss, average end-to-end delay, attendance success, and exact replay rejection. The analysis is deliberately descriptive: parameters and repeated-run statistics that were not retained with the original experiment are identified explicitly rather than reconstructed from simulator defaults.

5.1. Contiki-NG/Cooja Setup

The retained experiment contains three attendance network sizes: 5, 20, and 50 student motes communicating toward a validation endpoint. Re-examination of the archived setup screenshots recovered several configuration details that were not captured in the earlier text record. Cooja was launched from the Contiki-NG `tools/cooja` directory with `./gradlew run`; the source selection dialogue shows an `attendance-rpl` application containing `udp-client.c` and `udp-server.c`; and the client creation dialogues identify `Cooja Mote Type #2`. The 20- and 50-client setup screenshots record random positioning with $x, y \in [0, 100]$ and $z = 0$. A separate 5-mote setup screenshot records the same coordinate interval when the root/server mote type was added. These values are reported as archived configuration evidence rather than reconstructed defaults.

Figure 5 now uses only the retained Cooja screenshots from the original project archive. The three panels show the 5 20, and 50-mote runs and include the simulator network view together with the mote-output log. In each archived run the validation/root process visibly reports received attendance traffic and repeated "DUPLICATE ATTENDANCE REJECTED" events. These screenshots provide direct visual support for the presence of the duplicate rejection logic in the recorded experiments. They are not used to manufacture counts that are not visible in the raw archive, and they are not treated as evidence of a particular routing tree.

The 20 and 50-client creation dialogues use the same random placement envelope, $x, y \in [0, 100]$ with $z = 0$. Holding the placement area constant while increasing the client count from 20 to 50 raises the nominal client density from 20/10,000 to 50/10,000 clients per square simulation coordinate unit, i.e., a factor of 2.5. This is a useful secondary observation because the larger case is not only a larger population but also a denser nominal placement. It is not presented as a controlled density experiment because the realised coordinates, propagation model, transmission range, and random seed were not retained.

The visual packet arrows and radio message entries in Cooja show active communication among multiple simulated nodes, but they are not a retained history of RPL parent selection or hop depth. Route depth histories, selected parent changes, and per-route statistics were not preserved. The quantitative results are therefore interpreted as application and communication behaviour for the recorded topologies, not as a controlled measurement of multi-hop RPL routing efficiency. Figure 6 presents the archived Cooja setup evidence retained from the original experimental work. It documents the Contiki-NG/Cooja launch environment, the attendance-RPL client and server source files, and the mote-creation settings used for the 20 and 50-client scenarios, including random placement across the recorded x and y-coordinate range of 0–100 with $z = 0$.

Table 6 now distinguishes configuration details recovered directly from the archived screenshots from parameters that still cannot be established. This is important because a reproducibility table should not silently replace missing experimental settings with the defaults of a later Contiki-NG/Cooja release.

The modulation question requires a distinction between the protocol stack and the simulator abstraction. Cooja native motes execute Contiki-NG as native processes and do not precisely emulate a particular mote radio transceiver [5,6]. Consequently, the presence of an IEEE 802.15.4 link stack does not, by itself, establish that a specific RF waveform was explicitly simulated in the archived Cooja run. The selected Cooja radio medium and

operating channel are also absent from the retained evidence. Assigning UDGM, MRM, a particular channel, or a physical modulation waveform after the fact would therefore turn a reproducibility correction into a new experiment.

The two recent references suggested by the reviewer are retained because they illustrate why physical-layer assumptions matter: Ma et al. studied a multi-carrier DCSK scheme under multipath fading [33], whereas Gong et al. examined modulation choices in an IoT-oriented wireless-power system [34]. They are used to motivate explicit PHY reporting, not to reconstruct the original Cooja configuration. A future rerun should therefore save the `.csc` scenario, pin the Contiki-NG commit, record the radio-medium class and parameters, log the selected channel, and preserve the build-time configuration output together with the raw per-packet trace.

Table 7 separates the evidence actually present in the archive from quantities that would require a new experiment. This distinction was added because several reviewer concerns depend on information that cannot be recovered reliably from an image or from simulator defaults.

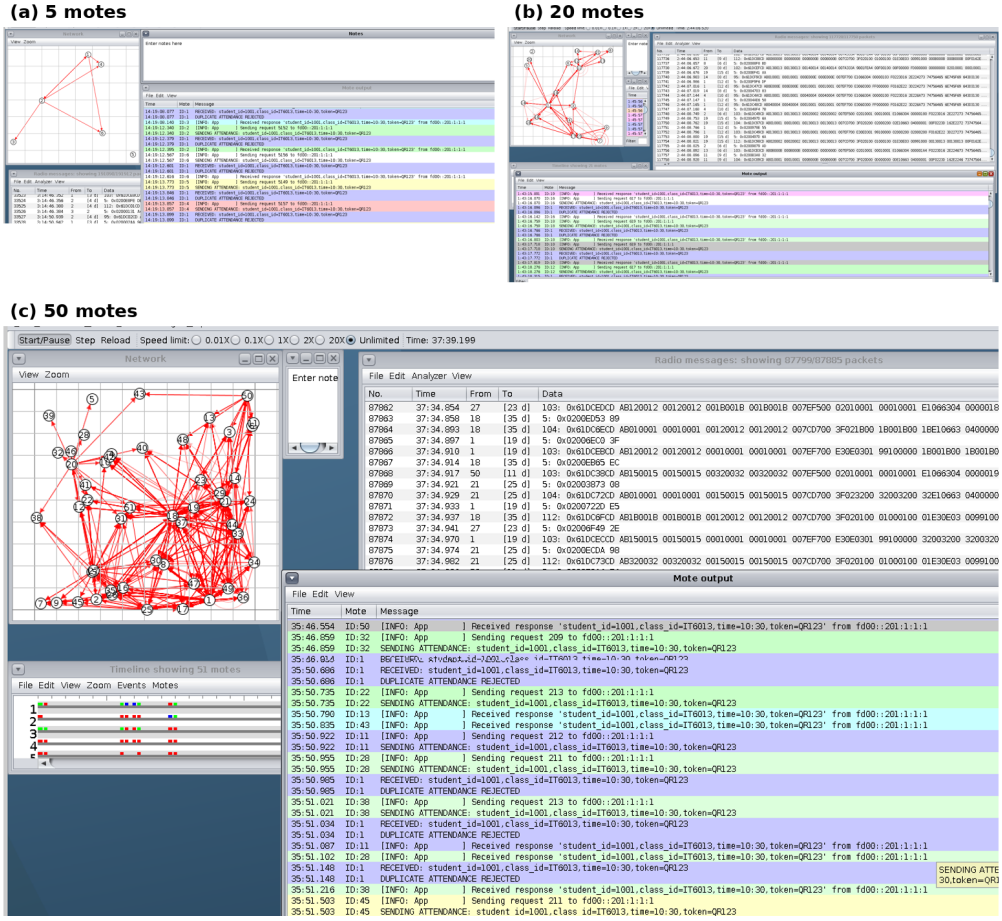


Figure 5. Archived Cooja evidence from the retained 5, 20, and 50-mote runs. The mote-output panes show repeated duplicate attendance rejection at the validation/root process, while the network/radio panes show active simulated communication. The screenshots are used as provenance evidence only; they are not used to infer unrecorded route depth, radio parameters, or additional numerical measurements.

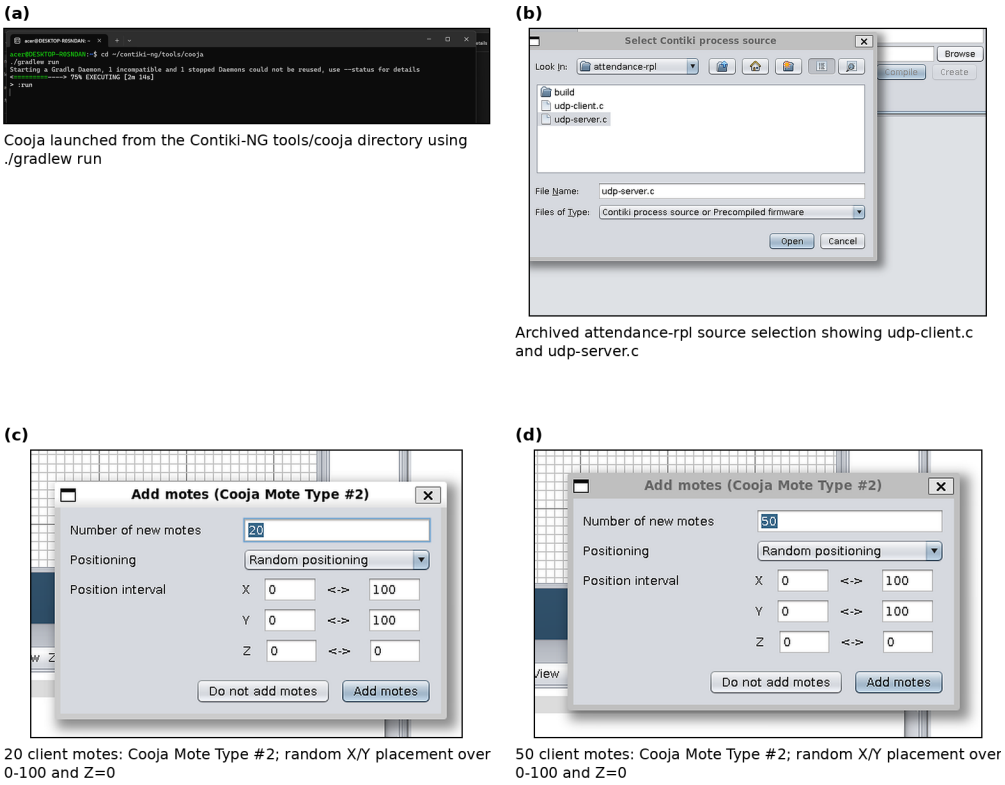


Figure 6. Archived Cooja setup evidence recovered during revision: (**a**) launch from the Contiki-NG Cooja directory using Gradle; (**b**) the `attendance-rpl` client/server source files; and (**c**,**d**) the 20- and 50-client Cooja-mote creation dialogues showing random x/y placement over 0–100 with $z = 0$.

Table 6. Experimental configuration and availability of reproducibility parameters after re-checking the archived Cooja setup evidence.

Element	Recorded Configuration/Evidence
Simulation platform	Contiki-NG with Cooja. Archived terminal evidence shows Cooja launched from `tools/cooja` with `./gradlew run`; the exact Contiki-NG release/commit was not retained.
Application source	Archived source selection dialogue shows the `attendance-rpl` project with `udp-client.c` and `udp-server.c`. The source files themselves are not present in the supplied revision archive.
Mote implementation	Cooja native mote types are shown in the setup dialogues: mote type #1 for the root/server setup and mote type #2 for clients.
Wireless/link stack	IEEE 802.15.4 with CSMA.
Adaptation/network/routing	6LoWPAN, IPv6, RPL Lite.
Transport	UDP.
Radio propagation model (e.g., UDGM/MRM)	The selected Cooja radio-medium class is not visible in the retained screenshots, and no `.csc` scenario file is present; therefore, it remains unverified.
PHY modulation/operating channel	The archived experiment uses Cooja native motes rather than a retained hardware-emulation profile. No explicit RF waveform/modulation setting or operating-channel value is visible in the archive; neither is assigned retrospectively.
Transmission power/communication radius	Not retained.

Table 6. *Cont.*

Element	Recorded Configuration/Evidence
Node placement	Random positioning is explicitly recorded. For the 20- and 50-client creation dialogues, x and y are each 0–100 and $z = 0$; the 5-mote setup archive shows the same range for the root/server creation step. Exact realised coordinates and random seeds were not retained.
Nominal client density in the common 20/50 placement envelope	20-client case: 0.002 clients per square simulation coordinate unit; 50-client case: 0.005 clients per square simulation coordinate unit (2.5× higher). This is a derived density indicator, not a controlled propagation experiment.
RPL control timing (DIO/DAO parameters)	Not retained beyond use of RPL Lite.
Application traffic interval/offered load	Not retained as a reproducible numerical schedule.
Application packet/payload size	Not retained as a reproducible byte count.
Network sizes	5, 20, and 50 student motes with one central validation/root endpoint.
Application abstraction	Post-scan tokenised attendance transaction.
Security cases	Normal attendance traffic, exact repeated attendance transactions, and unauthorised/spoofed identity behaviour.
Application checks	Token validation, attendance-rule validation, and duplicate/freshness state.
Recorded metrics	PDR, packet loss, average end-to-end delay, ASR, replay attempts, and replay rejections.
Random seeds/repeated independent runs	Not retained; one aggregate record is available for each network size.
Energy instrumentation	Not recorded; no Energest/Powertrace duty-cycle data or current/voltage conversion inputs are available.

Table 7. Evidence-to-claim boundary used in the revised analysis.

Question	Evidence Retained	Permitted Interpretation
Scale and nominal density	5-, 20-, and 50-mote runs; common 0–100 by 0–100 random-placement envelope is visible for 20 and 50 clients	Descriptive comparison across the three recorded populations; 50 clients have 2.5× the nominal density of 20 clients in the common placement envelope.
Delivery/application outcome	Scenario-level PDR, packet loss, mean delay, and ASR	Descriptive point estimates only; no significance claim without independent repeated runs.
Exact replay	5/5, 12/12, and 21/21 injected/rejected counts plus archived log evidence of duplicate rejection	Exact-replay rejection can be quantified with binomial intervals; modified replay and benign retransmission remain outside the measured claim.
Routing detail	Network/radio activity is visible in Cooja screenshots	No route depth, parent-change, or multi-hop efficiency claim because those histories were not retained.
Tail latency and saturation	Mean delay only	No p95/max-latency or saturation threshold claim.
Energy/resource cost	No Energest/Powertrace counters, current/voltage model, RAM/flash profile, or per-run duty cycles	No numerical energy efficiency claim; only the measurement method for a repeatable follow-on experiment is specified.

5.2. Parameter Scope and Reproducibility

The QR generation frequency, token lifetime, application timeout, and cryptographic parameters are not treated as measured independent variables because the recorded experiment begins at the post-scan transaction boundary and the baseline does not implement authenticated encryption. Assigning concrete values after the fact would create an artificial parameter set unrelated to the recorded outcomes. A deployment-oriented rerun should specify the QR/session lifetime, transaction freshness rule, retry timeout, authenticated

payload nonce construction, tag size, key lifetime, and enrolment procedure before data collection begins.

A fully reproducible rerun would additionally require executable application source, complete Cooja scenario files, the Contiki-NG version/commit, topology and radio parameters, random seeds, traffic schedule, and raw per-run logs. Those materials are not available as a complete executable set for the present evaluation. The manuscript therefore reports only the recorded 5-, 20-, and 50-mote evidence and transparent mathematical transformations of those values.

5.3. Evaluation Metrics

The communication and application metrics have different denominators and should not be treated as interchangeable. Let N_{int} denote the number of intended unique attendance events in the application workload. Let N_{tx} be the number of legitimate application-packet transmissions counted for the delivery measurement and N_{rx} the corresponding packets received at the validation endpoint. If retries are present, N_{tx} need not equal N_{int}. Likewise, a received packet is not automatically a successful attendance event because the application may reject it after token, authorisation, freshness, or duplicate checking.

Packet delivery ratio (PDR) and packet loss are

$$\mathrm{PDR}(\%) = \frac{N_{rx}}{N_{tx}} \times 100, \tag{3}$$

$$\mathrm{Loss}(\%) = 100 - \mathrm{PDR}. \tag{4}$$

For a successfully delivered packet j with send and receive times of $t_{tx}^{(j)}$ and $t_{rx}^{(j)}$, the recorded average end-to-end delay is

$$\overline{D} = \frac{1}{N_{rx}} \sum_{j=1}^{N_{rx}} \left(t_{rx}^{(j)} - t_{tx}^{(j)} \right). \tag{5}$$

If N_{acc} intended attendance events are durably accepted, the attendance success rate (ASR) is

$$\mathrm{ASR}(\%) = \frac{N_{acc}}{N_{int}} \times 100. \tag{6}$$

For replay testing, with N_{rep} injected exact replay events and N_{rej} rejected events, the replay rejection rate (RRR) is

$$\mathrm{RRR}(\%) = \frac{N_{rej}}{N_{rep}} \times 100. \tag{7}$$

The delivery-to-application gap is

$$G_{PA} = \mathrm{PDR} - \mathrm{ASR}. \tag{8}$$

PDR and ASR would be numerically identical only under restrictive conditions: one transmission per intended event, every delivered packet corresponding to one intended event, and every delivered event being accepted. The present architecture does not assume those conditions. Packet loss can reduce both measures, while application rejection, duplicate handling, retry behaviour, or differences between packet and event counting can create a residual gap after delivery. Because the retained data contain only aggregates rather than packet-to-event traces, the contribution of each mechanism to G_{PA} cannot be isolated.

5.4. Statistical Treatment and Metric Availability

The retained PDR, loss, delay, and ASR values are single scenario-level aggregates for each network size; independent repeated runs, per-packet samples, random seed distributions, and variances are not available. The differences between 5, 20, and 50 motes are therefore reported as descriptive changes rather than statistically significant effects. Confidence intervals, standard deviations, or hypothesis tests for those four metrics would require data that were not recorded and are not reconstructed. Replay testing is different because the attempt/rejection counts are available, allowing exact two-sided Clopper–Pearson 95% intervals to be reported for the replay rejection proportion.

Table 8 also makes clear that tail latency, direct energy, memory overhead, throughput, routing-depth information, and a saturation threshold are unavailable. In particular, the mean delay cannot reveal a maximum nor 95th-percentile delay nor can three network size points at one unrecorded offered-load schedule establish where a capacity knee or saturation point occurs. Those questions require an offered-load sweep, per-packet delay distributions, throughput/queue or retry evidence, and repeated independent seeds. These measures are requirements for the repeated-run validation stage rather than quantities inferred from the available aggregates.

Table 8. Availability and interpretation of performance and security metrics.

Metric	Status	Use/Limitation in This Paper
PDR/packet loss	Single aggregate per scale	Descriptive communication reliability across 5, 20, and 50 motes; no run-to-run CI or significance test.
Average end-to-end delay	Single aggregate per scale	Mean responsiveness only; maximum, distribution, and 95th percentile were not retained.
Attendance success rate	Single aggregate per scale	Application-level success after communication and validation; no run-to-run variance.
Replay attempts/rejections	Event counts	RRR with exact binomial confidence intervals.
Throughput	Not measured	Required in controlled offered-load experiments before throughput claims are made.
Energy consumption	Not measured	Requires radio/CPU duty-cycle counters and device power parameters before an energy claim can be made.
RAM/flash overhead	Not measured	Required to quantify the cost of authenticated encryption and routing defences.
Routing depth/parent changes	Not retained	Required to compare direct and multi-hop RPL behaviour.
Saturation threshold/capacity knee	Not measured	No saturation point is claimed. It requires controlled offered-load sweeps, per-packet tail latency, throughput/queue or retry traces, and repeated seeds.

6. Results and Evaluation

This section examines the behaviour of the proposed attendance architecture across the evaluated RPL-based network configurations. The analysis considers both communication-level performance and application-level outcomes so that changes in network behaviour can be related directly to the successful completion of attendance transactions. The reported results are based on the measurements obtained from the evaluated Contiki-NG/Cooja scenarios rather than on assumed or extrapolated values. The evaluation focuses on packet delivery ratio, packet loss, end-to-end delay, attendance success rate, and replay rejection behaviour. Particular attention is given to how these measures change as the number of participating motes increases and whether deterioration at the network layer is reflected in the attendance application. Replay experiments are analysed separately to determine how consistently repeated transactions are rejected under the recorded test conditions. Together,

these results describe the recorded scalability and security behaviour; they are not used as inferential evidence about unobserved network sizes or repeated runs.

6.1. Network and Application Outcomes

Table 9 reports the recorded communication, attendance, and replay rejection results for the three evaluated network sizes.

Table 9. Recorded communication, attendance, and replay rejection results.

Motes	PDR (%)	Loss (%)	Avg. Delay (ms)	ASR (%)	Replay Rejected/Injected
5	98.2	1.8	12	96.5	5/5
20	95.4	4.6	21	93.8	12/12
50	91.7	8.3	37	89.4	21/21

The results in Table 9 provide a consolidated view of both network-level and application-level behaviour across the three evaluated scenarios. Although performance decreases as the number of motes increases, the recorded values show that packet delivery and successful attendance processing remain closely aligned throughout the evaluated range. The following analysis considers these trends individually to clarify how network growth affects reliability, delay, and attendance completion. Figure 7 plots PDR and ASR on the same scale. The two measures remain close at each network size, suggesting that application success broadly follows the ability of the network to deliver the underlying transaction rather than diverging sharply because of application rejection.

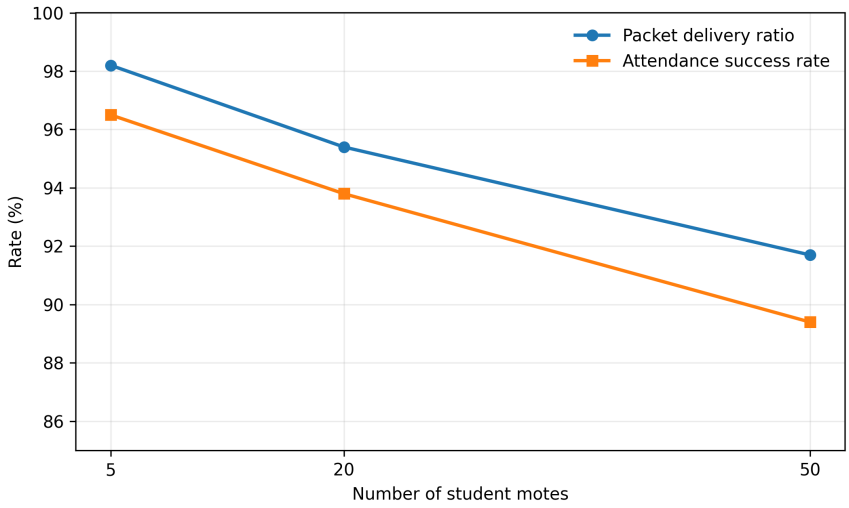


Figure 7. Recorded packet delivery ratio and attendance success rate as network population increases.

The PDR–ASR gaps are 1.7, 1.6, and 2.3 percentage points for 5, 20, and 50 motes, respectively. The decline in PDR can plausibly lower ASR because an attendance event that is not delivered cannot be accepted. The remaining PDR–ASR difference, however, should not be assigned automatically to radio collisions or to application logic. Radio collisions, queue pressure, and retransmissions primarily affect whether traffic is delivered, whereas token/authorisation/freshness checks can reject a transaction after delivery; retry and packet-to-event counting can also make the two denominators differ. The retained aggregate data contain no collision counters, queue traces, retry-to-event mapping, or per-transaction rejection reasons, so these causes cannot be separated quantitatively. The cross-metric result is therefore an observed association: ASR broadly follows PDR, with a small residual gap whose origin remains unresolved.

6.2. Packet Loss and Delay Scaling

Packet loss rises by 6.5 percentage points between the smallest and largest recorded scenarios, matching the absolute PDR decrease. Figure 8 shows this descriptive change. Greater contention is one possible explanation as more motes share the medium, but collision, queue overflow, selected parent, and per-packet retransmission traces were not retained. The manuscript therefore does not identify a single root cause for the loss increase or claim that the three aggregate points establish a general scaling law.

Average delay increases from 12 to 37 ms, a 25 ms absolute increase and an approximately 208.3% relative increase from the 5-mote baseline. Figure 9 shows the recorded means. These values do not describe tail latency: the per-packet delay distribution, maximum delay, and 95th percentile were not retained. Importantly, the manuscript does not claim that saturation was observed. The three recorded network sizes show only that the mean delay is higher at the larger recorded scenarios. Locating a capacity knee would require a controlled offered-load sweep together with throughput, queue/retry behaviour, and per-packet delay distributions over repeated seeds. The 50-mote result is therefore reported only as the largest observed operating point under the recorded workload, not as a saturation threshold or maximum supported size.

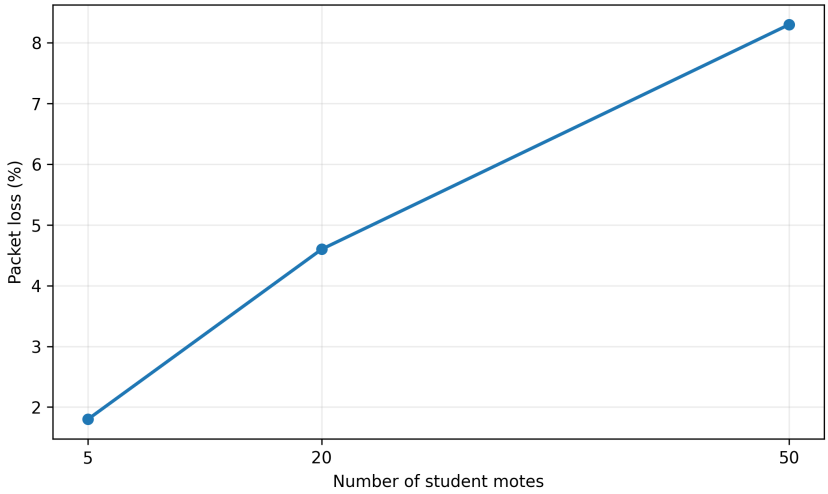


Figure 8. Recorded packet-loss percentage for the three evaluated network sizes.

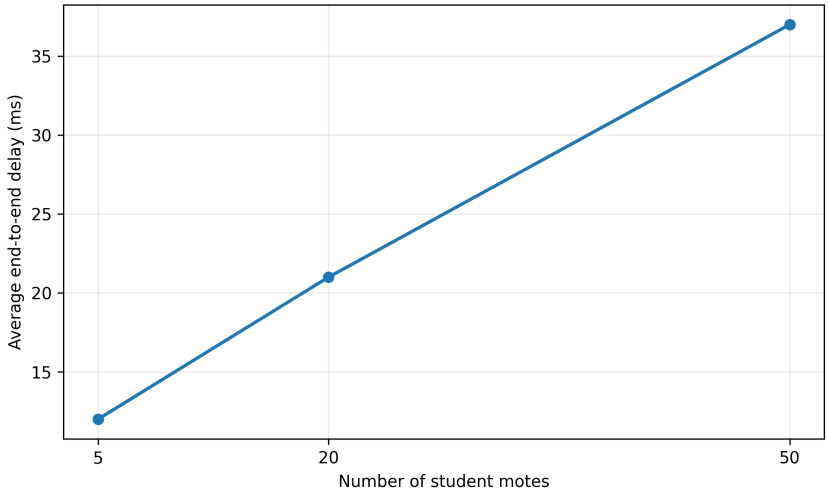


Figure 9. Recorded average end-to-end delay as network population increases.

Table 10 summarises the descriptive changes in network and application performance between the 5–50 mote scenarios. The comparison highlights the observed changes in packet delivery, packet loss, average end-to-end delay, attendance success rate, and the PDR–ASR gap across the evaluated range.

Table 10. Derived descriptive changes in network and application performance from 5 to 50 motes.

Performance Measure	Change (5–50 Motes)	Interpretation
Packet delivery ratio (PDR)	−6.5 percentage points (−6.62% relative)	Delivery decreases with network size, while remaining above 90% at the largest evaluated point.
Packet loss	+6.5 percentage points	Loss increases as the network grows; causal attribution requires per-run link and queue counters that are not available.
Average end-to-end delay	+25 ms (+208.3% relative)	Delay has the largest proportional change but remains below 40 ms in the 50-mote scenario.
Attendance success rate (ASR)	−7.1 percentage points (−7.36% relative)	Application success decreases at a rate close to the decline in PDR.
PDR–ASR gap	1.6–2.3 percentage points	Application success remains close to network delivery throughout the tested scale.

These results answer RQ1 within the measured range. Increasing population is associated with lower delivery and ASR and with higher delay, but the service remains operational at 50 motes. The study does not claim that 50 is the maximum practical network size; it is the largest demonstrated operating point available in the experiment.

6.3. Replay Rejection Analysis

Replay testing provides the strongest event-level security evidence. Table 11 reports both the observed rates and exact confidence intervals. The observed rejection rate is 100% in each scenario, but the intervals are intentionally shown because a finite number of successes does not prove perfect rejection for all future traffic. Across all 38 replay attempts, the exact two-sided 95% interval is 90.75–100%. Figure 10 provides a direct visual comparison of the replay attempts injected and rejected in the 5, 20, and 50-mote scenarios. The figure complements the numerical results in Table 11 and shows that every recorded replay attempt was rejected in each evaluated network size.

The replay result is consistent with the acceptance rule in Equation (2): once the transaction identifier is present in the retained duplicate state, resubmission of the same logical event fails the novelty check. This answers the demonstrated part of RQ2. It does not establish protection against a modified payload that produces a different identifier, a valid credential used by an imposter, a compromised endpoint, or RPL control plane manipulation. Those residual cases are addressed in the threat matrix and deployment extension.

Unauthorised or spoofed identity behaviour is represented by the application workflow, but a complete attempt/rejection count is not available for that class. Reporting a numerical spoofing detection percentage would therefore be unsupported. The defensible conclusion is qualitative: the baseline contains an authorisation predicate intended to reject identity/session combinations that do not satisfy the attendance rule; cryptographic proof of identity remains outside the measured design.

Table 11. Replay injections, rejections, observed rates, and exact two-sided 95% binomial confidence intervals.

Motes	Injected	Rejected	Observed RRR	Exact 95% CI
5	5	5	100%	47.82–100%
20	12	12	100%	73.54–100%
50	21	21	100%	83.89–100%
Total	38	38	{100%	90.75–100%

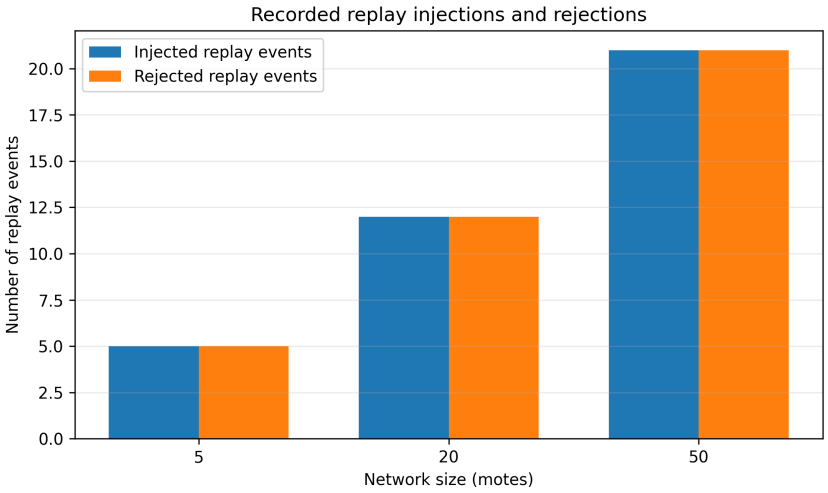


Figure 10. Recorded (replay injections and rejections. The plotted bars use the observed event counts only; no interpolated attack data are introduced.

6.4. Demonstrated Deployment Envelope

Table 12 states the strongest deployment claim supported by the measured data. This form is preferable to inferring a maximum population from three aggregate points. The 50-mote scenario demonstrates that the tested workload can sustain 91.7% PDR, 37 ms average delay, and 89.4% ASR. Whether a real deployment can support 75, 100, or several hundred simultaneous clients cannot be answered without additional network sizes, repeated seeds, and offered-load experiments.

Table 12. Demonstrated operating envelope and unresolved scalability questions.

Aspect	Largest Demonstrated Point	What Remains to Establish a Practical Maximum
Population	50 motes	Additional sizes and controlled offered-load sweeps
PDR	91.7%	Repeated runs, route depth analysis, link/collision statistics
Average delay	37 ms	Delay distribution, tail latency, timeout interaction
ASR	89.4%	Per-transaction failure attribution and application retry behaviour
Replay test	21/21 rejected at 50 motes	Crafted variants, benign retransmissions, longer attack campaigns
Energy/memory	Not measured	Direct radio/CPU energy and RAM/flash profiling on baseline and secured variants

7. Discussion and Results Interpretation

The recorded results show two related but distinct behaviours. First, PDR and ASR decline together as the scenario grows from 5 to 50 motes. This is consistent with the practical dependency of attendance success on delivery: an event that never reaches the receiving application cannot be accepted. Second, the residual PDR–ASR gap remains small (1.6–2.3 percentage points), but the available aggregates do not identify why a particular

event lies in that gap. Post-delivery rejection is one explanation; differences between packet and event counting, retries, and timing are others. Collisions or queue pressure may contribute to the reduction in PDR itself, but without the corresponding counters they cannot be isolated as the cause of the PDR–ASR difference. This distinction avoids attributing a cross-layer observation to a mechanism that was not measured.

The security result is similarly bounded. The retained duplicate state prevents the same accepted logical transaction from being counted twice in the exact-replay tests, but the QR image remains observable and is not cryptographically bound to a particular device. A copied QR value can therefore still be forwarded. Node impersonation and message modification also remain possible in principle because the baseline has no cryptographic proof that an apparently valid message originated from the claimed device. RPL manipulation is a separate availability problem: an attacker can disturb routing without ever creating an application duplicate. These threats were excluded from the quantitative claim because the implementation does not contain the mechanisms needed to test resistance to them, not because they are considered unimportant.

Authenticated encryption is a natural next step for the attendance payload. Ascon-AEAD128 is standardised by NIST for constrained devices [17], whereas AES-CCM is a mature authenticated encryption mode specified independently in RFC 3610 [35]. Either approach could provide confidentiality and integrity when combined with sound nonce management and secure key enrolment. Table 13 states the comparison only at the specification level; no performance winner is claimed because neither mechanism was implemented in the recorded experiment.

A meaningful implementation comparison must measure packet expansion, execution time, RAM/flash use, radio airtime, and energy on the target platform. This keeps the proposed cryptographic extension separate from the evidence obtained with the present non-cryptographic duplicate-state baseline.

The archived 20-mote material shows active radio exchanges among the 21 simulated motes, but the retained screenshots do not preserve RPL parent histories or per-packet hop counts. The red packet arrows visible in Cooja should therefore not be treated as a routing tree. The end-to-end architecture includes the complete chain from source, through any selected RPL forwarders, to the border router/gateway and server, but the present aggregate results are not a controlled measurement of multi-hop routing efficiency. A dedicated routing study should control hop depth and density, retain parent changes and route depth, and compare direct and multi-hop paths over independent seeds. This is also necessary before conclusions about larger classrooms, denser workplaces, or a capacity threshold can be supported.

Table 13. Specification-level comparison of candidate authenticated encryption options for the follow-on design.

Property	Ascon-AEAD128	AES-CCM	Relevance to This Work
Standard	NIST SP 800-232 (2025) [17]	RFC 3610 (2003) [35]	Both provide authenticated encryption; implementation cost must be measured on the selected platform.
Key/nonce	128-bit key and 128-bit nonce	AES key; nonce length is $15 - L$ octets, with $L = 2 \ldots 8$	Nonces must be unique under a key; nonce construction must be specified before replay/authentication claims are extended.
Authentication tag	128-bit tag in the base Ascon-AEAD128 specification	Configurable 4, 6, 8, 10, 12, 14, or 16 octets	Tag size changes packet expansion and therefore radio airtime.
Status in this study	Not implemented	Not implemented	No measured latency, memory, packet size, or energy comparison is reported.

Practical deployment adds effects that the simulation does not reproduce completely. Human occupancy can shadow or reflect radio paths, nearby Wi-Fi and Bluetooth activity can alter interference, devices may have different radios and power-management behaviour, and a gateway may be temporarily unavailable. A production design should distinguish "received" from "accepted" status, make duplicate handling idempotent during legitimate retries, protect stored attendance data, and provide a fallback procedure for connectivity failures.

Energy is another area where an estimate would be misleading without the underlying counters. Contiki-NG can expose software activity counters such as CPU, low-power, radio-transmit, and radio-listen time, but a numerical energy value also needs a declared platform voltage and state current model. For a future controlled run, the estimator would be

$$E = V(I_{\mathrm{CPU}}T_{\mathrm{CPU}} + I_{\mathrm{LPM}}T_{\mathrm{LPM}} + I_{\mathrm{TX}}T_{\mathrm{TX}} + I_{\mathrm{RX}}T_{\mathrm{RX}}), \tag{9}$$

with energy per accepted attendance event reported as

$$E_{\mathrm{acc}} = \frac{E}{N_{\mathrm{acc}}}. \tag{10}$$

The activity times in Equations (9) and (10) were not retained for the present runs, and the Cooja-native mote screenshots do not identify a physical radio whose current values could be assigned defensibly. A numeric energy estimate would therefore depend on assumptions that were not part of the recorded experiment. Instead of inserting a speculative value, the revised validation plan now requires the raw activity counters, the selected hardware current/voltage data, joules per accepted event, RAM/flash footprint, and packet size expansion for every security configuration.

8. Limitations, Threats to Validity, and Deployment Validation

The findings should be interpreted within the boundary of the recorded simulation. QR image capture and decoding are not simulated; the experiment begins with the post-scan attendance transaction. The 20-mote Cooja view is also a shallow, largely direct-to-root topology, and the available data do not include per-packet hop counts or route depth distributions. The reported results therefore describe the recorded communication/application behaviour, not a general measurement of multi-hop RPL routing performance.

8.1. Limitations and Threats to Validity

Statistical generalisation is limited because PDR, loss, delay, and ASR are represented by one aggregate value for each network size. Independent repeated runs, seed distributions, per-packet delays, and variances were not retained. Consequently, confidence intervals, standard deviations, effect-size estimates, and significance tests are not reported for those metrics. Replay events are an exception because attempt/rejection counts are available, permitting exact binomial confidence intervals. Even there, 38/38 exact-replay rejections remain a finite-sample result rather than proof of perfect future rejection.

External validity is bounded by the three recorded network sizes and the incomplete physical-layer configuration record. The evaluation stops at 50 motes and does not contain intermediate population points, controlled offered-load sweeps, repeated seeds, or tail latency distributions. The revision recovered the same random placement bounds for the visible 20- and 50-client creation dialogues ($x, y = 0$–100, $z = 0$), which means the 50-client case has 2.5 times the nominal client density of the 20-client case within that common placement envelope. This derived observation strengthens the description of the tested range, but it is not a substitute for controlled sparse-versus-dense experiments because the realised coordinates, random seeds, selected radio-medium class, operating

channel, and transmission power/range were not retained. These omissions prevent a defensible extrapolation to larger populations and prevent the present data from identifying a saturation threshold.

The attack evaluation is also intentionally narrow. Exact application replay is measured; modified replay variants, false positives under benign retransmission, cryptographic message modification, node impersonation, RPL control plane attacks, jamming, physical compromise, gateway/backend compromise, and side-channel attacks are not experimentally defeated by the baseline. Their exclusion reflects missing mechanisms and evidence, not low importance. Table 5 therefore treats several of these residual threats as high-impact deployment concerns.

Energy efficiency is not an experimental finding of the present study. Direct CPU activity, radio duty cycle, per-transaction energy, and memory footprint were not recorded, so neither measured nor simulated energy estimates are reported. The absence of hardware testing is likewise a material limitation. Before operational claims are made, the architecture should be evaluated on constrained devices and a real border router under representative interference, occupancy, and mobility conditions.

8.2. Deployment-Oriented Validation Plan

A stronger second-stage evaluation should preserve the existing 5-, 20-, and 50-mote points for continuity and add at least 10, 30, 40, 60, 75, and 100 clients, with controlled sparse/dense placement envelopes and explicit one-hop, two-hop, and deeper RPL topologies. Each population/topology/load combination should be repeated over at least five independent random seeds. The primary outputs should include PDR, loss, throughput, complete per-packet delay distributions (including maximum and 95th percentile), ASR, route depth/parent changes, radio and CPU activity, estimated energy, and RAM/flash footprint. A capacity threshold should be declared only after a pre-specified service criterion is crossed consistently across repeated runs, rather than inferred from a single point. Repetition would then support means, dispersion, confidence intervals, and effect sizes instead of one aggregate value per scale.

Security experiments should separate application and routing attacks. Application tests should include exact replay, replay with modified fields, forged/unauthorised identity, altered authentication tags, and benign retransmissions that resemble duplicates. Routing tests should include RPL replay, rank manipulation, and feasible flooding/DoS cases. The secured application configurations should compare the current baseline with authenticated payload protection (for example, Ascon-AEAD128 and an AES-CCM baseline), while keeping secure enrolment as an explicit prerequisite for origin authentication. Once the cryptographic message flow is fully specified, formal verification can test freshness and authentication properties independently of the network simulation.

The hardware stage should use constrained nodes and a real border router in a classroom-like environment. It should introduce controlled Wi-Fi/Bluetooth interference and human occupancy, record the same application and network outcomes, and measure power directly. The practical workflow should also test QR/session rotation, acknowledgement loss, benign retries, gateway outages, and user feedback so that duplicate protection does not degrade legitimate attendance submission. Table 14 lists the information that should be retained for a reproducible extension of the present evaluation. The listed artefacts cover the implementation files, Cooja configuration, traffic and routing settings, security parameters, attack conditions, energy measurements, and the raw logs and analysis scripts required to repeat and examine the experiments independently. Table 14 summarises the minimum reproducibility artefacts required for the extended evaluation. These include the implementation files, simulation settings, traffic and routing configuration, security

and attack parameters, energy measurements, and the raw logs and analysis scripts needed to reproduce the experiments.

Table 14. Minimum reproducibility artefacts for the extended evaluation.

Artefact	Required Content
Contiki-NG source	Client, gateway/server, attack generator, and security configurations.
Cooja scenarios	Node types, positions, radio model, run duration, seed, and controlled route depth.
Build metadata	Contiki-NG version/commit, toolchain, and compile flags.
Traffic definition	Message interval, attendance window, payload size/format, offered load, and retry rule.
PHY/link configuration	Channel, modulation/PHY assumptions, transmission power/range, and propagation model.
RPL configuration	DIO/DAO timing, objective function, parent-selection settings, and route depth logs.
Security parameters	Token/session lifetime, nonce/counter construction, key/tag sizes, and enrolment assumptions.
Attack definition	Attacker capability, start/stop time, injected counts, replay source, and benign-retry controls.
Energy instrumentation	CPU/radio/low-power counters, platform current/voltage assumptions, and joules per accepted event.
Raw logs and scripts	Per-run packet, route, validation, attack, energy logs, and reproducible analysis code.

9. Conclusions

This paper examined a post-scan QR attendance transaction carried over an RPL-based low-power IoT stack in Contiki-NG/Cooja. Across the three recorded network sizes, PDR decreases from 98.2% to 91.7%, packet loss rises from 1.8% to 8.3%, average end-to-end delay increases from 12 to 37 ms, and ASR decreases from 96.5% to 89.4%. The PDR–ASR gap remains between 1.6 and 2.3 percentage points, showing that application success broadly follows delivery in the recorded scenarios. Because only one aggregate record is available per network size, these values describe the observed test points rather than statistically significant scaling effects or a maximum supported population.

All 38 recorded exact replay attempts are rejected by the retained duplicate/freshness state. The aggregate exact 95% confidence interval is 90.75–100%, and the result is intentionally limited to exact replay under the recorded conditions. It does not establish immunity to modified replay variants, node impersonation, message modification, routing attacks, physical compromise, or backend compromise. The formal acceptance predicate makes that boundary explicit.

The revision also clarifies the topology and evidence boundary. The retained Cooja screenshots directly show the 5-, 20-, and 50-mote runs and duplicate rejection at the validation/root process. The visible 20- and 50-client setup dialogues use the same 0–100 by 0–100 random placement envelope, so the 50-client case represents 2.5 times the nominal client density of the 20-client case. Even so, the screenshots do not preserve parent histories or hop-count distributions, and the current data should not be interpreted as a quantitative evaluation of multi-hop RPL routing efficiency. Future work should therefore combine controlled multi-hop topologies and repeated seeds with tail latency, throughput, route depth, memory, and direct energy measurements. Authenticated payload protection and secure enrolment should then be evaluated alongside routing-aware defences, followed by hardware experiments under realistic interference and occupancy. This staged validation would show not only whether stronger protection works, but also what it costs on the constrained devices and network that carry the attendance transaction.

Author Contributions: Conceptualization, R.M. and N.I.S.; methodology, R.M.; simulation and investigation, R.M.; formal analysis, R.M.; validation, R.M. and N.I.S.; visualisation, R.M.; writing—

original draft preparation, R.M.; writing—review and editing, R.M. and N.I.S.; supervision, N.I.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. The study is based on network simulation and does not involve human participants or identifiable human subject data.

Data Availability Statement: Aggregate numerical results and representative Cooja figures used in the analysis are included with the manuscript materials. A complete executable reproduction package containing the application source, Cooja `.csc` scenarios, per-run raw logs, and random seed records is not currently available; accordingly, the paper does not claim independently reproducible reruns from those materials alone. Future experimental releases should deposit these artefacts together with analysis scripts in a public research repository.

Acknowledgments: During manuscript revision, the authors used AI-assisted tools for language editing and for redrawing a conceptual validation-flow diagram. No experimental measurements, Cooja screenshots, replay counts, network-performance values, or statistical results were generated or altered by these tools. The authors checked the revised technical content and references against the retained project evidence and source publications and take full responsibility for the accuracy, integrity, and final version of the work.

Conflicts of Interest: N.I.S. is listed as a Guest Editor of the Special Issue to which this manuscript is intended to be submitted. To avoid any editorial conflict, the authors request independent handling of the manuscript by an editor with no conflict of interest. The authors declare no other conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AEAD	Authenticated Encryption with Associated Data
AES	Advanced Encryption Standard
ASR	Attendance Success Rate
BLE	Bluetooth Low Energy
CI	Confidence Interval
CSMA	Carrier-Sense Multiple Access
IoT	Internet of Things
IPv6	Internet Protocol Version 6
LLN	Low-Power and Lossy Network
MITM	Man-in-the-Middle
NFC	Near-Field Communication
NIST	National Institute of Standards and Technology
OTP	One-Time Password
PDR	Packet Delivery Ratio
PUF	Physical Unclonable Function
QR	Quick Response
RFID	Radio-Frequency Identification
RPL	IPv6 Routing Protocol for Low-Power and Lossy Networks
RRR	Replay Rejection Rate
UDP	User Datagram Protocol
WSN	Wireless Sensor Network
6LoWPAN	IPv6 over Low-Power Wireless Personal Area Networks

References

1. Adam, M.; Hammoudeh, M.; Alrawashdeh, R.; Alsulaimy, B. A Survey on Security, Privacy, Trust, and Architectural Challenges in IoT Systems. *IEEE Access* **2024**, *12*, 57128–57149. https://doi.org/10.1109/ACCESS.2024.3382709.

2. Guşiţa, B.; Anton, A.A.; Stângaciu, C.S.; Stănescu, D.; Găină, L.I.; Micea, M.V. Securing IoT Edge: A Survey on Lightweight Cryptography, Anonymous Routing and Communication Protocol Enhancements. *Int. J. Inf. Secur.* **2025**, *24*, 149. https://doi.org/10.1007/s10207-025-01071-7.
3. Haider, Z.A.; Zeb, A.; Islam, A.K.M.M.; Rahman, T.; Arishi, A.; Ullah, I. Enhancing IoT Security with Resource-Efficient Cryptography: A Comprehensive Review of Lightweight and Hybrid Algorithms. *Comput. Sci. Rev.* **2026**, *59*, 100861. https://doi.org/10.1016/j.cosrev.2025.100861.
4. Winter, T.; Thubert, P.; Brandt, A.; Hui, J.; Kelsey, R.; Levis, P.; Pister, K.; Struik, R.; Vasseur, J.P.; Alexander, R. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. Technical Report RFC 6550, Internet Engineering Task Force, 2012. Available online: https://www.rfc-editor.org/info/rfc6550/ (accessed on 17 October 2026).
5. Oikonomou, G.; Duquennoy, S.; Elsts, A.; Eriksson, J.; Tanaka, Y.; Tsiftes, N. The Contiki-NG Open Source Operating System for Next Generation IoT Devices. *SoftwareX* **2022**, *18*, 101089. https://doi.org/10.1016/j.softx.2022.101089.
6. Österlind, F.; Dunkels, A.; Eriksson, J.; Finne, N.; Voigt, T. Cross-Level Sensor Network Simulation with COOJA. In Proceedings of the 2006 31st IEEE Conference on Local Computer Networks, Tampa, FL, USA, 14–16 November 2006; pp. 641–648. https://doi.org/10.1109/LCN.2006.322172.
7. Silva, M.; Cerdeira, D.; Pinto, S.; Gomes, T. Operating Systems for Internet of Things Low-End Devices: Analysis and Benchmarking. *IEEE Internet Things J.* **2019**, *6*, 10375–10383. https://doi.org/10.1109/JIOT.2019.2939008.
8. Albinali, H.; Azzedin, F. Replay Attacks in RPL-Based Internet of Things: Comparative and Empirical Study. *Comput. Netw.* **2025**, *257*, 110996. https://doi.org/10.1016/j.comnet.2024.110996.
9. Hkiri, A.; Alqurashi, S.; Ben Bahri, O.; Karmani, M.; Faraj, H.; Machhout, M. Performance Evaluation of Mobile RPL-Based IoT Networks under Hello Flood Attack. *Electronics* **2024**, *13*, 2226. https://doi.org/10.3390/electronics13112226.
10. Mokrani, S.; Belkadi, M.; Sadoun, T.; Lloret, J.; Aoudjit, R. LEA-RPL: Lightweight Energy-Aware RPL Protocol for Internet of Things Based on Particle Swarm Optimization. *Telecommun. Syst.* **2025**, *88*, 14. https://doi.org/10.1007/s11235-024-01254-y.
11. Irawan, A.I.; Santoso, I.H.; Istikmal.; Rahayu, M. Implementation of QR Code Attendance Security System Using RSA and Hash Algorithms. *J. Nas. Tek. Elektro Dan Teknol. Inf.* **2024**, *13*, 53–59. https://doi.org/10.22146/jnteti.v13i1.4395.
12. Annurrahma, K.D.; Suryawan, S.H.; Rahim, A. Design of a Meeting Attendance System Based on Dynamic QR Code with Universally Unique Identifier (UUID v4). *J. Ris. Inform.* **2025**, *7*, 336–346. https://doi.org/10.34288/jri.v7i4.416.
13. Oyerinde, E.I.; Idepefo, F.; Odo, F.; Onamade, A.A.; Adeoti, B.; Adediran, O.; Adelowo, O.; Laminga, M.; Isaac, M.; Akinyosoye, O.O. Design and Implementation of a Secure QR Code-Based Attendance Management System for Higher Education. *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol.* **2025**, *11*, 77–89. https://doi.org/10.32628/CSEIT25111699.
14. Siregar, R.; Maulana, B.; Isnaini, M.; Sabrina, E.; Hutahaean, H.D. Performance Evaluation of a Mobile Attendance System Using Dual-Factor Dynamic QR Code and GPS Geofencing. *J. Multimed. Dan Teknol. Inf. (Jatilima)* **2026**, *7*, 1126–1134. https://doi.org/10.54209/jatilima.v7i06.2281.
15. Prajapati, A.K.; Pilli, E.S.; Battula, R.B.; Varadharajan, V.; Verma, A.; Joshi, R.C. A Comprehensive Survey on RPL Routing-Based Attacks, Defences and Future Directions in Internet of Things. *Comput. Electr. Eng.* **2025**, *123*, 110071. https://doi.org/10.1016/j.compeleceng.2025.110071.
16. Panda, N.; Muthuraman, S. A Multiclass Machine Learning Framework for Detecting Routing Attacks in RPL-Based IoT Networks Using a Novel Simulation-Driven Dataset. *Future Internet* **2026**, *18*, 35. https://doi.org/10.3390/fi18010035.
17. Turan, M.S.; McKay, K.A.; Chang, D.; Kang, J.; Kelsey, J. *Ascon-Based Lightweight Cryptography Standards for Constrained Devices: Authenticated Encryption, Hash, and Extendable Output Functions*; Technical Report NIST SP 800-232; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2025. https://doi.org/10.6028/NIST.SP.800-232.
18. Fedrecheski, G.; Selander, G.; Watteyne, T.; Vučinić, M. ELA: Secure, Lightweight, and Zero-Touch Enrollment for IoT Devices. *Comput. Netw.* **2026**, *282*, 112258. https://doi.org/10.1016/j.comnet.2026.112258.
19. Lim, H.; Park, Y.; Park, Y. A PUF-Based Secure and Lightweight Authentication Protocol for Medical IoT Environments. *Sensors* **2026**, *26*, 3223. https://doi.org/10.3390/s26103223.
20. Arshad, D.; Asim, M.; Tariq, N.; Baker, T.; Tawfik, H.; Al-Jumeily, D. THC-RPL: A lightweight Trust-enabled routing in RPL-based IoT networks against Sybil attack. *PLoS ONE* **2022**, *17*, e0271277. https://doi.org/10.1371/journal.pone.0271277.
21. Dhifallah, W.; Tarhouni, M.; Moulahi, T.; Zidi, S. IoT-RPL 2021: Cyber Attack Dataset Based on RPL Routing for IoT. Mendeley Data, Dataset, Version 1, 2024. Available online: https://data.mendeley.com/datasets/4rcbbry2sc/1 (accessed on 17 October 2026).
22. Ayoub, K.; Abdelmajid, H.; Ayoub, T.; Soukaina, M. Enhancing IoT RPL Protocol Security Against Black Hole Attacks with Deep Learning Techniques. *J. Comput. Sci.* **2024**, *20*, 1530–1544. https://doi.org/10.3844/jcssp.2024.1530.1544.
23. Etheridge, A.; Anu, V. RPL Attack Detection in IoT Environments: An Ensemble Approach. In *Proceedings of IEMTRONICS 2024—International IoT, Electronics and Mechatronics Conference*; Lecture Notes in Electrical Engineering: Springer: Singapore, 2025; Volume 1228, pp. 113–122. https://doi.org/10.1007/978-981-97-4784-9_8.

24. Molose, R.; Isong, B. A Survey of Multi-Layer IoT Security Using SDN, Blockchain, and Machine Learning. *Electronics* **2026**, *15*, 494. https://doi.org/10.3390/electronics15030494.
25. Mustafa, R.; Sarkar, N.I.; Mohaghegh, M.; Pervez, S.; Morados, R. A Secure and Energy-Efficient Cross-Layer Network Architecture for the Internet of Things. *Sensors* **2025**, *25*, 3457. https://doi.org/10.3390/s25113457.
26. Mustafa, R.; Sarkar, N.I.; Mohaghegh, M.; Pervez, S.; Vohra, O. Cross-Layer Analysis of Machine Learning Models for Secure and Energy-Efficient IoT Networks. *Sensors* **2025**, *25*, 3720. https://doi.org/10.3390/s25123720.
27. Garcia Ribera, E.; Martinez Alvarez, B.; Samuel, C.; Ioulianou, P.P.; Vassilakis, V.G. An Intrusion Detection System for RPL-Based IoT Networks. *Electronics* **2022**, *11*, 4041. https://doi.org/10.3390/electronics11234041.
28. Matos, T.; Martins, M.F.; Rocha, J.; Dinis, H.; Martins, M.; Lopes, S.F.; Afonso, J.A.; Gonçalves, L. Design and Implementation of a Low-Cost and Low-Power Wireless Sensor Network for Estuary Environmental Monitoring. *Electronics* **2026**, *15*, 3155. https://doi.org/10.3390/electronics15143155.
29. Gang, Y.; Park, H.; Park, Y. A Lightweight and Secure End-to-End Authentication Protocol Using PUF for Internet of Drones. *Electronics* **2026**, *15*, 2535. https://doi.org/10.3390/electronics15122535.
30. Panimozhi, K.; Murthy, S.K. A Practical QR Code-based Attendance System with Proxy Mitigation Measures for Academic Institutions. In *Proceedings of the 2026 4th International Conference on Sustainable Computing and Smart Systems (ICSCSS)*; IEEE: Piscataway, NJ, USA, 2026; pp. 1976–1983. https://doi.org/10.1109/ICSCSS69635.2026.11645884.
31. Alabi, A.; Lateef, H.; Tewogbade, A.; Ozoh, P.; Balogun, F. Design and Implementation of an Enhanced QR-Code Based Attendance System. *J. Comput. Soc. Inform.* **2026**, *5*, 30–49. https://doi.org/10.33736/jcsi.10752.2026.
32. Abirami, R.K.; Rohann, S. Optimized QR Code-Based Authentication System for Attendance Management. *Acta Pedagog. Asiana* **2026**, *5*, 140–150. https://doi.org/10.53623/apga.v5iSI.1124.
33. Ma, H.; Tao, Y.; Fang, Y.; Chen, P.; Li, Y. Multi-Carrier Initial-Condition-Index-Aided DCSK Scheme: An Efficient Solution for Multipath Fading Channel. *IEEE Trans. Veh. Technol.* **2025**, *74*, 15743–15757. https://doi.org/10.1109/TVT.2025.3567450.
34. Gong, Y.; Zhang, Z.; Wang, K.; Gu, Y.; Wu, Y. IoT-Oriented Single-Transmitter Multiple-Receiver Wireless Charging Systems Using Hybrid Multi-Frequency Pulse Modulation. *IEEE Trans. Magn.* **2024**, *60*, 8401606. https://doi.org/10.1109/TMAG.2024.3417626.
35. Whiting, D.; Housley, R.; Ferguson, N. Counter with CBC-MAC (CCM). Technical Report RFC 3610, Internet Engineering Task Force, 2003. Available online: https://www.rfc-editor.org/info/rfc3610/ (accessed on 17 October 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.