

CAMPUS OF THINGS: A LORA-BASED IOT ARCHITECTURE FOR A SMART CAMPUS

Bernardo Milheiro

Department of Electrical and Computer Engineering, Polytechnic Institute of Setúbal, Setúbal, Portugal

ABSTRACT

This paper presents the design, implementation, and empirical evaluation of a LoRa-based IoT network deployed at the Setúbal campus of the Polytechnic Institute of Setúbal. The network leverages LoRa long-range communication integrated with open-source cloud infrastructure to create a research and development platform for IoT technologies while transforming the campus into a smart environment. A custom field Tester device was developed incorporating an RFM95 transceiver and integrated with a mobile application utilizing smartphone GPS for systematic coverage mapping. Extensive field testing across 33 locations demonstrated reliable coverage across all tested campus locations, with a maximum recorded communication distance of 17.5 km in one of the test sites. Path loss analysis revealed environmental coefficients ranging from 2.7 to 3.5 depending on terrain characteristics. The implementation validates the feasibility of cost-effective, scalable IoT infrastructure for educational institutions.

KEYWORDS

IoT, Smart Campus, LoRa, RSSI Mapping, Path Loss

1. INTRODUCTION

The convergence of ubiquitous connectivity, miniaturized sensing technologies, and cloud computing has catalyzed the rapid evolution of the Internet of Things (IoT) as a transformative paradigm in modern infrastructure management. Higher education institutions, characterized by their complex operational requirements and diverse stakeholder needs, represent ideal environments for IoT implementation, serving both as operational infrastructure and as living laboratories for technological innovation [1]. The Polytechnic Institute of Setúbal (IPS) identified the strategic imperative to establish comprehensive IoT infrastructure serving multiple complementary objectives beyond immediate operational efficiency. The "Campus das Coisas" (Campus of Things) project was conceived to create a platform for IoT research and education while enabling smart campus applications and extending services to surrounding communities through a LoRa-based wireless sensor network. The selection of LoRa technology reflects the specific requirements of campus-wide deployment. Unlike traditional wireless technologies that trade range for bandwidth, LoRa's Chirp Spread Spectrum (CSS) modulation enables long-range communication with minimal power consumption, making it particularly suitable for distributed sensor networks spanning large geographical areas. Operation in the unlicensed ISM band eliminates spectrum costs while strong security and high endpoint capacity align with anticipated scale and application diversity [2].

This implementation addresses critical challenges in educational IoT deployment, particularly reconciling campus network security with cloud-based platform integration. The architecture

employs virtual network segmentation and cloud-hosted servers to enable secure remote access while maintaining institutional cybersecurity posture. Adoption of open-source platforms ensures sustainability under budgetary constraints common to educational implementations. The research makes several contributions: demonstrating practical implementation methodology balancing security and accessibility; providing empirical coverage analysis across varied terrain; introducing an integrated mobile application framework leveraging smartphone GPS for automated Radio Frequency (RF) characterization; and establishing operational infrastructure serving both campus and surrounding communities.

2. RELATED WORK AND TECHNOLOGY BACKGROUND

This section surveys existing smart campus deployments to contextualize the Campus das Coisas project and provides a technical background on LoRa and LoRaWAN as the foundational communication technologies adopted in this work.

2.1. Smart Campus Implementations

Smart campus initiatives represent the application of IoT technologies to enhance educational environments through integrated sensing, data analytics, and intelligent automation. Successful implementations typically address environmental monitoring, energy management, space utilization, and security while serving as experimental platforms for student research [3].

The Tellus Space project exemplifies comprehensive smart campus integration, implementing distributed environmental sensors throughout learning spaces coupled with interactive digital objects and mobile-controlled intelligent lighting systems. The implementation demonstrated measurable improvements in energy efficiency while providing students hands-on experience with IoT system development and deployment methodologies [4]. The University of Coruña deployed targeted IoT infrastructure for monitoring CO₂ concentrations in examination halls, addressing the specific challenge of maintaining optimal air quality during high-occupancy assessment periods. Their analysis revealed significant CO₂ elevation during examinations, providing data-driven justification for enhanced ventilation strategies and demonstrating the value of application-specific IoT deployments in educational settings [5].

These case studies validate the technical feasibility and educational value of campus IoT implementations. However, they also reveal the importance of carefully planned network infrastructure capable of supporting diverse applications while maintaining accessibility for research and educational activities. The Campus das Coisas implementation builds upon these precedents while addressing the specific challenge of extended coverage beyond institutional boundaries to serve broader community needs.

2.2. Lora Technology

LoRa is a proprietary physical layer modulation technology developed by Semtech that employs CSS modulation [6]. Unlike narrowband modulation schemes such as Frequency-Shift Keying (FSK), CSS transmits signals whose frequency varies linearly over time, generating characteristic chirp waveforms that provide improved robustness to interference, multipath fading, and Doppler effects [7]. This robustness enables reliable communication under negative Signal-To-Noise Ratio (SNR) conditions, allowing reliable communication at very low SNR values and supporting the long-range and low-power operational characteristics that distinguish LoRa from alternative wireless technologies.

The technology operates in regional Industrial, Scientific and Medical (ISM) frequency bands, notably 868 MHz in Europe, where maximum transmit power is typically limited to 14 dBm Equivalent Isotropically Radiated Power (EIRP) under duty-cycle regulations. Depending on the spreading factor and bandwidth configuration, LoRa gateway receivers can achieve sensitivities down to approximately -142 dBm under SF12 and 125 kHz configuration. Under favourable rural line-of-sight conditions, communication ranges of up to 20 km are achievable. These characteristics make LoRa appropriate for wide-area sensing applications requiring minimal infrastructure deployment.

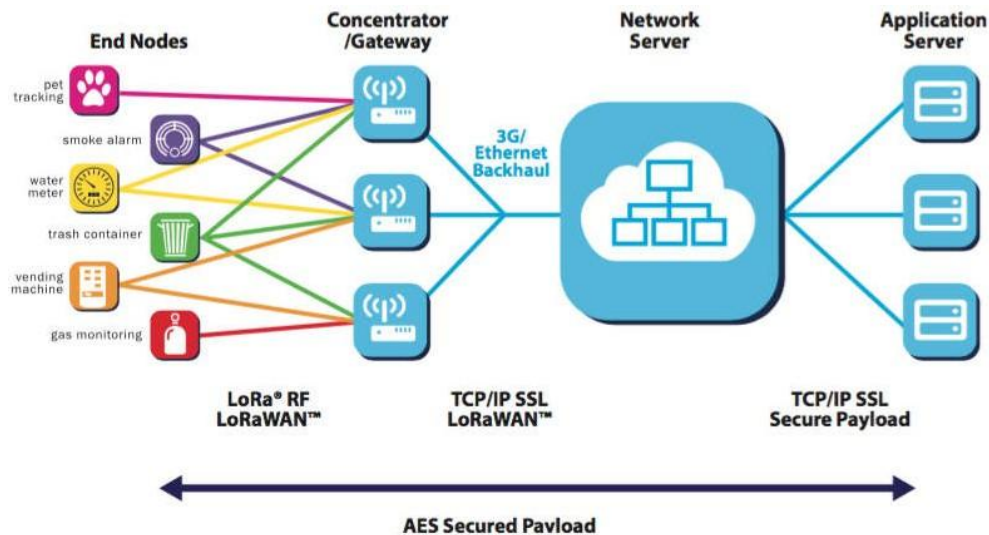


Figure 1: LoRa Network Architecture [8].

The LoRaWAN network protocol, standardized by the LoRa Alliance, builds upon the LoRa physical layer to provide a complete communication stack optimized for battery-powered sensor networks. The protocol defines three device classes (Class A, B, and C), offering different trade-offs between power consumption and downlink latency. Class A devices, employed in this implementation, achieve minimal power consumption by opening short receive windows only after each uplink transmission. Class B devices introduce scheduled additional receive windows synchronized by network beacons, reducing downlink latency at the expense of increased power usage. Class C devices maintain nearly continuous reception, enabling low-latency downlink communication but with significantly higher power consumption [9].

LoRa modulation is characterized by several configurable parameters that significantly influence communication performance. The spreading factor (SF), ranging from 7 to 12, defines the number of chips per symbol, with higher values resulting in lower data rates but increased receiver sensitivity, extended communication range, and improved robustness to interference at the cost of longer airtime. Bandwidth selection (125 kHz, 250 kHz, or 500 kHz) affects both data rate and receiver sensitivity. The coding rate determines the level of forward error correction overhead, improving resilience to interference and packet loss at the expense of reduced effective throughput. These parameters can be dynamically optimized through the Adaptive Data Rate (ADR) mechanism, which adjusts data rate and transmission power according to link conditions. However, ADR was deliberately disabled in this implementation to ensure consistent transmission settings across all coverage testing locations [6].

3. SYSTEM DESIGN AND IMPLEMENTATION

This section describes the campus environment and network requirements that shaped the architectural decisions, followed by a detailed description of the four-tier architecture adopted for the implementation.

3.1. Campus And Network Requirements

The Setúbal campus of IPS covers approximately 0.258 km² and is located on gently sloping terrain, with an average elevation of about 10 meters above sea level. The campus comprises four primary academic buildings housing the Schools of Technology (ESTS), Education (ESE), Business Sciences (ESCE), and supporting facilities. Building heights are relatively uniform at approximately 15 to 20 meters, with ESTS building representing the tallest structure and offering advantageous positioning for gateway placement.

Topographical analysis revealed that complete campus coverage would require minimum communication range of 450 meters from a centrally positioned gateway. However, the project scope explicitly extended beyond institutional boundaries to encompass surrounding educational facilities, commercial enterprises, and industrial zones within the greater Setúbal metropolitan area, necessitating significantly greater communication range targeting 5--10~km radius coverage. The geographical context presents both opportunities and challenges for wireless propagation. The city of Setúbal and its surrounding areas, such as Arrábida, Palmela, and Azeitão, feature significant variations in altitude and a highly diverse geography. To the south and east of the city, the elevation is lower and the area is predominantly aquatic, while to the north and west, mountainous terrain and higher elevations can be found. This geographical asymmetry suggested that coverage patterns would deviate substantially from simple circular models, necessitating empirical field testing to characterize actual performance.

Network requirements extended beyond mere coverage to encompass capacity, security, accessibility, and sustainability considerations. The anticipated deployment scale, potentially encompassing hundreds of sensor nodes across diverse applications, demanded gateway infrastructure capable of managing thousands of devices simultaneously. The campus network security policy, appropriately restricting external access to institutional IT infrastructure, created the challenge of enabling remote IoT management while maintaining security posture. The requirement for community accessibility, allowing students and external researchers to deploy devices without administrative intervention, further complicated access control requirements. Finally, long-term sustainability demanded minimal operational costs and dependence on maintainable open-source platforms rather than proprietary vendor solutions.

3.2. Architectural Design

The implemented architecture (Fig. 2) was designed to address coverage, scalability, security, and accessibility requirements through a four-tier hierarchy comprising device, gateway, network, and application layers.

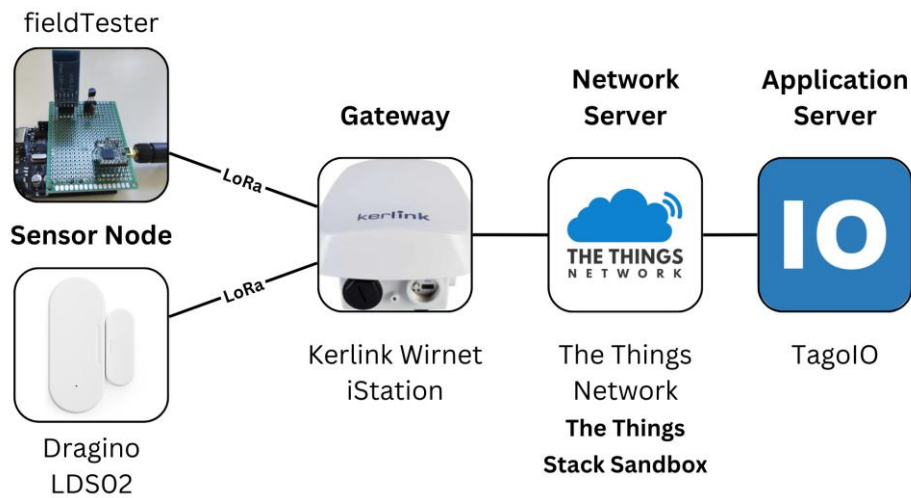


Figure 2: IoT Network Architecture for Setúbal Campus.

The device layer consists of sensor nodes implementing the LoRaWAN Class A protocol to ensure minimal power consumption. Two device types were deployed during the initial phase: a custom-developed fieldTester for coverage characterization and a commercial Dragino LDS02 door sensor used to validate integration of off-the-shelf LoRaWAN devices while providing operational building monitoring. Future expansions include environmental sensing nodes for temperature, humidity, CO₂ concentration, particulate matter, and occupancy detection.

The gateway layer functions as the bridge between the low-power wide-area network and IP-based infrastructure. A Kerlink Wirnet iStation eight-channel outdoor LoRaWAN gateway was selected due to its industrial-grade reliability, IP67 environmental protection, and multi-channel demodulation capability across multiple spreading factors. The gateway was installed at approximately 40 meters above ground level on the ESTS building, maximizing line-of-sight conditions across the campus and surrounding areas. Connectivity to the institutional network is provided via Gigabit Ethernet with Power over Ethernet (PoE), simplifying deployment and power provisioning.

The network layer is implemented through The Things Stack Sandbox (TTSS), a cloud-hosted LoRaWAN network server. Instead of deploying a local network server within the campus infrastructure, the gateway establishes secure outbound HTTPS connections to TTSS. This approach ensures remote accessibility while preserving institutional cybersecurity constraints. The gateway operates within a dedicated Virtual Local Area Network (VLAN) configured for IoT traffic, with firewall rules allowing only outbound communication to predefined TTSS endpoints. This configuration creates a controlled, unidirectional data flow that prevents external access to internal campus systems while enabling full remote network management.

The application layer employs TagoIO as the primary platform for data visualization, storage, and device management. Integration between TTSS and TagoIO is achieved through webhook mechanisms that forward decoded uplink payloads to the application server. TagoIO provides dashboard customization, automated analysis scripts, and free-tier storage capacity of up to 800,000 data points. This separation of concerns ensures that LoRaWAN protocol management remains within TTSS, while application-level logic and user interaction are handled independently by TagoIO.

4. IOT ARCHITECTURE IMPLEMENTATION

This section details the hardware and software components developed and configured for the system, covering the custom fieldTester device, the mobile application, and the cloud platform integration.

4.1. Fieldtester Device Development

The fieldTester (Fig. 3) represents a custom-developed coverage testing platform designed to characterize LoRa communication performance across diverse environments and distances. The device integrates several functional subsystems: an Arduino microcontroller providing computational resources and peripheral interfacing; an RFM95 LoRa transceiver module implementing the LoRa physical layer; a DS18B20 digital temperature sensor providing auxiliary data for testing; and an HC-06 Bluetooth module enabling wireless data transfer to the companion mobile application.

The RFM95 transceiver was selected for its complete implementation of LoRa modulation compatible with LoRaWAN networks, wide availability in the maker community facilitating future replication, and cost-effectiveness at approximately €11 per module. The module operates in the 868 MHz European ISM band with configurable output power up to 20 dBm and receiver sensitivity reaching -148 dBm at SF12, though practical sensitivity is constrained by noise floor considerations to approximately -136 dBm under SF9 configuration employed for testing. Device firmware implements LoRaWAN Class A communication using Activation by Personalization (ABP) rather than Over-the-Air Activation (OTAA) to simplify field deployment and eliminate dependence on network availability during initial power-up.

Although ABP simplifies field deployment during experimental campaigns, OTAA would be preferable in long-term production deployments due to its improved security and session key management mechanisms. The device transmits uplink messages at 30-second intervals, each containing GPS-derived latitude and longitude coordinates obtained from the companion mobile application along with current temperature readings. The payload encoding employs compact representation with coordinates encoded as 32-bit floating-point values and temperature as a 16-bit integer, yielding total payload size of 12 bytes well within LoRaWAN constraints.




Figure 3: fieldTester and Dragino LDS02 devices.

A critical aspect of the fieldTester design addresses the challenge of measuring received signal strength at the sensor node rather than the gateway. LoRaWAN specifications define RSSI (Received Signal Strength Indicator) and SNR (Signal-to-Noise Ratio) as metrics characterizing downlink quality from the perspective of the end device.

To obtain these measurements, the device must receive downlink transmissions from the gateway. The implementation triggers downlink transmission through user interaction with the mobile application, which invokes a cloud-hosted script on TagoIO that in turn posts a downlink message to TTSS. Upon receiving the next uplink from the fieldTester, the network server appends the queued downlink payload, enabling the device to measure and report RSSI values.

The firmware processes received downlink messages to extract RSSI values using library functions provided by the LMIC (LoRa MAC in C) implementation. These values are transmitted via Bluetooth to the mobile application in ASCII format for logging alongside GPS coordinates. The device incorporates visual feedback through an LED indicator illuminating upon successful downlink reception, providing field operators immediate confirmation of successful measurement cycles.

Power supply employs a standard 9V battery providing sufficient capacity for several hours of continuous operation during field testing campaigns. While inappropriate for long-term deployment due to limited battery life, this power source provides adequate portability and simplicity for coverage characterization purposes. Production sensor deployments would employ optimized power management with efficient DC-DC conversion and deep sleep modes enabling multi-year operation from primary lithium cells.

4.2. Mobile Application For Coverage Mapping

A critical innovation in the coverage testing methodology involves the development of a mobile application that automates the association of RSSI measurements with geographical coordinates. Previous approaches to coverage mapping typically required manual logging of GPS coordinates and signal measurements, introducing opportunities for transcription errors and substantially increasing field testing effort.

The developed application eliminates these limitations by integrating real-time GPS acquisition, Bluetooth communication, and automated data logging in a unified interface. The application was implemented using MIT App Inventor, a visual programming environment enabling rapid development of Android applications through drag-and-drop interface design and block-based logic programming. While limited in comparison to native development environments, App Inventor provides sufficient capability for the relatively straightforward requirements of coverage testing while dramatically reducing development effort and enabling future modification by users without specialized mobile development expertise.

The application architecture centres on three parallel processes: GPS coordinate acquisition, Bluetooth data reception, and downlink command transmission. The GPS subsystem continuously monitors the device's LocationChanged event, updating displayed latitude and longitude values whenever the smartphone's position changes. This approach leverages the smartphone's integrated GPS receiver, eliminating the need for dedicated GPS hardware in the fieldTester device and reducing both cost and complexity. Modern smartphones typically achieve GPS accuracy of 3-10 meters under clear-sky conditions, more than adequate for coverage mapping purposes where measurement locations are discrete points separated by hundreds of meters.

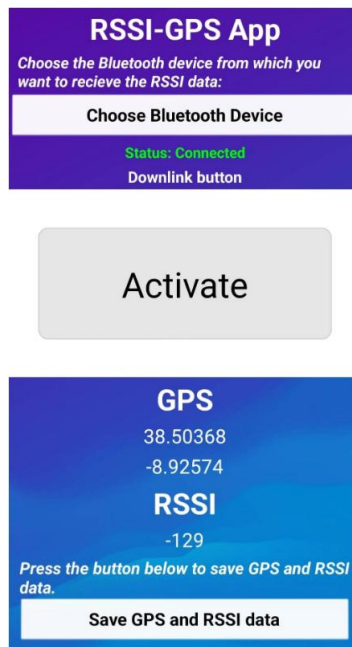


Figure 4: Mobile App.

The Bluetooth subsystem establishes connection to the fieldTester device upon user selection from a list of paired devices. Once connected, the application monitors the Bluetooth data stream, parsing incoming ASCII-encoded RSSI values transmitted by the fieldTester at 20-second intervals corresponding to successful downlink reception. Received RSSI values are displayed in the application interface, providing immediate feedback to field operators regarding signal quality at the current location.

The downlink trigger subsystem integrates the TagoIO webhook mechanism enabling remote command transmission without requiring direct platform access. During TagoIO dashboard configuration, a trigger button was configured with an associated webhook URL. This URL endpoint, when invoked via HTTP POST request, executes a server-side script that posts a downlink message to the TTSS network server. The mobile application embeds this webhook invocation capability through a WebView component configured to invoke the URL when the user presses the downlink button. This architecture enables field operators to trigger RSSI measurements through a single button press rather than requiring mobile browser access to the TagoIO dashboard. Data logging functionality stores each RSSI measurement alongside its associated GPS coordinates in CSV format within the smartphone's local storage. Each log entry comprises timestamp, latitude, longitude, and RSSI value, formatted as comma-separated values compatible with spreadsheet applications and geospatial analysis tools. The resulting data file can be exported via USB connection or cloud synchronization for subsequent processing and visualization.

4.3. Network Server And Application Configuration

The Things Stack Sandbox deployment required creation of an application container named "IPS-LoRa" to manage campus devices. Within this application, gateway integration began with registration of the Kerlink Wirnet iStation using its globally unique EUI-64 identifier and appropriate frequency plan configuration for the European 868 MHz band. Gateway registration generated API keys enable secure authenticated communication between the physical gateway and network server.

Device integration followed distinct processes for the fieldTester and Dragino LDS02 sensors. The commercial Dragino device benefited from pre-existing device profile definitions within TTSS specifying appropriate MAC layer parameters and payload decoding functions. Integration

required only EUI registration and selection of the appropriate device profile. In contrast, the custom fieldTester necessitated manual definition of device parameters including DevAddr, NwkSKey, and AppSKey for ABP activation, along with development of custom payload formatter functions.

Payload formatters were implemented in JavaScript, executing within the TTSS platform to decode raw byte arrays received from devices into structured JSON objects with named fields suitable for application consumption. The fieldTester formatter extracts latitude and longitude as IEEE 754 floating-point values and temperature as a signed 16-bit integer, presenting these as decimal degree and Celsius measurements respectively. This server-side decoding simplifies application development by providing pre-processed data rather than requiring each application to implement device-specific decoding logic. Webhook integration connects TTSS to TagoIO through HTTP POST requests transmitted for each uplink message, device status change, and downlink acknowledgment. The webhook payload comprises JSON-formatted data including device identifier, decoded payload fields, gateway metadata, and radio parameters. TagoIO parses this payload using its TTN/TTSS connector, automatically routing data to the corresponding device within the TagoIO platform based on device EUI matching.

Critical to coverage testing functionality is the configuration of ADR (Adaptive Data Rate) and frequency parameters. ADR was explicitly disabled to ensure consistent spreading factor across all test locations, eliminating this variable from coverage analysis.

Manual frequency selection restricted operation to three specific channels (868.1, 868.3, and 868.5 MHz) matching the gateway configuration, improving reliability by concentrating transmissions on actively monitored channels rather than distributing across the full eight-channel LoRaWAN specification.

4.4. Application Server Implementation

TagoIO was configured using the Custom TTI/TTN connector type, linking it to TTSS via webhooks that automatically route decoded uplink payloads to the corresponding device based on EUI matching. The platform creates time-series storage for each data field, enabling historical visualization through customizable dashboards displaying real-time sensor values and device status. A critical dashboard element is the downlink trigger button, which posts commands to TTSS when clicked, enabling RSSI measurement triggering from either the dashboard or mobile application. Under TagoIO's free tier, the 800,000 data point and 1,000,000 monthly API call limits provide adequate capacity for pilot deployments, guiding transmission interval and device density decisions for future production scaling.

5. COVERAGE TESTING AND ANALYSIS

This section presents the field-testing methodology employed to characterise network coverage, followed by the link budget analysis framework and the results obtained across campus and extended area locations.

5.1. Testing Methodology

Coverage characterization employed systematic field testing across 33 locations selected to represent diverse propagation environments and application scenarios. Testing locations were categorized into campus sites (15 locations within and immediately adjacent to institutional property) and extended area sites (18 locations distributed across Setúbal municipality and

surrounding communities). Campus sites included interior locations within buildings, exterior locations on campus grounds, and perimeter locations at campus access points. Extended area sites encompassed educational facilities, commercial areas, industrial zones, transportation hubs, and elevated terrain providing varying line-of-sight conditions.

At each location, the field-testing procedure followed a standardized protocol to ensure measurement consistency. The fieldTester device was positioned in an unobstructed location, typically at chest height (approximately 1.5 meters) to simulate typical sensor deployment scenarios. After powering the device, Bluetooth connection was established with the mobile application, and GPS coordinate acquisition was verified through the application display. The operator then triggered downlink transmission through the application interface and waited for the fieldTester to transmit its next uplink message, typically within 20 seconds. Upon downlink reception, the fieldTester calculated RSSI from the received packet and transmitted this value via Bluetooth to the mobile application, which automatically logged the measurement along with GPS coordinates and timestamp.

During the entire testing campaign, packet delivery reliability proved highly consistent. All successfully established communication links resulted in correct packet decoding, yielding an observed Packet Delivery Ratio (PDR) of 100% for all recorded measurements. Communication behaviour was therefore characterized by a binary outcome: packets were either received without error or not received at all. Multiple measurements were collected at locations exhibiting borderline signal quality or unusual propagation characteristics to verify repeatability.

The testing protocol specifically employed Spreading Factor 9 with 125 kHz bandwidth for all measurements, representing a balanced configuration between range and data rate commonly employed in LoRaWAN deployments. Transmit power was configured at 20 dBm (100 mW), the maximum permitted under European regulations. This consistent parameterization ensures that coverage maps accurately reflect expected performance for typical sensor deployments rather than optimized configurations applicable only to specialized scenarios.

5.2. Link Budget Analysis

The theoretical foundation for coverage analysis employs the link budget equation characterizing received signal power as a function of transmitter parameters, antenna gains, and propagation losses:

$$P_{Rx} = P_{Tx} + G_{Tx} + G_{Rx} - L_{path} - L_{misc} \quad (1)$$

where all values are in dB units, L_{path} represents distance-dependent path loss and L_{misc} encompasses cable and connector losses. For this implementation, miscellaneous losses were assumed negligible. Path loss estimation employed the long-distance path loss model:

$$L_{path} = 10n \log_{10}(d) + 20 \log_{10}(f) - 27.55 \quad (2)$$

where n is the path loss exponent (typically 2.0 for free space, 2.7-3.5 for urban environments), d is distance in meters, and f is frequency in MHz.

The path loss exponent n can be derived from a measured RSSI value by rearranging the link budget and path loss equations:

$$n = \frac{P_{TX} + G_{TX} + G_{RX} - P_{RX} - 20 \log_{10}(868) + 27.55}{10 \log_{10}(d)} \quad (3)$$

For this implementation, $P_{TX} = 20$ dBm, $G_{TX} = G_{RX} = 3$ dBi (antennas gain), and $f = 868$ MHz, giving a constant term of $20 \cdot \log_{10}(868) - 27.55 = 31.22$ dB. Applied to the Capela das Necessidades measurement ($d = 12,700$ m, $P_{RX} = -120$ dBm):

$$n = \frac{146 - 31.22}{41.04} = 2.797 \quad (4)$$

This aligns well with suburban environments. Similarly, for Covas de Coina ($d = 17,500$ m, $P_{RX} = -121$ dBm):

$$n = \frac{147 - 31.22}{42.43} = 2.729 \quad (5)$$

The lower path loss exponent suggests favourable propagation with limited obstruction. Theoretical maximum range is estimated when P_{RX} equals receiver sensitivity (-136 dBm at SF9) with $n = 2.797$:

$$d_{max} = 10^{(162 - 31.22)/(10 \times 2.797)} \approx 47.4 \text{ km} \quad (6)$$

This theoretical maximum range significantly exceeds achieved measurements, suggesting that factors beyond pure path loss limit practical range. These factors likely include multipath fading, interference from co-channel users, and gateway antenna limitations. Nevertheless, the close agreement between measured RSSI values and link budget predictions validates the propagation model and suggests that coverage in untested locations can be reliably estimated through computational modelling.

All 15 campus test locations demonstrated reliable connectivity with RSSI values ranging from -85 dBm to -110 dBm, providing 26 to 51 dB link margin above the receiver sensitivity threshold. This substantial margin ensures robust communication even in the presence of temporal fading, interference, or device component variations. The best signal quality was observed at the ESTS entrance parking area (100 meters, -85 dBm), benefiting from clear line-of-sight to the gateway without building penetration. The weakest on-campus signal occurred at the northwest campus entrance (424 meters, -110 dBm), attributable to greater distance combined with several intervening buildings.

Interior locations exhibited systematically higher path loss compared to exterior locations at similar distances, with the differential ranging from 8 to 15 dB. This building penetration loss aligns with theoretical expectations for reinforced concrete construction typical of modern educational buildings. The ESTS cafeteria interior (130 meters, -104 dBm) provides a representative example, showing 19 dB higher path loss than the ESTS entrance exterior location (67 meters, -92 dBm) despite only doubling the distance. This observation confirms that interior sensor deployments will experience reduced link margins and may require consideration of spreading factor optimization or strategic placement near windows to mitigate building attenuation.

The comprehensive campus coverage demonstrated by these measurements validates the gateway placement strategy and confirms that future sensor deployments can be positioned throughout campus facilities without coverage concerns. The substantial link margins enable operation at lower spreading factors (SF7 or SF8) for applications requiring higher data rates or reduced airtime or alternatively support increased deployment density without congestion concerns.

5.3. Extended Area Coverage Results

Extended area testing revealed a maximum communication range of 17.5 km at the Covas de Coina location, with an RSSI of -121 dBm and a link margin of 15 dB above the SF9 sensitivity threshold. This measurement was obtained under non-line-of-sight (NLoS) conditions: the path between the ESTS rooftop gateway and Covas de Coina crosses the Setúbal urban area and mixed terrain at low elevation, with no direct visual clearance between transmitter and receiver. Meteorological conditions during the test were stable, with clear skies, low humidity, and no precipitation, minimizing atmospheric absorption. The terrain profile at this southern location is relatively flat with low obstacle density, producing a path loss exponent of approximately 2.69, consistent with open suburban propagation, which accounts for the strong signal despite the absence of line-of-sight. This represents the greatest distance achieved during testing and suggests that further range may be possible under optimal conditions. The relatively strong signal at maximum range indicates that receiver sensitivity rather than fundamental path loss limitations determined this endpoint. Analysis of extended area results reveals several propagation phenomena characteristic of the local geography. Locations at similar distances but different elevations exhibited substantially different signal strengths, demonstrating the critical influence of terrain profile. The Palmela Castle location (7.15 km, 219 m elevation, -106 dBm) achieved 16 dB better signal than the Railway Station (4.20 km, 20 m elevation, -122 dBm) despite greater distance, attributed to elevated positioning providing clear line-of-sight over intervening obstacles.

The Serra da Arrábida mountain range creates a pronounced shadow region for locations immediately north of the ridgeline. Testing at several locations in this shadow zone (data not shown) revealed communication failures or extremely poor signal quality despite distances under 10 km. This phenomenon illustrates the fundamental limitation of line-of-sight propagation at UHF frequencies, where terrain obstacles create diffractive losses that can exceed 20-30 dB. Gateway deployments serving areas with significant terrain relief may require multiple gateway sites positioned to provide overlapping coverage from different directions.

Urban density effects are observable through comparison of locations at similar distances but different building densities. The Professional School of Setúbal (730 m, -92 dBm) in a residential area achieved superior signal compared to the Commercial Zone of Monte Belo (3.04 km, -120 dBm) despite the latter's greater altitude (44 m vs 16 m), attributable to higher building density and associated multipath effects in the commercial district. Path loss exponent calculations for near-campus locations reveal higher values (3.5-4.0) than distant suburban locations (2.7-2.9), reflecting the greater obstacle density and increased multipath propagation characteristic of dense urban environments. This spatial variation in path loss characteristics necessitates site-specific propagation modelling for precise coverage prediction rather than simple distance-based estimation.

6. DISCUSSION

The following subsections interpret the network performance results, reflect on implementation lessons, and identify current limitations alongside planned future enhancements.

6.1. Network Performance and Coverage Characteristics

The implemented network successfully achieved complete campus coverage while extending service availability substantially beyond institutional boundaries. The 17.5 km maximum communication range encompasses numerous educational institutions, commercial enterprises, and residential communities, enabling collaborative IoT projects and community engagement initiatives that transcend traditional campus boundaries. This extended reach validates LoRa technology selection and demonstrates that institutional IoT investments can generate broader societal benefit through strategic network design.

The comprehensive coverage mapping reveals distinct propagation regimes corresponding to different environmental characteristics. Within the dense campus core (distances under 500 m), path loss exponents exceed 3.0, reflecting multiple building penetrations and extensive multipath propagation. In the suburban ring (500 m to 5 km), path loss exponents decrease to 2.7-3.0, characteristic of mixed residential-commercial environments with moderate building density. Beyond 5 km in rural areas, path loss exponents approach free-space values near 2.0, limited primarily by terrain profile rather than man-made obstacles.

These propagation characteristics inform deployment strategies for different application scenarios. High-reliability applications requiring guaranteed connectivity should employ SF9 or higher to ensure adequate link margin across all propagation regimes. Applications tolerating occasional packet loss can employ SF7 or SF8 to increase data rate and reduce airtime, particularly for sensors in favourable propagation environments. The measured link margins support deployment of several hundred sensor nodes without capacity concerns, given typical IoT application duty cycles of 1% or less.

Contextualising these results against existing LoRa deployments in the literature reinforces the validity of the implementation. Petäjajarvi et al. [7] reported maximum ranges of 15–30 km in rural flat terrain using SF12, while Augustin et al. [2] characterised typical urban deployments achieving 2–5 km with path loss exponents between 2.7 and 3.5.

The Campus das Coisas deployment achieves a maximum range of 17.5 km under NLoS suburban conditions at SF9, a more demanding configuration than SF12, which compares favourably with the literature benchmarks. The measured path loss exponents (2.69–3.5 depending on environment) fall within the ranges reported for mixed urban-suburban deployments, confirming that the propagation model generalises beyond this specific campus context. The 100% PDR across all connected sites further aligns with findings from comparable single-gateway LoRa testbeds, where binary connectivity (full reception or no reception) is the dominant behaviour rather than partial packet loss.

6.2. Implementation Methodology And Lessons Learned

The integration of smartphone GPS capabilities through the mobile application proved critical to efficient coverage characterization. Previous methodologies requiring manual coordinate logging introduced substantial overhead and error potential. The automated approach reduced field testing time by approximately 50% while eliminating transcription errors, validating the investment in mobile application development. This methodology could be readily adapted to other wireless technologies and deployment scenarios requiring empirical coverage characterization.

The security architecture employing cloud-hosted network servers while maintaining campus network isolation represents a generalizable pattern applicable to other institutional IoT deployments facing similar constraints. Many educational and enterprise environments impose strict network access controls that complicate IoT integration. The demonstrated approach of

virtual network segmentation combined with outbound-only connectivity to cloud platforms provides a template for resolving this architectural challenge.

The total implementation cost of approximately €1,300 demonstrates the accessibility of LoRa infrastructure for resource-constrained educational institutions. This investment provides production-grade infrastructure capable of supporting hundreds of sensor nodes while serving both operational and educational purposes. The reliance on open-source platforms and community-supported services ensures long-term sustainability without recurring licensing fees, though institutions should recognize the need for technical expertise in system administration and network management.

6.3. Limitations And Future Enhancements

Several limitations of the current implementation suggest directions for enhancement. The single-gateway architecture, while adequate for current coverage requirements, creates a single point of failure and limits redundancy. Installation of secondary gateways at the ESE and ESCE buildings would provide backup connectivity and enable geolocation through time-difference-of-arrival calculations. The additional gateways would also improve indoor coverage in currently marginal areas.

The reliance on The Things Stack Sandbox (TTSS) introduces cloud dependency risks, as its community-tier service provides no SLA guarantees. To mitigate potential disruption, recommended measures include periodic local data buffering at the gateway during outages, regular automated export of device configurations to institutional storage, and evaluation of a self-hosted ChirpStack instance for production deployments where data sovereignty or uptime guarantees are required.

The coverage testing methodology, while systematic, sampled only 33 locations across the approximately 1000 km² metropolitan area encompassed by the signal's maximum range. This area contrasts with the institutional campus footprint of 0.258 km², with the extended coverage zone representing the broader Setúbal municipality and surrounding communities potentially served by the network. Interpolation between measurement points involves assumptions about propagation behaviour that may not reflect local terrain variations.

Enhanced coverage characterization could employ drive-testing methodologies with continuous RSSI logging to provide higher spatial resolution. However, such intensive testing must balance improved accuracy against increased resource requirements.

The current deployment includes only two sensor types, limiting demonstration of application diversity. Planned sensor additions include environmental monitoring stations measuring temperature, humidity, CO₂, and particulate matter; occupancy sensors for space utilization analysis; and equipment monitors for predictive maintenance. These additional sensors will validate the platform's capability to support heterogeneous applications while providing operational benefits through improved building management.

7. CONCLUSIONS

This paper presented the design, implementation, and empirical evaluation of a LoRa-based IoT network serving the Polytechnic Institute of Setúbal campus and surrounding communities. The Campus das Coisas project successfully established production-grade infrastructure supporting

research, education, and operational applications through strategic technology selection, thoughtful architectural design, and systematic empirical validation.

The implementation demonstrates that educational institutions can deploy sophisticated IoT infrastructure at modest cost through judicious use of open-source platforms and community-supported services. The measured performance characteristics validate LoRa technology for campus-scale deployments while providing empirical data on propagation behaviour in mixed urban-suburban environments. The innovative integration of smartphone GPS capabilities for coverage mapping represents a methodological contribution applicable to other wireless technology evaluations.

The extended coverage reaching 17.5 km enables collaborative applications with external educational institutions, government agencies, and commercial enterprises, demonstrating how institutional IoT investments can generate broader community benefit. The substantial link margins ensure robust operation and support future expansion to hundreds of sensor nodes across diverse applications. Future development will focus on sensor deployment across priority applications including environmental monitoring, energy management, and space utilization optimization. Additional gateway installations will improve redundancy and enable geolocation capabilities. The operational experience gained through these deployments will inform curriculum development and student projects, fulfilling the platform's educational mission while generating operational benefits for the institution.

REFERENCES

- [1] P. Gokhale, B. Omkar and B. Sagar, (2018) "Introduction to IOT", *International Advanced Research Journal in Science, Engineering and Technology*, Volume 5, Issue 1, pp. 41-44
- [2] Augustin, J. Yi, T. Clausen and W. M. Townsley, (2016) "A Study of LoRa: Long Range and Low Power Networks for the Internet of Things", *Sensors*, Volume 16, Issue 9, Article n° 1466.
- [3] Valks, M. H. Arkesteijn, A. Koutamanis and A. C. den Heijer, (2021) "Towards a smart campus: supporting campus decisions with Internet of Things applications", *Building Research and Information*, Volume 49, Issue 1, pp. 1-20.
- [4] E. Gilman, S. Tamminen, R. Yasmin, E. Ristimella, E. Peltonen, M. Harju, L. Lovén, J. Rieki and S. Pirttikangas, (2020) "Internet of Things for Smart Spaces: A University Campus Case Study", *Sensors*, Volume 20, Issue 13, Article n° 371.
- [5] T. Domínguez-Bolaño, V. Barral, C. J. Escudero and J. A. García-Naya, (2024) "An IoT system for a smart campus: Challenges and solutions illustrated over several real-world use cases", *Internet of Things*, Volume 25, Article n° 101099.
- [6] M. C. Bor, J. Vidler and U. Roedig, (2016) "LoRa for the Internet of Things", *European Conference on Wireless Sensor Networks*, Volume 16, pp. 361-366.
- [7] J. Petajarvi, K. Mikhaylov, M. Pettissalo, J. Janhunen and J. Iinatti, (2017) "Performance of a low-power wide-area network based on LoRa technology: Doppler robustness, scalability, and coverage", *International Journal of Distributed Sensor Networks*, Volume 13, pp. 1-16.
- [8] Evans, (2019) "LoRaWAN Network Server – Integrated or Cloud-hosted?", *RS DesignSpark*.
- [9] J. Haxhibeqiri, E. De Poorter, I. Moerman and J. Hoebeke, (2018) "A Survey of LoRaWAN for IoT: From Technology to Application", *Sensors*, Volume 18, Issue 11, Article n° 3995.

AUTHORS

Bernardo Milheiro received Master of Electrical and Computer Engineering from the Polytechnic Institute of Setúbal (Setúbal, Portugal). He also did Bachelor of Science in Electrical and Computer Engineering from the same institution. Currently, he is pursuing his PhD in Digitalisation Engineering from the Polytechnic Institute of Leiria (Leiria, Portugal) and concluding a Master of Information and Enterprise Systems from Universidade Aberta (Lisbon, Portugal). He is a Guest Lecturer in the School of Technology of Setúbal, Polytechnic Institute of Setúbal. His research interests include IoT, LoRa, Microcontrollers and AI.

