


HTWK

Hochschule für Technik,
Wirtschaft und Kultur Leipzig



Bachelor Thesis

Practical Performance Analysis of NB-IoT across Terrestrial and Non-Terrestrial Networks

Dominic Richter

2026

Abstract

The standardization of 3GPP Release 17 Non-Terrestrial Networks (NTN) enables standardized Narrowband IoT (NB-IoT) devices to communicate directly with satellite constellations, theoretically addressing the 85% global cellular coverage gap. However, treating these high-latency, power-constrained satellite links as direct equivalents to terrestrial networks introduces significant system inefficiencies. This thesis presents a comprehensive empirical end-to-end performance analysis comparing terrestrial NB-IoT with geostationary Earth orbit (GEO) NB-NTN using commercial off-the-shelf hardware. The research systematically evaluates protocol efficiency, power consumption, antenna polarization, and mobility impacts to establish validated performance baselines for hybrid IoT architectures.

Extensive field measurements reveal that the increased latency of GEO NTN significantly degrades connection-oriented protocols. For example, TCP generated 147% protocol overhead and experienced frequent timeouts, whereas the Constrained Application Protocol (CoAP) operating over UDP achieved reliable transmission with no packet loss observed in the evaluated scenarios.

Power profiling further demonstrates substantial energy overheads for satellite connectivity: while attaching to the NTN link consumes approximately the same energy as a suboptimal terrestrial attach, NTN message transmissions require more than 10 times the power of terrestrial NB-IoT transmissions. These characteristics necessitate optimization techniques such as payload compression or reduced messaging frequency.

Furthermore, environmental evaluations reveal that standard meteorological conditions, including rain and cloud cover, have no measurable impact on NTN performance. However, dynamic mobility campaigns expose a critical hardware trade-off in antenna architecture. While specialized circularly polarized antennas improve the link budget under line-of-sight conditions, they exhibit reduced reliability in reflection-heavy environments such as urban areas due to the physical phenomenon of polarization reversal. In such multipath scenarios, standard linearly polarized antennas demonstrate greater robustness and functional reliability.

Overall, this research concludes that GEO NB-NTN represents a viable connectivity option for globally deployable low-power telemetry systems. By leveraging connectionless application protocols, minimizing payload sizes, and utilizing context-aware antenna configurations, IoT device manufacturers can effectively harness standardized satellite connectivity for global data acquisition on battery-powered devices.

Table of Contents

Abstract	II
List of Acronyms	VI
1. Introduction	1
1.1 Motivation	1
1.2 Problem Statement	1
1.3 Scope and Objectives	2
1.4 Research Questions	4
2. NB-IoT over Satellite: Standards and Technical Constraints	5
2.1 NB-IoT and Release 17 NTN Adaptations	5
2.1.1 GEO Specifics: Propagation Delay and Path Loss	5
2.1.2 Timing Advance Pre-Compensation	6
2.1.3 HARQ Disabling	6
2.2 IoT Protocols in High-Latency Networks	6
2.2.1 Protocol Overhead: Payload vs. Total Transmission	6
2.2.2 TCP Handshake Overhead and Latency Degradation	7
2.2.3 UDP and CoAP: Efficient Implementations for Hybrid Networks	7
2.3 Power Consumption Profiles: TN vs. NTN	8
2.3.1 Energy Dynamics in NB-IoT	8
2.3.2 Impact of Path Loss and Repetitions on Transmit Power	8
2.3.3 Energy Penalty of Latency and Synchronization	9
2.4 Antenna Fundamentals: Linear vs. Circular Polarization	9
2.4.1 Principles of Polarization	9
2.4.2 Satellite Downlink Polarization Constraints	9
2.4.3 Polarization Mismatch Loss	10
2.4.4 Multipath Propagation and Signal Reflection	10
2.5 End-to-End Cellular IoT Architecture	11
2.5.1 Device Domain	11
2.5.2 Network Domain	12
2.5.3 Cloud Platform Domain	12
	III

3. System Design & Measurement Setup	13
3.1 Solution Architecture	13
3.1.1 Device Domain: Hardware and Firmware	13
3.1.2 Network Domain: Connectivity and Transport	14
3.1.3 Cloud Domain: IoT Platforms and Integration	14
3.2 Device Architecture	14
3.2.1 Hardware Components	14
3.2.2 Firmware Architecture and Concurrency Model	15
3.2.3 RAT Management and GNSS Data Synchronization	15
3.2.4 Protocol Abstraction and Campaign Orchestration	16
3.3 Payload Design and Data Modeling	18
3.3.1 Positional Schema and Serialization Efficiency	18
3.3.2 Dynamic Payload Padding and Compression Resistance	18
3.3.3 Cellular Connection Parameters and Signal Metrics	19
3.3.4 Spatial and Environmental Data Aggregation	19
3.3.5 Time Synchronization and Latency Profiling	19
3.4 Receiving Server Architecture	20
3.4.1 Multi-Protocol Ingestion Layer	20
3.4.2 Network Overhead and Latency Analysis	20
3.4.3 Asynchronous Data Enrichment	21
3.4.4 Storage Architecture and Real-Time Distribution	21
3.5 Measurements	21
3.5.1 Measurement Campaigns	22
3.5.2 Measurement Setups	24
4. Measurement Results & Analysis	28
4.1 Basic Analysis of NB-IoT and NB-NTN	28
4.1.1 NTN Cell Variability and Network Allocation	28
4.1.3 Impact of Meteorological Conditions on NTN Performance	32
4.1.4 Impact of Payload Size on Transmission Latency	34
4.1.5 Network Reliability and Message Drop Rates	36
4.2 Transmission Protocol Comparison	37
4.2.1 Protocol Overhead Analysis	37

4.2.2 Network Reliability and Drop Rates	38
4.3 Power Consumption Measurements	40
4.3.1 Network Attach Energy Profile	40
4.3.2 Active Transmission Power and Payload Scaling	41
4.3.3 Average Current Draw and Extrapolation	42
4.3.4 Summary of Energy Expenditure Multipliers	43
4.4 Antenna Performance	44
4.5 Mobility & Automotive Performance	47
4.5.1 Impact of Velocity	47
4.5.2 Antenna Polarization and Environmental Attenuation	47
4.5.3 Data Analysis: Polarization Reversal Phenomenon	48
5. Discussion	50
5.1 Interpretation of Results and Implications for IoT Device Manufacturers	50
5.1.1 Redefining the Application Stack	50
5.1.2 Payload Sizing	51
5.1.3 Hardware Architecture and Environmental Adaptability	51
5.1.4 Optimizing the Power Budget	52
5.2 Operational Boundaries of GEO-Based NB-NTN	52
5.2.1 Latency and Half-Duplex Characteristics	52
5.2.2 Return-Link Energy Dynamics	53
5.2.3 Urban Shadowing	53
5.2.4 Cloud-Side Integration and Protocol Translation	53
6. Conclusion and Outlook	54
6.1 Conclusion	54
6.2 Outlook	55
Bibliography	VII
Appendix	IX

List of Acronyms

3GPP	3rd Generation Partnership Program	PSM	Power Saving Mode
ACK	Acknowledgment	RAN	Radio Access Network
API	Application Programming Interface	RAT	Radio Access Technology
CE	Coverage Enhancement	RF	Radio Frequency
CoAP	Constrained Application Protocol	RHCP	Right-Hand Circular Polarization
CP	Circularly Polarized	RLC	Radio Link Control
E-UTRA	Evolved Universal Terrestrial Radio Access	RQ	Research Question
E-UTRAN	Evolved Universal Terrestrial Radio Access Network	RRC	Radio Resource Control
E2E	End-to-End	RSRP	Reference Signal Received Power
eDRX	extended Discontinuous Reception	RSRQ	Reference Signal Received Quality
EIRP	Equivalent Isotropically Radiated Power	RSSI	Received Signal Strength Indicator
eMTC	enhanced Machine Type Communication	RTOS	Real-Time Operating System
eNodeB	Evolved Node B	RTT	Round-Trip Time
EPC	Evolved Packet Core	SDK	Software Development Kit
FOTA	Firmware Over-The-Air	SIM	Subscriber Identity Module
FSPL	Free Space Path Loss	SINR	Signal-to-Interference-plus-Noise Ratio
GEO	Geostationary Earth Orbit	SiP	System-in-Package
GNSS	Global Navigation Satellite System	TA	Timing Advance
GSMA	Global System for Mobile Communications Association	TAC	Tracking Area Code
HARQ	Hybrid Automatic Repeat Request	TCP	Transmission Control Protocol
HTTP	HyperText Transfer Protocol	TN	Terrestrial Networks
IoT	Internet of Things	UDP	User Datagram Protocol
JSON	JavaScript Object Notation	UE	User Equipment
LEO	Low Earth Orbit	UTC	Coordinated Universal Time
LHCP	Left-Hand Circularly Polarized		
LOS	Line-of-Sight		
LP	Linearly Polarized		
LPWAN	Low-Power Wide-Area Network		
LwM2M	Lightweight M2M		
MAC	Medium Access Control		
MCL	Maximum Coupling Loss		
MCU	Microcontroller Unit		
MQTT	Message Queueing Telemetry Transport		
NAT	Network Address Translation		
NB-IoT	Narrowband IoT		
NB-NTN	NarrowBand Non-Terrestrial Networks		
NR	New Radio		
NTN	Non-Terrestrial Networks		
NTP	Network Time Protocol		
OS	Operating System		
PLF	Polarization Loss Factor		
PLMN	Public Land Mobile Network		
PPK2	Nordic Semiconductor Power Profiler Kit II		

1. Introduction

1.1 Motivation

The deployment of Internet of Things (IoT) devices is integral to data acquisition in industries such as logistics, agriculture, and environmental monitoring. However, the primary connectivity infrastructure for these devices is constrained by geographic and economic limitations. Currently, terrestrial cellular networks provide coverage to approximately 15% of the Earth's surface. The remaining 85% constitutes a coverage gap that restricts the continuous operation of IoT sensors in remote, rural, and maritime environments [1].

Historically, addressing this coverage gap required proprietary satellite communication systems. These legacy solutions typically involve high hardware costs, larger form factors, and elevated power consumption, which are incompatible with the constraints of battery-powered Low-Power Wide-Area Network (LPWAN) devices [1]. The standardization of Non-Terrestrial Networks (NTN) in the 3GPP (3rd Generation Partnership Project) Release 17 specifications in 2021 introduces a framework to address this limitation [2]. This release enables standard Narrowband IoT (NB-IoT) protocols to operate over satellite constellations. Consequently, standard cellular hardware can establish direct communication links with Geostationary Earth Orbit (GEO) and Low Earth Orbit (LEO) satellites, enabling IoT connectivity in remote areas without the need for proprietary radio modules.

While the 3GPP Rel-17 specifications establish compatibility at the hardware and modulation levels, the practical implementation of Narrowband IoT over Non-Terrestrial Networks (NB-NTN) introduces significant physical and protocol-level challenges. A prevalent approach in current IoT systems engineering is to treat NTN as a direct equivalent to Terrestrial Networks (TN). As a result, standard terrestrial application stacks, network protocols, and conventional hardware designs are copied in satellite environments without architectural adaptation.

This direct porting of terrestrial architecture fails to account for the physical constraints inherent to satellite communications. An NB-IoT module communicating with a GEO satellite at an altitude of approximately 35,786 kilometers experiences channel conditions fundamentally different from those of a terrestrial cell tower link. These conditions include severe free-space path loss, high signal propagation delays, and specific radio frequency phenomena such as polarization reversal [3]. When standard terrestrial configurations such as connection-oriented Transmission Control Protocol (TCP) layers or standard linearly polarized antennas are utilized under these constraints, overall system efficiency degrades. This degradation manifests as excessive protocol overhead, increased power consumption, and frequent connection timeouts at the edge of the link budget [4], [5].

To effectively utilize the capabilities introduced by 3GPP Rel-17, it is necessary to empirically evaluate the physical and protocol-level boundaries of NB-NTN. This thesis is motivated by the requirement to quantify the discrepancies between standard terrestrial IoT application design and the operational realities of high-latency, power-constrained satellite networks, thereby providing a data-driven foundation for future IoT implementations utilizing satellite links.

1.2 Problem Statement

While the 3GPP Rel-17 specifications provide the theoretical and regulatory framework for Non-Terrestrial Networks, the practical deployment of NB-NTN using commercial off-the-shelf hardware introduces

variables that are currently under-researched. The central problem addressed by this thesis is the critical deficit of empirical, real-world performance data comparing TN and NTN under identical hardware and software conditions.

Currently, the majority of available literature and technical documentation regarding NB-NTN performance relies heavily on software simulations, mathematical link budget models, or proprietary laboratory testbeds. There is a distinct lack of publicly available, quantitative measurements utilizing standard, commercially available cellular modules operating in true hybrid TN/NTN environments. This absence of empirical baselines creates significant uncertainty for systems engineers tasked with designing global IoT solutions. Preliminary industry requirement analyses, including surveys conducted among IoT developers and architects, reveal a widespread underestimation of the physical limitations imposed by satellite links, further highlighting the need for objective parameter validation.

Specifically, this uncertainty manifests across the critical domains of network protocol selection, antenna architecture, and the resulting overall power consumption profiles.

First, regarding network protocols, the transport and application layers heavily influence the energy efficiency and data viability of constrained devices. Although connection-oriented protocols are fundamentally unsuitable for low-power, constrained environments, TCP remains frequently utilized in terrestrial IoT applications. However, the impact of NTN's extreme latency and potential packet loss on TCP's three-way handshake, acknowledgment mechanisms, and keep-alive transmissions remains inadequately quantified in real-world deployments. The efficiency and protocol overhead of connectionless alternatives, such as the User Datagram Protocol (UDP) and the Constrained Application Protocol (CoAP), must be empirically validated under specific NTN constraints to determine the ratio of useful payload to total transmission size, and to directly compare the energy expenditure of TN versus NTN during network attach procedures and active message transmissions.

Second, the transition to multi-orbit connectivity introduces conflicting antenna design requirements. Standard terrestrial cellular devices predominantly utilize linearly polarized (LP) antennas due to their compact form factor and cost-effectiveness. In contrast, because satellite downlinks typically utilize circularly polarized (CP) signals, IoT devices require matching CP antennas to optimize the already tight link budget. It remains unclear to what extent standard LP antennas can maintain a reliable link budget for GEO satellite communication under real-world conditions and whether specialized CP hardware is strictly mandatory for hybrid IoT devices.

Third, the aggregate effect of these protocol inefficiencies and hardware constraints culminates in a highly unpredictable power consumption profile. The core value proposition of NB-IoT is an extended battery lifespan, typically achieved through extended sleep cycles and brief active transmission windows. However, the extreme Free Space Path Loss (FSPL) and latency of GEO NTN force the modem to operate at maximum transmission power for significantly longer durations, compounded by the energy cost of acquiring Global Navigation Satellite System (GNSS) timing synchronization. Without precise, empirical power profiling of these prolonged active states, it is impossible to accurately dimension the battery requirements for commercial satellite IoT deployments.

1.3 Scope and Objectives

To ensure that this research addresses the most pressing technical challenges within the current IoT landscape, the preliminary scope was validated through open consultations with industry experts across the telecommunications, satellite, and Radio Frequency (RF) engineering sectors [6]. Guided by this industry feedback, the primary objective of this thesis is to conduct a comprehensive, empirical performance evaluation of 3GPP Rel-17 NB-NTN utilizing commercial off-the-shelf hardware. By

transitioning from theoretical models to real-world field testing, this research aims to quantify the physical and protocol-level limitations of hybrid Terrestrial Network and Non-Terrestrial Network deployments. The resulting data is intended to provide systems engineers with validated parameters for optimizing power consumption, protocol selection, and hardware architecture in satellite-enabled IoT devices.

To achieve this objective, the scope of the thesis encompasses the design, implementation, and execution of systematic measurement campaigns. The investigation focuses specifically on the interaction between transport and application-layer protocols and the physical layer constraints of high-latency satellite links. Furthermore, the scope includes a comparative hardware analysis evaluating the viability of standard LP antennas against specialized circularly polarized antennas under various environmental conditions. Across all these evaluations, a comparative power consumption analysis is conducted to profile the energy expenditure of typical TN and NTN operations. By measuring the power draft during network attach procedures, message transmissions, and under varying connection parameters, the thesis quantitatively isolates the additional energy burden introduced by satellite connectivity.

To ensure experimental control and feasibility, specific boundaries are established for this research. The empirical measurements are delimited to GEO satellite networks, utilizing the Skylo network infrastructure as the NTN provider. While LEO constellations constitute a significant portion of the broader NTN ecosystem, empirical testing on LEO networks is excluded from this scope due to the current state of network maturity.

Additionally, the hardware scope is restricted to the Nordic Semiconductor nRF9151 module to provide a consistent baseline for all baseband and modem operations. This specific module was selected because Nordic Semiconductor is an industry leader in highly optimized, power-efficient modem architectures, and the nRF9151 is currently the only commercial module natively supporting LEO next to GEO and terrestrial NB-IoT connectivity. Furthermore, the Nordic hardware ecosystem includes the Power Profiler Kit II, which allows simple but high-resolution, empirical power measurements to answer the research objectives of this thesis.

The cellular connectivity analysis is exclusively limited to Narrowband IoT. Other LPWAN technologies, such as LTE-M, are not evaluated. By clearly defining these boundaries, the thesis ensures a rigorous, focused, and reproducible analysis of the core factors influencing NB-NTN performance in constrained IoT environments

1.4 Research Questions

To systematically address the problem statement and fulfill the defined objectives, this thesis is guided by four primary research questions (RQ). These questions are designed to isolate and quantify the physical, protocol-level, and energy-related discrepancies between terrestrial networks and non-terrestrial networks.

RQ1 - Protocol Efficiency and Reliability: How does the selection of transport and application-layer protocols impact transmission reliability, protocol overhead, and latency in high-path-loss GEO NB-NTN environments compared to terrestrial NB-IoT?

RQ2 - Environmental and Operational Impact: To what extent do environmental and operational factors degrade the link stability and transmission success rate of NB-NTN compared to equivalent terrestrial networks?

RQ3 - Hardware Architecture and Antenna Polarization: Can commercial off-the-shelf hardware, utilizing linearly polarized antennas and sub-optimal device placement, maintain a reliable GEO NTN link budget, or is the integration of specialized hardware such as circularly polarized antennas and optimized antenna placement mandatory for robust connectivity?

RQ4 - Energy Profiling: What is the quantitative difference in power consumption between terrestrial NB-IoT and satellite NB-NTN operations, specifically regarding the energy expenditure required for initial network attach procedures and subsequent payload transmissions?

Answering these questions will provide a comprehensive, data-driven framework for evaluating the commercial and technical viability of 3GPP Rel-17 NB-NTN for low-power IoT applications

2. NB-IoT over Satellite: Standards and Technical Constraints

2.1 NB-IoT and Release 17 NTN Adaptations

Narrowband IoT was initially introduced in 3GPP Release 13 as a terrestrial LPWAN standard. It was designed to provide extended coverage, deep indoor penetration, and extreme energy efficiency for constrained devices by operating within a narrow 180 kHz bandwidth. To achieve reliable communication in challenging RF environments, NB-IoT relies on a high Maximum Coupling Loss (MCL) of up to 164 dB, achieved primarily through extensive repetitions at the physical layer rather than complex modulation schemes [7].

While these mechanisms are highly effective for TN, porting NB-IoT to NTN introduces severe physical layer and Medium Access Control (MAC) layer challenges. 3GPP Release 17 standardizes the necessary protocol adaptations to allow NB-IoT to operate over satellite links, addressing the fundamental differences in signal propagation delay and attenuation [2].

2.1.1 GEO Specifics: Propagation Delay and Path Loss

The most significant physical distinction between TN and NTN lies in the deployment altitude of the base station. GEO satellites are positioned at an altitude of approximately 35,786 kilometers. This immense distance fundamentally alters the physical channel in two ways:

1. Absolute Latency:

In terrestrial cellular networks, the propagation delay between User Equipment (UE) and the Evolved Node B (eNodeB) is typically measured in microseconds. In a GEO NTN architecture, electromagnetic waves must travel a minimum of 71,572 kilometers for a single Round-Trip Time (RTT). Limited by the speed of light, the one-way propagation delay is approximately 119 to 135 milliseconds, depending on the satellite's elevation angle relative to the UE. Consequently, the minimum physical RTT exceeds 240 ms, and when combined with baseband processing, network routing, and core network latency, the operational RTT frequently exceeds 500 ms to 600 ms [8].

2. Free Space Path Loss:

The link budget in NTN is severely constrained by Free Space Path Loss, which scales logarithmically with distance d and frequency f , calculated as:

$$FSPL = 20\log_{10}(d) + 20\log_{10}(f) + 20\log_{10}\left(\frac{4\pi}{c}\right)$$

For an S-band or L-band satellite link (e.g., 2.0 GHz), the FSPL approaches 190 dB. This exceeds the Rel-13 terrestrial MCL limit of 164 dB. To close this link budget, GEO NTN relies on highly directional satellite spot beams with substantial Equivalent Isotropically Radiated Power (EIRP) and high antenna gain-to-noise-temperature (G/T) figures, coupled with maximum physical layer repetitions from the UE [2]

2.1.2 Timing Advance Pre-Compensation

In synchronous cellular networks like LTE and NB-IoT, the eNodeB dictates a Timing Advance (TA) value to the UE. This ensures that uplink transmissions from various devices arrive at the eNodeB perfectly aligned within their designated time slots, avoiding inter-symbol interference. In TN, the eNodeB measures the delay and sends dynamic TA update commands to the UE.

In a GEO NTN environment, the standard closed-loop TA mechanism fails. The 500 ms RTT renders real-time TA commands from the eNodeB obsolete before they reach the UE. To resolve this, 3GPP Rel-17 introduces GNSS-based TA pre-compensation. The UE is required to determine its own geolocation, for example via an integrated GNSS receiver and utilize broadcasted satellite ephemeris data which includes the satellite's exact orbital position and velocity vector. The UE mathematically calculates the propagation delay to the satellite and autonomously pre-compensates its uplink transmission timing before transmitting the Random Access Preamble (NPRACH) [9].

2.1.3 HARQ Disabling

Hybrid Automatic Repeat Request (HARQ) is a fundamental error-correction mechanism in terrestrial MAC layers. It relies on a rapid "Stop-and-Wait" protocol, where the UE transmits a transport block and waits for a rapid Acknowledgement (ACK/NACK) from the eNodeB before transmitting the next block or retransmitting.

In a GEO NTN, the 500 ms+ RTT fundamentally breaks the efficiency of HARQ. Waiting half a second for every acknowledgment stalls the MAC processes, drastically reduces data throughput, and forces the UE to keep its radio active for listening, causing massive energy consumption. Consequently, 3GPP Rel-17 specifies that HARQ feedback can be disabled for NTN deployments [10]. Instead of relying on MAC-layer HARQ, NTN relies on Radio Link Control Acknowledged Mode (RLC AM) for reliability, combined with blind physical layer repetitions which are pre-configured retransmissions without waiting for an ACK.

2.2 IoT Protocols in High-Latency Networks

The selection of transport and application-layer protocols dictates the signaling behavior, energy expenditure, and overall data efficiency of an IoT device. In high-latency, bandwidth-constrained environments such as GEO NB-NTN, protocol inefficiencies inherent to standard IT architectures are severely magnified. While industry documents, such as the GSMA (Global System for Mobile Communications Association) *IoT Device Connection Efficiency Guidelines* (TS.34) [11], provide best-practice frameworks for efficient implementation, the necessity for optimized protocols is fundamentally driven by the physical constraints of the satellite link and the energy limitations of the sensor itself [1].

2.2.1 Protocol Overhead: Payload vs. Total Transmission

To quantitatively evaluate protocol efficiency in constrained networks, it is necessary to differentiate between the application-layer data ("Payload") and the total data transmitted over the radio interface ("Total Transmission"). Every OSI layer appends header bytes to the payload, which consume limited radio resources and transmit power.

The Total Transmission size S_{total} in bytes can be defined as:

$$S_{total} = S_{payload} + S_{app} + S_{transport} + S_{network}$$

Where:

$S_{payload}$ represents the actual sensor data (e.g., a 10-byte GPS coordinate).

S_{app} , $S_{transport}$ and $S_{network}$ represent the headers for the application (e.g., Message Queueing Telemetry Transport (MQTT)/CoAP), transport (TCP/UDP), and network (IPv4/IPv6) layers, respectively.

Protocol efficiency η is calculated as the ratio of useful payload to the total transmission size:

$$\eta = \frac{S_{payload}}{S_{total}}$$

Academic studies on LPWAN energy consumption emphasize that maximizing η is critical for battery-operated devices [12]. Reducing S_{total} directly minimizes the duration the radio module remains in its high-power active state (Radio Resource Control (RRC) Connected) resulting in lower power consumption. Therefore, limiting protocol overhead is not merely an industry recommendation, but a physical requirement for viable NTN deployments.

2.2.2 TCP Handshake Overhead and Latency Degradation

TCP is a connection-oriented transport protocol designed to ensure reliable data streams in stable, low-latency terrestrial networks. To establish a connection, TCP requires a three-way handshake (SYN, SYN-ACK, ACK) before any application data can be transmitted.

In a terrestrial NB-IoT network, this handshake introduces a marginal delay. However, in a GEO NTN environment with a RTT of approximately 600 ms, the TCP handshake alone introduces a minimum delay of 1.2 to 1.8 seconds. This "handshake tax" forces the UE to transmit and receive empty signaling packets across a 71,572-kilometer link solely to open a socket.

Furthermore, research evaluating TCP performance over satellite links highlights that the protocol's reliance on dynamic RTT estimations and acknowledgment mechanisms leads to severe performance degradation [13]. The extreme and variable latency of GEO networks frequently causes TCP retransmission timers to expire prematurely, triggering spurious retransmissions. Additionally, to maintain stateful TCP connections through Network Address Translation (NAT) gateways, devices are often forced to transmit periodic keep-alive messages. As noted in GSMA TS.34, frequent network polling prevents the NB-IoT module from entering critical energy-saving states, such as Power Saving Mode (PSM) or extended Discontinuous Reception (eDRX), thereby rapidly depleting the battery [11].

2.2.3 UDP and CoAP: Efficient Implementations for Hybrid Networks

To optimize energy efficiency and mitigate the latency penalties of connection-oriented protocols, connectionless transport is broadly recommended for constrained IoT networks [4]. UDP eliminates the three-way handshake and the associated connection state maintenance. A UDP datagram utilizes an 8-byte header, drastically reducing $S_{transport}$ compared to TCP's 20-byte minimum header.

However, because UDP is stateless and provides no inherent guarantee of delivery, the application layer must assume responsibility for reliability which is a critical requirement in NTN environments where MAC-layer HARQ may be disabled.

The CoAP, standardized in RFC 7252, is specifically engineered to resolve this issue in lossy networks [14]. CoAP operates over UDP and implements essential reliability features without the connection-oriented overhead. Key architectural advantages include:

Confirmable (CON) Messages: CoAP allows the UE to request an application-layer acknowledgment (ACK) only when necessary, minimizing the volume of return traffic over the radio link.

Minimal Header Overhead: CoAP utilizes a compact, binary header of merely 4 bytes S_{app} , significantly improving the protocol efficiency η compared to text-based protocols like MQTT or HyperText Transfer Protocol (HTTP).

Controlled Retransmissions: Unlike TCP's aggressive window scaling, CoAP employs a strict, randomized exponential backoff algorithm for retransmissions. If a message fails over the satellite link, the device waits for a progressively longer duration before retrying. This built-in backoff mechanism aligns with good implementation practices to prevent network congestion and conserves device energy during temporary coverage losses.

Empirical studies comparing IoT application protocols consistently demonstrate that CoAP yields significantly lower overhead and faster transaction times in high-delay networks [4], making it structurally superior for satellite NB-IoT deployments.

2.3 Power Consumption Profiles: TN vs. NTN

The fundamental value proposition of Narrowband IoT is its extreme energy efficiency, theoretically enabling battery lifespans of up to 10 years for constrained terrestrial devices. However, the architectural shifts required to bridge the 85% global coverage gap via satellite connectivity fundamentally alter the energy expenditure profile of the UE. To establish a theoretical baseline for empirical comparison, it is necessary to examine how the physical constraints of NTN inflate the overall energy budget per transaction.

2.3.1 Energy Dynamics in NB-IoT

The total energy E_{total} consumed by an IoT device during a single communication cycle is the integral of its power draft over time:

$$E_{total} = \int_0^{T_{cycle}} P_{state}(t) dt$$

Where P_{state} represents the power consumption during specific operational states, including RRC Idle, RRC Connected (Transmit/Receive), and PSM. In LPWAN design, the objective is to minimize the time t spent in high-power states, specifically the active Transmit and Receive phases, and maximize the time spent in deep sleep (PSM), where current draw is typically measured in single-digit microamperes.

2.3.2 Impact of Path Loss and Repetitions on Transmit Power

In terrestrial networks, the eNodeB utilizes dynamic power control to instruct the UE to lower its transmit power when the radio link is strong. A device operating near a cell tower may successfully transmit utilizing less than 0 dBm of power. Furthermore, in optimal TN conditions, a transmission requires only a single Physical Uplink Shared Channel (NPUSCH) subframe.

Conversely, GEO NTN deployments operate at the absolute limits of the NB-IoT Link Budget due to the extreme Free Space Path Loss, approaching 190 dB for L-band and S-band. Consequently, an NTN UE is almost exclusively forced to operate at its maximum transmission power at 23 dBm to close the link [2].

Furthermore, to combat the high signal attenuation, the network dictates extreme Coverage Enhancement levels, requiring the UE to execute dozens or hundreds of blind physical layer repetitions. Academic evaluations of NB-IoT over satellite indicate that the energy consumed by the power amplifier during these prolonged, maximum-power repetition cycles constitutes the largest individual energy penalty in the NTN architecture, frequently resulting in a 300% to 500% increase in energy-per-bit compared to nominal TN operations [15].

2.3.3 Energy Penalty of Latency and Synchronization

The time domain t is equally critical to the energy equation. As established in Section 2.1, GEO NTN environments exhibit RTT exceeding 500 *ms*. This latency forces the UE to remain in the active Rx state, listening for network responses, for significantly longer durations than in a terrestrial network.

Additionally, the 3GPP Rel-17 requirement for GNSS-based TA pre-compensation introduces a major, unavoidable power drain. In a TN environment, the UE synchronizes blindly with the eNodeB using minimal energy. In an NTN environment, the device must activate its internal GNSS receiver, acquire a satellite lock, and download ephemeris data before it can even initiate the cellular attach procedure [16]. The Time-To-First-Fix (TTFF) for a cold-start GNSS module can range from 30 to 60 seconds, during which the module continuously draws significant current. However, for static IoT deployments, this substantial energy overhead can be entirely circumvented by pre-provisioning the device's known geographic coordinates directly into the Microcontroller Unit (MCU) firmware, thereby bypassing the GNSS acquisition phase.

2.4 Antenna Fundamentals: Linear vs. Circular Polarization

In terrestrial cellular networks, antenna design for UE primarily prioritizes cost, compactness, and multi-band support. Standard terrestrial base stations and mobile devices predominantly utilize LP antennas. However, the transition to NTN introduces specific propagation and geometric challenges that fundamentally conflict with standard terrestrial antenna architectures.

2.4.1 Principles of Polarization

Polarization describes the orientation and trace of the electric field vector \vec{E} of an electromagnetic wave as it propagates through space.

Linear Polarization: For linear polarized signals the electric field vector oscillates strictly within a single, fixed plane, either vertical or horizontal, along the axis of propagation. Standard cellular IoT devices use LP antennas, such as standard chip antennas or Planar Inverted-F Antennas, because they are highly space-efficient and cost-effective. In terrestrial environments, severe multi-path scattering randomizes the polarization of the arriving wave, making the strict geometric alignment of transmit and receive antennas less critical.

Circular Polarization: For circular polarized signals the electric field vector consists of two orthogonal linear components of equal magnitude, with a phase difference of 90°. As the wave propagates, the tip of the electric field vector traces a circle, rotating either clockwise (Right-Hand Circular Polarization, RHCP) or counter-clockwise (Left-Hand Circular Polarization, LHCP).

2.4.2 Satellite Downlink Polarization Constraints

In GEO satellite communications, maintaining a stable radio link geometry is a primary engineering constraint. Unlike terrestrial base stations, which communicate with devices in relatively fixed terrestrial planes, satellites must establish links with devices across vast geographic areas and in unpredictable orientations.

If a GEO satellite were to transmit a linearly polarized signal, the receiving antenna on the ground would need to be physically aligned with the incoming signal's polarization plane to avoid severe cross-polarization attenuation. For mobile IoT applications, such as asset trackers on maritime vessels or automotive telematics modules, maintaining this strict physical orientation relative to a satellite 35,786 kilometers away is impossible. As the device moves, turns, or tilts, the geometric misalignment would cause massive fluctuations in the received signal strength.

To ensure consistent signal reception regardless of the terrestrial device's physical orientation, GEO satellite downlinks are predominantly transmitted using Circular Polarization. Because a circularly polarized wave rotates continuously by design, a receiving antenna will capture the same relative signal profile regardless of how it is rotated around the axis of propagation [3].

2.4.3 Polarization Mismatch Loss

While transmitting CP from the satellite resolves the orientation problem, it creates a fundamental hardware conflict for hybrid Terrestrial/Non-Terrestrial IoT devices. To capture the maximum energy of a CP wave, the receiving antenna should also be CP, perfectly matching the handedness (e.g., RHCP to RHCP).

If a standard, commercial off-the-shelf LPWAN device attempts to receive a CP satellite signal using its standard LP cellular antenna, a physical phenomenon known as Polarization Mismatch Loss occurs. An LP antenna can only couple with one of the two orthogonal components of the arriving CP wave.

The Polarization Loss Factor (PLF) represents the ratio of the power received by an antenna to the power of the incident wave. When an LP antenna receives a CP wave, the theoretical PLF is exactly 0.5. In decibels, this mismatch loss L_{pol} is calculated as:

$$L_{pol} = 10 \log_{10}(0.5) \approx -3 \text{ dB}$$

This intrinsic penalty means that a standard terrestrial IoT device permanently loses half of the incoming satellite signal power purely due to antenna geometry [17].

2.4.4 Multipath Propagation and Signal Reflection

In real-world IoT deployments, a clear Line-of-Sight (LOS) to the satellite is frequently obstructed by buildings, foliage, or the device's own enclosure. Under these harsh operational conditions, the receiver must rely on multipath propagation to establish a connection. While the direct LOS wave is attenuated or blocked, the reflected waves scatter and converge at the receiving antenna. The efficiency with which an antenna captures these reflected signals is heavily dependent on its polarization characteristics [18].

When a circularly polarized electromagnetic wave reflects off a surface, its geometric trace fundamentally alters. According to the boundary conditions of electromagnetic reflection on a conductive or semi-conductive surface, the phase of the electric field component parallel to the surface is inverted.

Consequently, the handedness of the circular wave is reversed upon reflection—a phenomenon formally known as polarization reversal. If a satellite transmits an RHCP signal, the primary direct LOS wave arrives at the ground as RHCP. However, any portion of that signal that bounces off the ground, a vehicle roof, or a nearby structure undergoes this phase inversion and will arrive at the receiving antenna as a LHCP wave [19].

If the IoT device is equipped with a specialized RHCP antenna designed to perfectly match the satellite's direct downlink, it will inherently reject the reflected LHCP signals. This phenomenon is known as Cross-Polarization Isolation. The theoretical PLF for an RHCP antenna receiving an LHCP wave is 0, meaning infinite attenuation. In practical engineering, a well-designed CP antenna typically exhibits a cross-polarization discrimination of 15 dB to 20 dB. Thus, the CP antenna effectively filters out the reflected signals, treating them as noise rather than usable energy [3].

The physical rejection of reflected waves by CP antennas creates a significant disadvantage in Non-Line-of-Sight environments where the direct RHCP wave is heavily obstructed. In such scenarios, the reflected LHCP waves constitute the majority of the available signal energy.

A standard LP antenna does not discriminate between RHCP and LHCP waves. Because linear polarization consists of a single vector plane, it couples with any circular or elliptically polarized wave with a constant Polarization Loss Factor of 0.5.

This creates a mathematically definable advantage for standard cellular hardware in multipath environments.

Therefore, in environments dominated by reflections, an LP antenna will capture significantly more signal energy than an RHCP antenna [20]. Furthermore, as multipath waves scatter through dense environments they frequently depolarize into random linear components. An LP antenna can constructively harvest these linear reflections, whereas a strict CP antenna suffers continuous mismatch losses against depolarized waves.

Consequently, while CP antennas are strictly superior for unobstructed LOS satellite links, as long as the link budget allows standard LP antennas present a highly resilient, cost-effective compromise for harsh IoT environments where devices rely heavily on signal scattering and reflections.

2.5 End-to-End Cellular IoT Architecture

The deployment of a cellular IoT solution extends far beyond the physical radio link. To successfully transport a physical sensor reading to a cloud-based application, a systematic orchestration of hardware, telecommunication networks, and software protocols is required. As defined by the GSMA Mobile IoT Deployment Guidelines and 3GPP specifications, a complete End-to-End cellular IoT architecture is theoretically modeled across three distinct topological tiers: the Device Domain, the Network Domain, and the Cloud Platform Domain [21], [22].

Understanding the distinct responsibilities and constraints of each tier is critical, as inefficiencies at any level cascade through the system, ultimately degrading the device's battery life [11].

2.5.1 Device Domain

The Device Domain represents the physical edge node, commonly referred to in 3GPP terminology as the UE. In LPWAN architectures, this domain is characterized by severe constraints regarding computational power, memory, and, most importantly, available energy.

The primary objective of the UE is to wake up, acquire physical sensor data (e.g., GNSS coordinates or environmental metrics), transmit the data as rapidly as possible, and return to a deep sleep state. A standard theoretical UE architecture consists of:

Microcontroller Unit: The application processor responsible for executing the embedded firmware, managing sensor I/O, and serializing the payload.

Cellular Baseband Modem: The dedicated radio processor that handles the complex 3GPP physical layer, MAC layer, and RRC signaling.

UICC / SIM: The Universal Integrated Circuit Card hosting the Subscriber Identity Module (SIM), responsible for cryptographic network authentication.

According to the GSMA IoT Device Connection Efficiency Guidelines, the software architecture on the UE must be explicitly designed to leverage network power-saving features like PSM and eDRX [11]. The UE must offload all computationally heavy tasks (such as data enrichment or complex cryptography) to the cloud domain to conserve its limited energy budget.

2.5.2 Network Domain

The Network Domain acts as the secure transit layer, bridging the constrained UE with the public internet. This domain is subdivided into the Radio Access Network (RAN) and the Evolved Packet Core (EPC) [22].

Radio Access Network: The RAN provides the physical wireless connection. In a hybrid Terrestrial and Non-Terrestrial Network (TN/NTN) architecture, the RAN encompasses both standard terrestrial eNodeBs and satellite constellations. According to 3GPP specifications, these satellites operate using one of two primary payload architectures: a transparent payload often referred to colloquially as a "bent-pipe" configuration, which simply filters, amplifies, and reflects the RF signal down to a terrestrial gateway connected to an on-ground eNodeB or a regenerative payload, wherein the eNodeB baseband processing capabilities are fully integrated onboard the spacecraft itself [2].

Regardless of the payload type, the RAN dictates the physical link budget, handles HARQ retransmissions (if enabled), and manages the Coverage Enhancement (CE) levels required to service edge-of-cell devices.

Evolved Packet Core: Once the data successfully traverses the RAN, it enters the telecommunications core. The EPC is responsible for mobility management, authenticating the SIM, and routing the data. For IP-based NB-IoT traffic, the data is routed through the Serving Gateway (SGW) and the Packet Data Network Gateway (PGW), which translates the cellular bearer traffic into standard TCP/UDP packets formatted for the public internet [22].

The Network Domain is a "black box" for the IoT system engineer; it strictly adheres to 3GPP Release 13-17 standards and cannot be modified by the end-user. Therefore, the device and cloud domains must be engineered to tolerate the latency and packet loss inherent to the Network.

2.5.3 Cloud Platform Domain

The final tier is the Cloud Platform Domain, which serves as the destination for the UE's telemetry. Because the Device Domain is strictly optimized to send the smallest possible binary or array-based payloads, the raw data arriving from the Network Domain is typically unstructured and lacking context.

The Cloud Domain is responsible for:

Ingestion and Protocol Translation: Terminating the CoAP, UDP, or TCP sockets originating from the device.

Data Decoding and Enrichment: Expanding the highly compressed edge payloads back into standardized IT formats (such as fully qualified JavaScript Object Notation (JSON)). This often includes timestamping the arrival to calculate network latency and correlating the payload with external Application Programming Interfaces (APIs), e.g. weather or mapping services.

Storage and Application Logic: Persisting the telemetry in a database and triggering the final business logic, e.g., alerting a logistics manager that a tracking module has deviated from its route.

By concentrating the computational burden and complex data structuring within the Cloud Domain, the architecture ensures that the remote, battery-powered UE expends energy solely on maintaining the physical radio link [21]. This theoretical division of labor forms the structural blueprint for the empirical measurement system designed and implemented in Chapter 3

3. System Design & Measurement Setup

3.1 Solution Architecture

To empirically evaluate the theoretical physical and protocol-level boundaries established in Chapter 2, a comprehensive end-to-end cellular IoT measurement system was designed and implemented. The deployment of a hybrid Terrestrial/Non-Terrestrial cellular IoT solution requires the orchestration of constrained embedded hardware, optimized transport protocols, and cloud services.

In systems engineering, this lifecycle is modeled across a multi-tier architecture, dividing the solution into the Device Domain, the Network Domain, and the Cloud Platform Domain [21]. To ensure high external validity and accurately reflect a production-ready deployment, the measurement setup was engineered to mirror a real-world cellular IoT solution as closely as possible. Consequently, every hardware, software, and network component within this architecture was systematically selected and configured in strict accordance with public industry best practices and standardization recommendations, specifically adhering to the GSMA Mobile IoT Deployment Guidelines for hardware architecture [23] and the IETF Architectural Considerations for Smart Object Networking for protocol stack implementation [24].

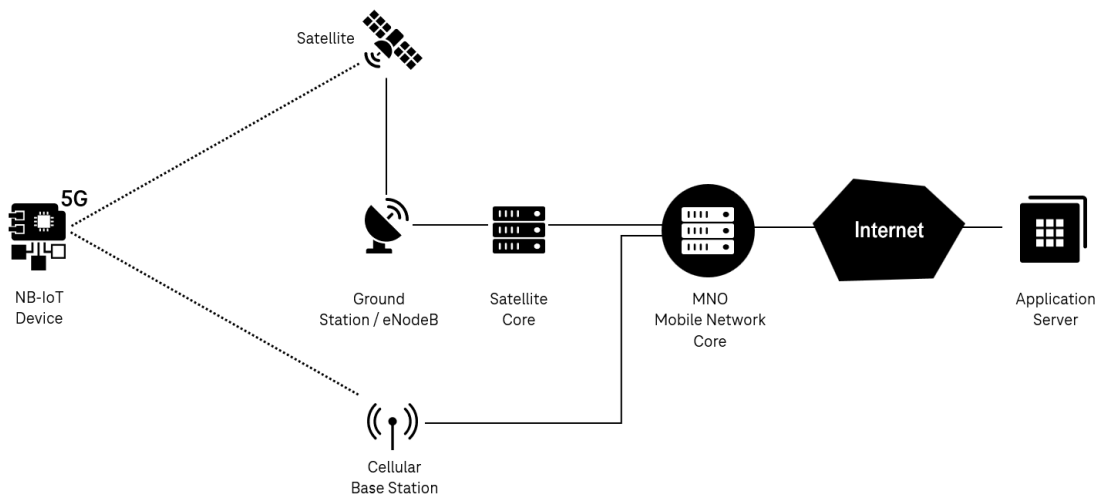


Figure 1: End-to-End IoT Solution Architecture for TN and NTN Converged Devices

3.1.1 Device Domain: Hardware and Firmware

The foundation of the measurement architecture is the physical edge device. In constrained NB-IoT deployments, the hardware architecture must strictly adhere to the energy budget defined by the intended deployment lifespan. For this thesis, the device domain acts as the active testing probe, designed to operate at the absolute limits of the link budget.

The critical hardware components integrated into this design include the MCU, the cellular baseband modem, the SIM, and modular antennas to evaluate the impact of polarization as theoretically defined in Section 2.3. At the firmware level, the application logic was specifically tailored to the constraints of the hardware. This includes the implementation of efficient power management features such as PSM and

eDRX which allow the modem to sleep deeply while maintaining its network registration context [25]. The specific hardware component selection and the Zephyr Operating System (OS) embedded firmware implementation are detailed in Section 3.2.

3.1.2 Network Domain: Connectivity and Transport

The Network Domain encompasses the RAN, utilizing both terrestrial eNodeBs and GEO satellites, and the Core Network that routes the data to the public internet. The primary engineering challenge addressed within this domain is the selection, implementation, and empirical comparison of the communication stack.

As analyzed theoretically in Section 2.2, standard internet protocols are frequently too heavy for LPWAN links. Therefore, the implementation phase of this architecture integrates both connection-oriented (TCP) and lightweight, connectionless application protocols (UDP, CoAP). This allows for a direct, real-world comparison of the protocols.

3.1.3 Cloud Domain: IoT Platforms and Integration

The final tier of the architecture is the Cloud Domain, which acts as the bridge between the highly constrained edge device and the data analysis layer. Because NB-IoT devices transmit optimized, highly compressed payloads these payloads cannot be directly utilized without an intermediary processing stage.

Instead of relying on a commercial black-box IoT platform, a custom Receiving Server Architecture was engineered for this thesis. This server acts as the network gateway, required to ingest the TCP, UDP, and CoAP packets, decode the binary payloads into standardized formats, and mathematically evaluate the transmission metrics. Crucially, this cloud layer is responsible for calculating the end-to-end latency by comparing device-generated timestamps against server-side reception timestamps. The engineering of this cloud architecture and the data enrichment processes are detailed in Section 3.3, while the structural design of the payload itself is analyzed in Section 3.4.

3.2 Device Architecture

The device domain serves as the active measurement probe for all empirical campaigns conducted in this thesis. To accurately characterize the performance of hybrid terrestrial and satellite networks, the measurement device was engineered as an autonomous, event-driven IoT sensor node. The architecture consists of a state-of-the-art cellular System-in-Package (SiP), modular antenna configurations, and a custom-developed firmware stack operating on a Real-Time Operating System (RTOS).

3.2.1 Hardware Components

Cellular Modem

The core computing and communication module selected for this architecture is the Nordic Semiconductor nRF9151. This SiP was chosen because it is currently the world's first and only commercial cellular IoT module to integrate hardware support for both Geostationary and Low Earth Orbit NB-NTN connectivity, alongside standard terrestrial LTE-M and NB-IoT. The SiP incorporates an ARM Cortex-M33 application processor and a highly energy-efficient baseband modem.

A critical physical constraint of the nRF9151, common in highly integrated cellular IoT modules to reduce hardware cost and footprint, is that the GNSS receiver and the cellular modem share a single radio receiver. Consequently, concurrent operation of GNSS and cellular transmission is physically impossible. This time-multiplexed hardware architecture strictly dictates the design of the embedded firmware, requiring precise orchestration to prevent module faults during measurement campaigns.

The nRF9151-SMA-DK with the modem firmware version nrf_9151-ntn_1.0.0-1.alpha was utilized to provide a robust and proven hardware foundation for the measurement campaigns. This development platform was specifically chosen for its modular RF front-end, which enables the efficient interchange of various antenna configurations.

Antenna Configurations

To empirically answer Research Question 3 regarding the necessity of specialized hardware, two distinct antenna configurations were integrated into the test setup:

Linear Reference Antenna: The Kyocera AVX X9003334 was utilized as the baseline terrestrial antenna. Designed as a broadband external 5G antenna, it provides a standard linearly polarized radiation pattern across a wide spectrum. While highly versatile across standard cellular frequencies, it exhibits a peak gain of 2.9 dBi and an average efficiency of 59% in the 1400–1500 MHz range, and a peak gain of 3.57 dBi with 62% efficiency in the 1700–2400 MHz range. Its flexible, commercial off-the-shelf architecture makes it the ideal control variable for evaluating how standard terrestrial hardware performs when subjected to constrained satellite link budgets.

Circularly Polarized Antenna: The Kyocera AVX 9002418L0-L16L was integrated to provide a highly directional, circularly polarized reference. Manufactured using Laser Direct Structuring (LDS) technology, this "Cap" antenna is explicitly optimized for the NTN L-Band, precisely covering 1525–1559 MHz and 1626.5–1660.5 MHz (Band n255). It features Right-Hand Circular Polarization with an exceptionally pure axial ratio of < 0.8 dB at 1540 MHz and < 1.2 dB at 1640 MHz. The hardware delivers a peak gain between 1.6 dBi and 1.7 dBi, maintaining an average efficiency ranging from 47% to 54% within the target satellite bands.

Comparing the empirical performance of these two antennas strictly quantifies the operational advantage of CP over LP hardware in practical NTN deployments.

3.2.2 Firmware Architecture and Concurrency Model

The embedded application was developed utilizing the nRF Connect Software Development Kit (SDK), built upon the Zephyr Real-Time Operating System. To accurately profile the baseline power consumption of Terrestrial Networks versus Non-Terrestrial Networks (RQ4), the software architecture must not introduce artificial energy overhead. Consequently, the firmware is strictly designed around an event-driven, cooperative concurrency model.

Following system initialization, the application relies exclusively on Zephyr's deferred execution model via kernel work queues. By eliminating busy-wait polling loops, the MCU can instantly enter low-power sleep states whenever the cellular modem is inactive. This architectural choice ensures that the power measurements captured during the campaigns reflect the true energy expenditure of the modem's network attach and message sending procedures untainted by inefficient application-layer CPU cycles.

3.2.3 RAT Management and GNSS Data Synchronization

To evaluate the impact of environmental factors and high-speed mobility on satellite link stability (RQ2), and to accurately measure end-to-end network latency (RQ1), the measurement device relies on the GNSS subsystem as its absolute source of truth. The GNSS receiver is utilized to extract three critical variables for every measurement cycle:

1. **Geographic Coordinates:** To correlate specific signal drops or attenuation events with physical environments.
2. **Velocity and Heading Vectors:** To analyze the impact of automotive mobility and directional movement on the satellite link's stability.

3. **Absolute UTC Time:** The GNSS is also used to acquire an absolute Coordinated Universal Time (UTC) time reference. This synchronized time is used to stamp every payload at the moment of transmission, enabling the server to calculate end-to-end latency with millisecond precision.

To resolve the cellular and GNSS multiplexing constraint, a custom Radio Access Technology (RAT) Manager was engineered. The RAT Manager operates as a centralized state machine that enforces a strict single-active-RAT invariant while seamlessly orchestrating context switches between the satellite network and the GNSS subsystem.

To prevent connection timeouts during these transitions, the RAT Manager utilizes event-driven RRC state tracking. The system suppresses GNSS acquisition while the modem is in an active *RRC Connected* state, ensuring that ongoing payload transmissions or network handshakes are never interrupted.

Once the modem transitions to *RRC Idle* the RAT Manager safely executes a highly specific context switch:

1. It suspends the active cellular context without dropping the IP session.
2. It reconfigures the modem to switch to the GNSS receiver.
3. After acquiring the geolocation fix and extracting the updated UTC time and velocity vectors, the GNSS subsystem is explicitly powered down.
4. The baseband multiplexer is switched back, and the cellular or satellite link is restored.

This software-controlled multiplexing architecture ensures that every transmitted payload is accurately geo-tagged and synchronized, providing the necessary spatial and temporal data to analyze how physical location and automotive movement degrade the NTN link budget without violating the physical limitations of the SiP.

3.2.4 Protocol Abstraction and Campaign Orchestration

A core objective of this thesis is to quantify how the choice of transport and application-layer protocols impacts transmission reliability and latency in high-path-loss GEO environments (RQ1). To facilitate a scientifically rigorous comparison, the software architecture implements a Protocol Abstraction Layer using a polymorphic design pattern.

A virtual method table (vtable) defines standardized lifecycle interfaces (connect, send, disconnect) for UDP, TCP and CoAP. The overarching Campaign Framework utilizes these abstracted interfaces to generate identical payloads and timing sequences regardless of the underlying protocol. This architectural separation of concerns guarantees that any measured variations in protocol overhead, connection time, or latency are attributable to the protocol specifications themselves, rather than discrepancies in the application logic.

Supported Custom Command API (ATAPP%)

To provide granular control over the hardware during isolated testing and debugging, the custom command parser supports a comprehensive suite of single-action instructions. These are categorized by subsystem:

Network and RAT Control:

- ATAPP%LTEM[=apn], ATAPP%NBIOT[=apn], ATAPP%NBNTN[=apn]: Initiates an attach procedure to the specified radio access network with an optional custom Access Point Name (APN).
- ATAPP%LTEMd, ATAPP%NBIOTd, ATAPP%NBNTNd: Disconnects the modem from the respective network.

Geolocation (GNSS) Control:

- ATAPP%GNSS: Acquires a single geolocation fix.
- ATAPP%GNSSP , ATAPP%GNSSS: Starts periodic GNSS tracking or stops the active GNSS subsystem.

Protocol Transmission:

- ATAPP%UDP, ATAPP%TCP, ATAPP%COAP, ATAPP%HTTP: Dispatches a single payload over the specified protocol. These commands accept optional overrides for destination server, port, payload content, and byte padding (e.g., ATAPP%UDP=192.168.1.100,8080).

System and Logging Operations:

- ATAPP%LOGDUMP / ATAPP%LOGCLEAR: Extracts or deletes measurement logs stored in the device's non-volatile LittleFS flash memory.

Automated Campaign Parameterization (ATAPP%MEAS)

While single commands are utilized for debugging, the primary operational mode relies on the ATAPP%MEAS command. This single execution string instantiates the autonomous state machine, configured by eight comma-separated integer parameters:

ATAPP%MEAS=<gnss>,<rat>,<proto>,<msg>,<pad>,<edrx>,<psm>,<term>

The arguments directly map to the research methodologies required to answer the thesis objectives:

- **<gnss>**: Defines the geolocation refresh rate. Valid states are -1 (continuous tracking), 0 (off), 1 (single initial fix), or >1 periodic refresh interval in seconds. *Critical for RQ2 (Mobility/Environment).*
- **<rat>**: Selects the active cellular layer, where 1 = LTE-M, 2 = NB-IoT, and 3 = NB-NTN.
- **<proto>**: Dictates the application stack, where 1 = UDP, 2 = TCP, 3 = CoAP, 4 = CoAPS (CoAP & DTLS), and 5 = HTTP. *Critical for RQ1 (Protocol Efficiency).*
- **<msg>**: The messaging frequency in seconds, it dictates the wait time between payload transmissions.
- **<pad>**: The number of pseudo-random bytes injected into the payload. This artificially stresses the constrained satellite link budget to test throughput limits.
- **<edrx>**: A boolean flag to disable or enable extended Discontinuous Reception.
- **<psm>**: A boolean flag to disable or enable Power Saving Mode.
- **<term>**: A boolean flag (0/1) to toggle serial logging. Disabling terminal output allows the MCU to power down the UART peripherals, eliminating idle current draw during power profiling campaigns.
- ATAPP%MEASSTATS / ATAPP%MEASSTOP: Retrieves real-time statistics of an active campaign or safely terminates the running state machine.

By encapsulating these parameters into a modular framework, the device seamlessly pivots between Static attenuation tests, Dynamic automotive measurements, and controlled Power Profiling without requiring firmware recompilation. This architectural automation guarantees high consistency, strict timing, and reproducibility across the thousands of data points required to answer the research questions.

3.3 Payload Design and Data Modeling

In highly constrained LPWAN the structural design of the application payload is generally as critical as the underlying transport protocol. Every transmitted byte consumes finite energy, prolongs the active radio state, and increases the probability of packet collision or corruption during the extended satellite propagation delay.

The primary objective of these measurement campaigns is not to engineer the most perfectly optimized IoT application, but rather to conduct systematic limit testing on the underlying Radio Access Technologies. To accurately evaluate the boundary conditions of the NTN link budget, such as protocol overhead penalties (RQ1) and maximum throughput under varying weather conditions (RQ2), the system requires highly flexible, easily modifiable payload sizes and transparent data structures. Consequently, the telemetry payload was designed to be "sufficiently optimized" to function over satellite links, intentionally sacrificing absolute binary efficiency in favor of the modularity and observability required for rigorous empirical research.

3.3.1 Positional Schema and Serialization Efficiency

To achieve this balance between observability and constraint compliance, the measurement device avoids standard, fully qualified JSON objects with descriptive key-value pairs (e.g., {"temperature": 22.5, "humidity": 40}).

To systematically eliminate this inefficiency without resorting to the aforementioned opaque binary blobs, the architecture implements a Positional Array Schema. The primary transmission payload is an ordered JSON array. By stripping all key identifiers and relying strictly on index position for data mapping, the baseline payload size is compressed to approximately 120 bytes. The schema is defined as:

```
[IMEI, MsgID, "Content", RAT, RSRP, RSRQ, RSSI, SINR, "Protocol", PLMN, TAC, CID, "SentAt", "Position"]
```

This architecture provides the optimal balance for the measurement setup: it maintains human-readable debuggability for the Node.js backend while ensuring the radio module spends its energy transmitting actual telemetry rather than structural formatting.

3.3.2 Dynamic Payload Padding and Compression Resistance

While the 120-byte baseline is good for standard mobility and signal tracking, evaluating protocol overhead and energy-consumption (RQ1 and RQ4) requires the ability to transmit standardized, larger data sizes. To accommodate this, the firmware features a dynamically configurable Payload Padding Mechanism, accessible via the ATAPP%MEAS campaign command.

When padding is requested (e.g., forcing a 200-byte or 500-byte payload), the device dynamically injects a calculated sequence of characters into the "Content" string field until the total serialized JSON length matches the exact target size.

Crucially, this padding does not consist of repeating static characters (e.g., a string of zeros). Modern cellular modems and core networks frequently implement Packet Data Convergence Protocol (PDCP) layer optimizations, such as Robust Header Compression (RoHC) or payload compression. If the device transmits a 500-byte payload of repeating zeros, the network layer might compress this into 20 bytes over the air.

To actively prevent this, the firmware utilizes the RTOS random number generator (`sys_rand32_get()`) to construct a high-entropy string of pseudo-random alphanumeric characters. Because this high-entropy

data strongly resists compression algorithms, it forces the baseband modem to physically transmit the exact number of bytes requested over the radio interface.

3.3.3 Cellular Connection Parameters and Signal Metrics

To evaluate physical link stability (RQ2) and the polarization mismatch loss of the antennas (RQ3), the payload must carry real-time radio metrics captured at the exact moment of the transmission attempt. A measurement taken 10 seconds prior is invalid in a mobile automotive context. Prior to assembling the JSON array, the application immediately queries the modem's baseband (AT%XMONITOR) to extract the following critical parameters:

Signal Metrics: Reference Signal Received Power (RSRP, in dBm), Reference Signal Received Quality (RSRQ, in dB), Received Signal Strength Indicator (RSSI, in dBm), and Signal-to-Interference-plus-Noise Ratio (SINR, in dB). In satellite communications, the thermal noise floor is exceptionally high; therefore, tracking SINR alongside RSRP is critical for accurately quantifying the attenuation experienced by the linear versus circular antenna configurations.

Network Identity: Public Land Mobile Network (PLMN) ID, Tracking Area Code (TAC), and Cell ID (CID). Because the device operates in a multi-orbit roaming environment, these identifiers prove definitively whether a specific packet traversed the terrestrial cellular network or successfully handed over to the satellite NTN infrastructure.

3.3.4 Spatial and Environmental Data Aggregation

To assess the impact of high-speed mobility and the physical environment on the satellite link budget (RQ2), high-resolution spatial data is mandatory. However, splitting coordinates into multiple float fields consumes valuable payload space.

To optimize this, the payload parser consolidates the geographic data into a single, comma-separated string within the final array index: "lat,lon,alt,speed,heading".

Velocity Vectors: The speed and heading are extracted directly from the modem's native GNSS hardware solver. This provides the continuous vector data necessary to correlate physical automotive maneuvers.

Environmental Correlation: The receiving server utilizes these precise coordinates to asynchronously fetch meteorological data. By embedding the GNSS string in the payload, the design inherently supports the analysis of atmospheric attenuation.

3.3.5 Time Synchronization and Latency Profiling

Calculating end-to-end latency in a hybrid multi-orbit network is a architectural challenge.

Internal Real-Time Clocks (RTC) on IoT devices suffer from clock drift, making them unreliable for accurate profiling over long measurement campaigns. To resolve this, the firmware utilizes the `time_sync` module to extract the absolute UTC time directly from the GNSS constellation during the initial location fix. This highly accurate time is used to generate an ISO 8601 timestamp at the exact moment the socket transmission is initiated by the application thread.

When the payload arrives at the receiving server which is strictly synchronized to Network Time Protocol (NTP) stratum servers, the backend derives the transmission delay:

$$Latency_{E2E} = T_{Server(Received)} - T_{Device(Sent)}$$

It must be noted that application-layer timestamping is not the optimal solution for strict, sub-millisecond accurate latency profiling. The measurement inherently includes minor systemic jitter introduced by RTOS thread scheduling, AT command processing, and modem baseband queues. While this systemic jitter

would be a limiting factor for benchmarking ultra-low-latency terrestrial New Radio (NR) networks, it remains a highly efficient and sufficient methodology for the scope of this research. Because the NB-IoT and NB-NTN messaging delays observed in these campaigns typically span several seconds or even tens of seconds, a millisecond-range timing variance is statistically negligible for the comparative analysis.

3.4 Receiving Server Architecture

To ingest, analyze, and store the telemetry generated by the device domain, a custom multi-protocol receiving server was engineered. Because commercial black-box IoT platforms obscure transport-layer metadata, they are insufficient for the granular protocol analysis required by this thesis. Therefore, the cloud backend was developed as a event-driven Node.js application, utilizing asynchronous I/O to handle high-latency satellite connections without blocking the main execution thread.

The server architecture is structured around a four-stage asynchronous pipeline: Reception, Immediate Response, Data Enrichment, and Storage. This non-blocking design ensures that the server immediately acknowledges incoming packets before offloading the computationally expensive enrichment and database operations to the background.

3.4.1 Multi-Protocol Ingestion Layer

To systematically evaluate how different application and transport-layer protocols affect transmission efficiency (RQ1), the server implements independent protocol listeners operating simultaneously. Each listener is tailored to the specific state mechanics of its respective protocol:

UDP: Connectionless UDP datagrams are parsed immediately upon reception.

TCP: To accommodate the high latency and potential fragmentation of the satellite link, the TCP listener maintains persistent connections with per-client state. It utilizes a dynamic brace-balancing algorithm to reassemble partial JSON segments across multiple TCP frames before triggering the parsing logic.

CoAP: Due to the payload size variations introduced by the measurement device's padding feature, the server implements manual block-wise transfer reassembly. It intercepts raw UDP packets, indexes incoming blocks and concatenates the complete payload upon receiving the final block marker.

3.4.2 Network Overhead and Latency Analysis

A primary objective of the measurement campaigns is to quantify the exact protocol overhead, the "handshake tax" and header sizes, and the true end-to-end latency introduced by the GEO NTN architecture. To extract this data at the network layer, the server integrates a real-time Wireshark (tshark) packet capture subsystem.

The packet capture subsystem operates as a persistent child process, monitoring all active protocol ports. It utilizes an EventEmitter pattern to correlate captured IP frame lengths and TCP/UDP port metadata with the application-layer messages. The overhead is dynamically computed as:

$$\text{Overhead} = \text{TotalFrameSize} - \text{DecodedPayloadSize}$$

To measure latency, the server requires strict chronological alignment. The server synchronizes its internal clock against a designated NTP pool (pool.ntp.org) every 15 minutes. Upon receiving the final packet of a transmission, the server captures a completedAt timestamp and compares it against the sentAt timestamp embedded within the device's payload. This differential provides an accurate measurement of the one-way transmission delay across the end-to-end link.

3.4.3 Asynchronous Data Enrichment

To evaluate the extent to which environmental factors degrade satellite link stability (RQ2), the measurement data must be correlated with meteorological conditions at the exact time and location of transmission.

Once the payload parser extracts the geographic coordinates from the device's telemetry string, the server initiates an asynchronous request to the OpenWeatherMap REST API. This enrichment module retrieves real-time meteorological data and appends it to the message object. Because weather enrichment relies on an external API, it is engineered as a non-critical, fail-safe operation: If the API times out or the device lacks a valid GNSS fix, the server proceeds with data storage using null weather values, ensuring no measurement data is lost.

3.4.4 Storage Architecture and Real-Time Distribution

The persistence layer is implemented using a PostgreSQL relational database. To support the execution of distinct measurement campaigns, the storage architecture utilizes a Dual-Write Pattern with fault isolation.

All incoming messages are strictly written to a mandatory primary table, guaranteeing a continuous, unbroken chronological log of all thesis measurements. Simultaneously, the server executes a secondary write to an active "Custom Table" defined dynamically by the system operator via the web UI. This allows campaign-specific data to be partitioned cleanly. The dual-write mechanism is fault-isolated: if a custom table write fails due to schema mismatch or misconfiguration, the failure is logged, but the primary write succeeds, preventing irrecoverable data loss.

Finally, to monitor campaigns in the field, the server integrates a Socket.IO WebSocket layer. Upon successful database insertion, the fully enriched measurement object is broadcast to all authenticated web clients.

3.5 Measurements

To systematically evaluate the performance boundaries of the NB-NTN architecture against terrestrial baselines, a comprehensive measurement methodology was established. This chapter is bifurcated into two primary components: Section 3.5.1 defines the Measurement Campaigns, detailing the operational tests designed to capture empirical data. Subsequently, Section 3.5.2 outlines the Measurement Setups, documenting the physical hardware configurations and environmental conditions required to execute these tests.

3.5.1 Measurement Campaigns

The empirical data collection is organized into three distinct operational domains: Static Campaigns, Dynamic Mobility Campaigns, and Power Profiling Campaigns. Each category is specifically tailored to isolate and quantify different physical and protocol-level constraints of the cellular networks.

Static Campaigns

The static campaigns provide the foundational baseline for network reliability, latency, and protocol efficiency (RQ1 and RQ3). These measurements were conducted from a fixed, open-sky rooftop location in an urban environment, eliminating mobility-induced Doppler shifts and frequent handovers.

Table 1: Overview of the Static Measurement Campaigns

Nr.	Duration	Location	RAT	Payload Size	Antenna Setup	Protocols
1	7d, every 10m	City, roof	NB-IoT	200 Byte	external linear polarized	UDP
2	7d, every 10m	City, roof	GEO NB-NTN	200 Byte	external linear polarized	UDP
3	24h, every 10m	City, roof	GEO NB-NTN	200 Byte	circular polarized	UDP
4	24h, every 10m	City, roof	GEO NB-NTN	500 Byte	circular polarized	UDP
5	24h, every 10m	City, roof	NB-IoT	200 Byte	external linear polarized	CoAP
6	24h, every 10m	City, roof	GEO NB-NTN	200 Byte	external linear polarized	CoAP
7	24h, every 10m	City, roof	NB-IoT	200 Byte	external linear polarized	TCP
8	24h, every 10m	City, roof	GEO NB-NTN	200 Byte	external linear polarized	TCP
9	24h, every 10m	City, roof	NB-IoT	200 Byte	external linear polarized -53db	CoAP
10	24h, every 10m	City, roof	GEO NB-NTN	200 Byte	external linear polarized -6db	CoAP
11	24h, every 10m	City, roof	NB-IoT	200 Byte	external linear polarized -40db	UDP
12	24h, every 10m	City, roof	GEO NB-NTN	200 Byte	external linear polarized -6db	UDP
13	24h, every 10m	City, roof	NB-IoT	200 Byte	external linear polarized -40db	TCP
14	24h, every 10m	City, roof	GEO NB-NTN	200 Byte	external linear polarized -6db	TCP
15	24h, every 10m	City, roof	NB-IoT	500 Byte	external linear polarized	CoAP
16	24h, every 10m	City, roof	GEO NB-NTN	500 Byte	external linear polarized	CoAP
17	24h, every 10m	City, roof	NB-IoT	500 Byte	external linear polarized	UDP
18	24h, every 10m	City, roof	GEO NB-NTN	500 Byte	external linear polarized	UDP
19	24h, every 10m	City, roof	NB-IoT	500 Byte	external linear polarized	TCP
20	24h, every 10m	City, roof	GEO NB-NTN	500 Byte	external linear polarized	TCP

Baseline Reference (Campaigns 1 & 2): To establish a statistically significant baseline, the device transmitted a 200-Byte UDP payload every 10 minutes over a continuous period. This was executed identically on both the terrestrial NB-IoT network and the GEO NB-NTN network using the standard linear polarized antenna.

Circular Polarized Antenna Performance (Campaigns 3 & 4): To isolate and quantify the impact of antenna polarization on the NTN link budget (addressing RQ3), the standard linear antenna was replaced with the circularly polarized antenna.

Protocol and Payload Stress Testing (Campaigns 5-8, 15-20): To quantify protocol performance and maximum throughput limits (directly addressing RQ1), the transport protocol was systematically rotated between UDP, CoAP, and TCP. TCP timeouts were specifically tuned for the high-latency NTN environment. Payload sizes were scaled between 200 Bytes and 500 Bytes to artificially stress the constrained satellite link budget and measure the resulting protocol overhead and latency.

Artificial Attenuation and Edge-of-Cell Testing (Campaigns 9-14): To simulate poor deployment locations (e.g., partial indoor or obstructed views) and support the environmental baseline for RQ2, inline RF attenuators were utilized. The terrestrial network was attenuated by up to -53 dB, and the satellite link by -6 dB, forcing the modem into extreme CE repetition levels to compare protocol resilience near the absolute cell edge.

Dynamic Campaigns

To address RQ2 regarding the viability of GEO NB-NTN for automotive tracking and logistics, the methodology includes a suite of dynamic mobility campaigns. These tests evaluate the modem's ability to maintain satellite synchronization and execute payload transmissions while subjected to Doppler shifts and rapidly changing physical obstructions in real world environments.

Table 2: Overview of the Dynamic Measurement Campaigns

Nr.	Duration	Location	RAT	Payload Size	Antenna Setup	Protocols
21	1h, every 2.5m	Autobahn and urban	NB-IoT	200 Byte	external linear polarized, in-car behind windshield	UDP
22	1h, every 2.5m	Autobahn and urban	GEO NB-NTN	200 Byte	external linear polarized, in-car behind windshield	UDP
23	1h, every 2.5m	Autobahn and urban	GEO NB-NTN	200 Byte	circular polarized, in-car behind windshield	UDP
24	1h, every 2.5m	Autobahn and urban	GEO NB-NTN	200 Byte	circular polarized, rear bench	UDP
25	1h, every 2.5m	Autobahn and urban	GEO NB-NTN	200 Byte	external linear polarized, rear bench	UDP
26	1h, every 2.5m	Autobahn and urban	GEO NB-NTN	200 Byte	circular polarized, roof	UDP
27	1h, every 2.5m	Autobahn and urban	GEO NB-NTN	200 Byte	external linear polarized, roof	UDP

Mobility Baselines (Campaigns 21 & 22):

The measurement node was deployed in a vehicle traveling on the German Autobahn at speeds ranging from 80 km/h to 150 km/h. Continuous GNSS tracking was enabled, and a 200-Byte UDP payload was dispatched every 2.5 minutes to compare TN and NTN link stability, specifically isolating the impact of velocity and Doppler shifts (RQ2).

Different Placement: A critical factor in automotive IoT is device placement. To evaluate penetration loss through the vehicle chassis (addressing the sub-optimal placement variable of RQ3 and the environmental factors of RQ2), the device was rotated through three distinct physical locations:

1. Mounted externally on the vehicle roof
2. Mounted internally behind the windshield
3. Placed deep within the cabin on the rear bench

Antenna Directivity in Dynamic Campaigns: Throughout the mobility testing, the external linear polarized antenna and the circularly polarized antenna were swapped to definitively answer RQ3, determining whether the high gain and optimized polarization of the CP antenna can overcome the severe penetration losses of an automotive chassis in motion.

Power Profiling Campaigns

To definitively answer RQ4, the methodology isolates the exact energy expenditure of the cellular modem. Because actual battery life is highly dependent on the network-dictated repetition levels, these measurements were conducted in a controlled RF environment using variable attenuators to simulate precise, repeatable link budgets.

The power profiling is divided into two distinct phases, comparing TN and NTN across varying RSRP thresholds:

1. **Network Attach Energy:** The total milliampere-second mAs charge required to transition the modem from a cold boot into a fully registered RRC Connected state. This isolates the attach procedure and the power to send messages required by GEO NTN.
2. **Messaging Energy:** Once attached, the payload sizes were scaled across 250, 500, 750, and 1,000 Bytes. The total active radio time and power consumption were measured for each transmission. By cross-referencing these payload sizes against the artificial signal degradation, the campaigns provide the exact empirical data required to calculate the optimal energy-per-bit ratio for satellite IoT deployments.

Specifically, the terrestrial NB-IoT measurements were executed under two distinct RF conditions: an optimal scenario (RSRP > -100 dBm) and a severely attenuated cell-edge scenario (RSRP < -130 dBm). This methodology allows the inherently constrained link budget of the NTN connection to be directly evaluated against both best-case and worst-case terrestrial deployments.

3.5.2 Measurement Setups

To execute the campaigns defined in Section 3.5.1, controlled physical testing environments were constructed. The structural arrangement of the hardware was strictly tailored to the specific objectives of the static, dynamic, and power profiling methodologies, ensuring that external RF interference was mitigated and real-world attenuation factors were accurately simulated.

Static and Environmental Setup

The long-term baseline and environmental stress tests were deployed on the rooftop of the Faculty of Digital Transformation at Hochschule für Technik, Wirtschaft und Kultur Leipzig, providing an open-sky, urban Line-of-Sight to the GEO satellite.




Figure 2: Static Measurement Setup

To ensure scientific rigor, it was imperative that the measured radio metrics (RSRP, SINR) reflected only the performance of the external antennas rather than parasitic reception from the nRF9151-SMA-DK PCB traces. Consequently, the development kit was encapsulated within an MTS MSB-0200 series RF Shielding Box.

To prevent near-field metallic interference from the shielding enclosure itself, a non-conductive polymer spacer was mounted on top of the MTS box, elevating the GNSS and cellular antennas into free space. This entire hardware stack was then sealed within a larger enclosure. This outer box acted as a functional case.

The measurement node was tethered via USB through ports on the shielding box to a laptop located indoors. This workstation provided continuous power and acted as the serial command interface. Utilizing the custom ATAPP%MEAS firmware API detailed in Section 3.2.4, the operator could remotely instantiate autonomous, week-long state machines without requiring physical access to the sealed rooftop node.

Power Profiling Setup

For the power profiling campaigns (RQ4), the static rooftop setup was augmented with precision measurement hardware to calculate the exact power consumption of the NTN link.

A Nordic Power Profiler Kit II (PPK2) with firmware version 1.0.1 was integrated directly into the physical enclosure. The PPK2 was wired inline with the power supply of the nRF9151 module to capture high-resolution, continuous milliampere (mA) readings. Through the tethered control laptop, shorter campaigns were executed while the PPK2 measured the power draw of the exact operational states of the modem. This setup enabled the exact calculation of the total milliampere-second (mAs) charge consumed during the extended NTN attach procedures and the transmission of the messages.

Dynamic Mobility Setup

To evaluate the viability of NB-NTN for automotive logistics and asset tracking (RQ2 and RQ3), the hardware was transitioned from the controlled rooftop environment into a mobile vehicular setup. For these dynamic campaigns, the MTS shielding box was removed to accommodate the spatial constraints of the vehicle.

To ensure strict comparability across all mobility tests, every dynamic campaign was executed along a consistently maintained, standardized test route traversing the greater Leipzig metropolitan area. The circuit commenced on the Autobahn A9 (west of Leipzig), transitioned northbound onto the A14, and exited in the East at junction 25 (Leipzig-Nordost). From there, the route proceeded directly through the dense

urban core of Leipzig, passing the central railway station, before returning to the point of origin. This standardized route covered a total distance of 52.2 kilometers and required approximately 50 to 60 minutes to complete per measurement campaign.

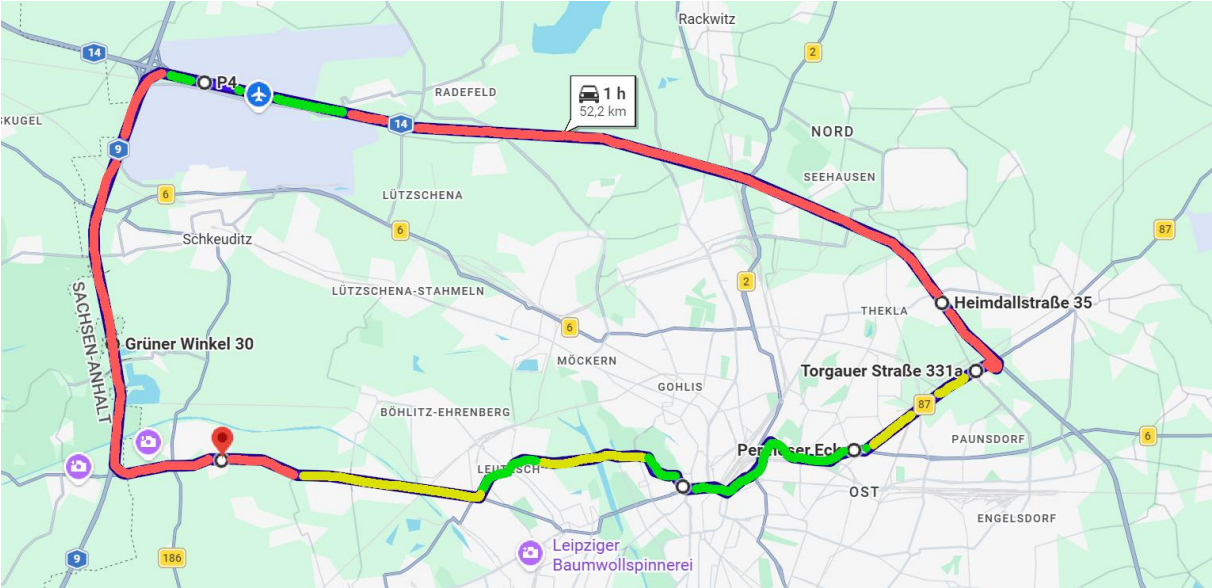


Figure 3: Color-coded Map of the Driving Route for the Dynamic Measurement Campaigns. Map data © 2026 Google, annotated by the author.

This specific trajectory was deliberately selected to expose the satellite link to a comprehensive spectrum of operational environments. To systematically evaluate network performance across varying degrees of structural obstruction and multipath fading, the route is geographically divided into three distinct segments, corresponding to the color-coding on the map:

Rural (Red on the map): Open highway segments characterized by high vehicle speeds and minimal structural obstruction, providing optimal Line-of-Sight to the GEO satellite.

Suburban (Yellow on the map): Intermediate transit corridors featuring moderate building density, foliage, and varying vehicle speeds.

Urban (Green on the map): Deep urban canyons within the city core, where severe structural shadowing and dense multipath fading fundamentally disrupt the direct satellite link.

To systematically profile the link's behavior under these varying automotive and environmental movements, the device and its antennas were rotated through three distinct physical placements during the test drives:



Figure 4: Dynamic Measurement Setups

Optimal Exterior (Roof-Mounted): The antennas were placed externally on the vehicle roof, providing an unobstructed Line-of-Sight to the satellite to simulate a best-case scenario and isolate the effects of high-speed Doppler shifts.

Partial Attenuation (Windshield): The setup was mounted on the dashboard immediately behind the windshield glass, representing a standard aftermarket tracking deployment with partial dielectric and metallic obstruction.

Severe Attenuation (Cabin Interior): The device was placed deep within the vehicle on the rear passenger bench, surrounding the antennas with heavy chassis metal to test the absolute limits of the NTN link budget.

Across all three placement scenarios, the campaigns alternated between the standard linearly polarized antenna and the specialized circularly polarized antenna to empirically quantify the operational advantage of optimized polarization in high-mobility environments.

4. Measurement Results & Analysis

4.1 Basic Analysis of NB-IoT and NB-NTN

Before evaluating specific optimization options for NTN it is necessary to establish the fundamental physical and network-layer baseline of the connection type.

4.1.1 NTN Cell Variability and Network Allocation

A fundamental characteristic of utilizing an aggregated NTN service provider, such as Skylo, is the dynamic allocation of satellite resources. For static deployments rather than connecting to a fixed, known eNodeB as in terrestrial NB-IoT, the measurement node dynamically attached to multiple distinct satellite Cell IDs across the duration of the campaigns. These Cell IDs represent different Geostationary satellites (e.g., from the Viasat or Echostar constellations) or overlapping spot beams covering the Leipzig test area.

Latency and RF Metric Discrepancies

The following figures show the results of the Cell ID specific analysis for latency, SINR and RSRP representing the most important connection parameters for IoT devices.




Figure 5: Distribution of NB-NTN Latency by Cell ID




Figure 6: Distribution of NB-NTN SINR by Cell ID




Figure 7: Distribution of NB-NTN RSRP by Cell ID

An analysis of the distribution data demonstrates extreme variance in end-to-end transmission delays depending on the allocated cell. For example, Cell ID 00201AD5 demonstrated the most efficient link, achieving a median latency of 11.73 s, with a highly compressed interquartile range indicating stable delivery times. Conversely, Cell ID 001A2B49 suffered from a median latency of 30.07 s with upper-whisker delays stretching to a severe 64.14 s.

Interestingly, the empirical data proves that this latency variance does not strictly correlate with physical signal strength.

The highly efficient Cell ID 00201AD5 actually operated near the absolute limit of the link budget, exhibiting the poorest median RSRP (-127 dBm) and a negative median SINR (-5 dB).

In stark contrast, the slowest network, Cell ID 001A2B49, provided the best physical radio conditions among all recorded cells, with a median RSRP of -123 dBm and a median SINR of 0 dB.

This inverse relationship indicates that in hybrid GEO NTN architectures, application-layer latency is often bottlenecked by satellite-specific performance, rather than pure Radio Frequency attenuation.

Despite the satellite specific differences observed between different cells, an aggregated analysis of all transmitted payloads reveals a definitive, overarching correlation between the physical SINR and end-to-end latency. When plotting the entirety of the measurement data across all static and dynamic NTN campaigns, the resulting scatter chart demonstrates a strong inverse relationship between signal quality and transmission delay.

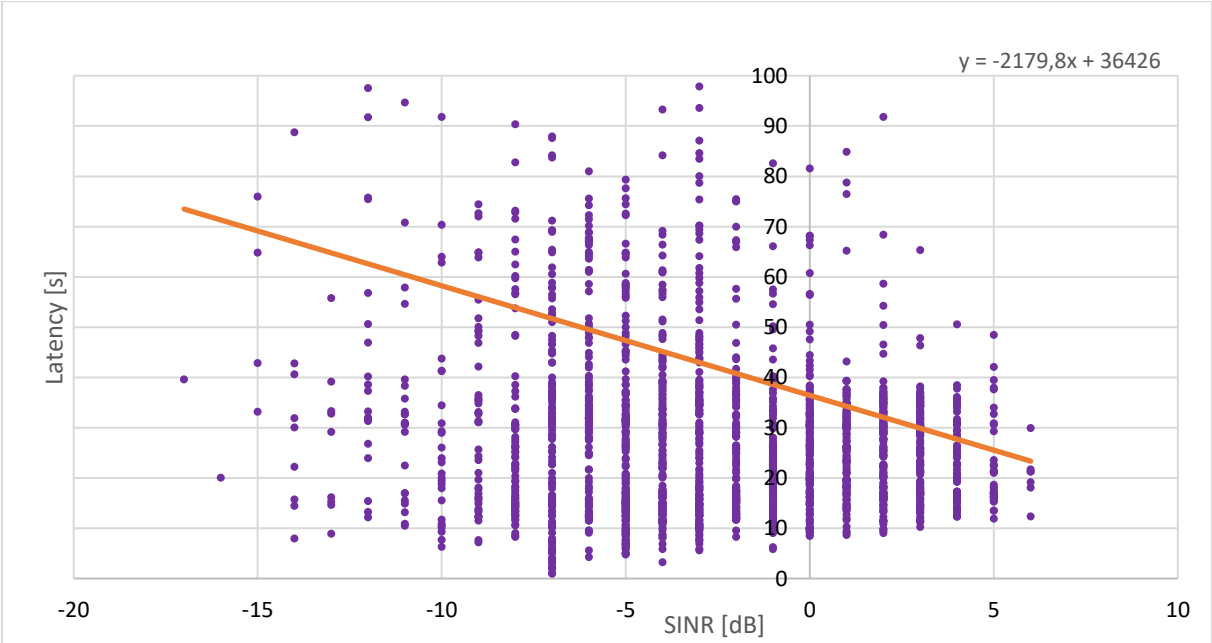


Figure 8: Correlation Between SINR and NB-NTN Message Latency

The linear regression of this dataset yields the following trendline equation:

$$Latency = -2037.4 * SINR + 37050$$

This equation indicates that for every 1 dB improvement in the Signal-to-Interference-plus-Noise Ratio, the end-to-end application latency decreases by approximately 2.03 s. The baseline latency at an SINR of 0 dB sits at approximately 37 s.

To contextualize the NTN findings, an identical latency-to-SINR analysis was performed on the terrestrial NB-IoT network. Under normal urban deployment conditions, such as the static rooftop setup, the TN link naturally exhibited highly favorable RF metrics. The measured RSRP consistently ranged between -80 dBm and -90 dBm, resulting in a stable and high SINR between 4 dB and 6 dB. Under these optimal conditions, transmission latency remained minimal and highly predictable.

However, to create a fair comparative baseline against the constrained satellite link, specific measurement campaigns were executed using inline RF attenuators to simulate extreme terrestrial cell-edge environments. By artificially depressing the RSRP down to the absolute threshold of -135 dBm, the corresponding SINR was forced as low as -8 dB.

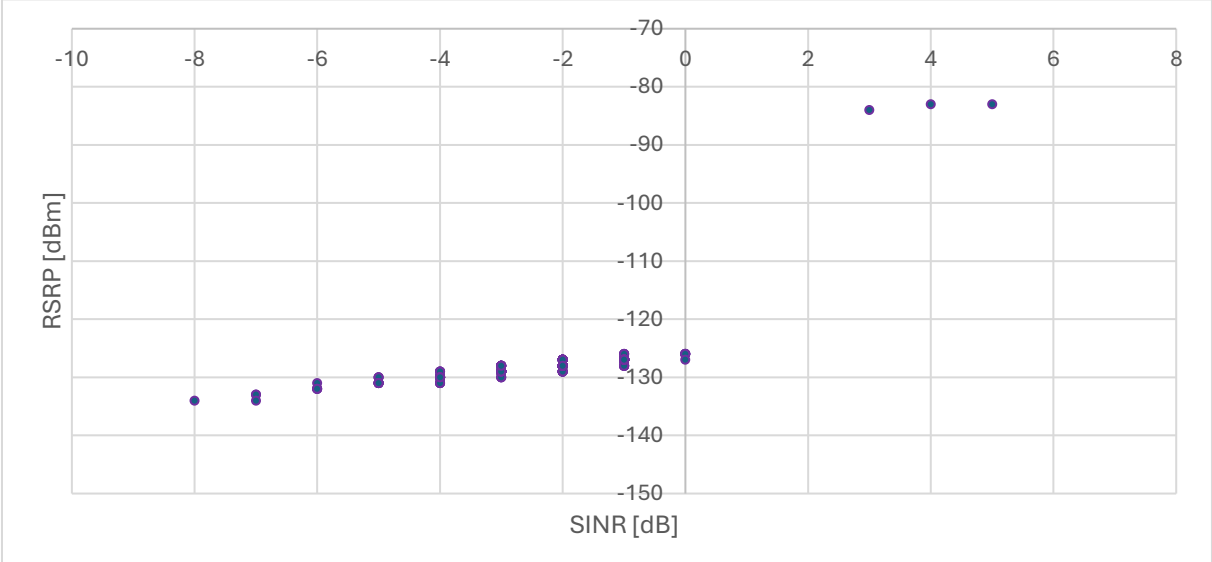


Figure 9: Correlation Between SINR and RSRP for NB-IoT

Plotting the TN latency against these attenuated SINR values reveals a correlation structurally similar to the NTN behavior. The empirical data clearly demonstrates that under severe signal degradation, terrestrial NB-IoT is behaving similar to NB-NTN. Transmission latencies were observed spiking up to 60 s at the absolute edge of the link budget. The linear regression for the terrestrial dataset yields the following trendline equation:

$$Latency = -1050.2 * SINR + 21127$$

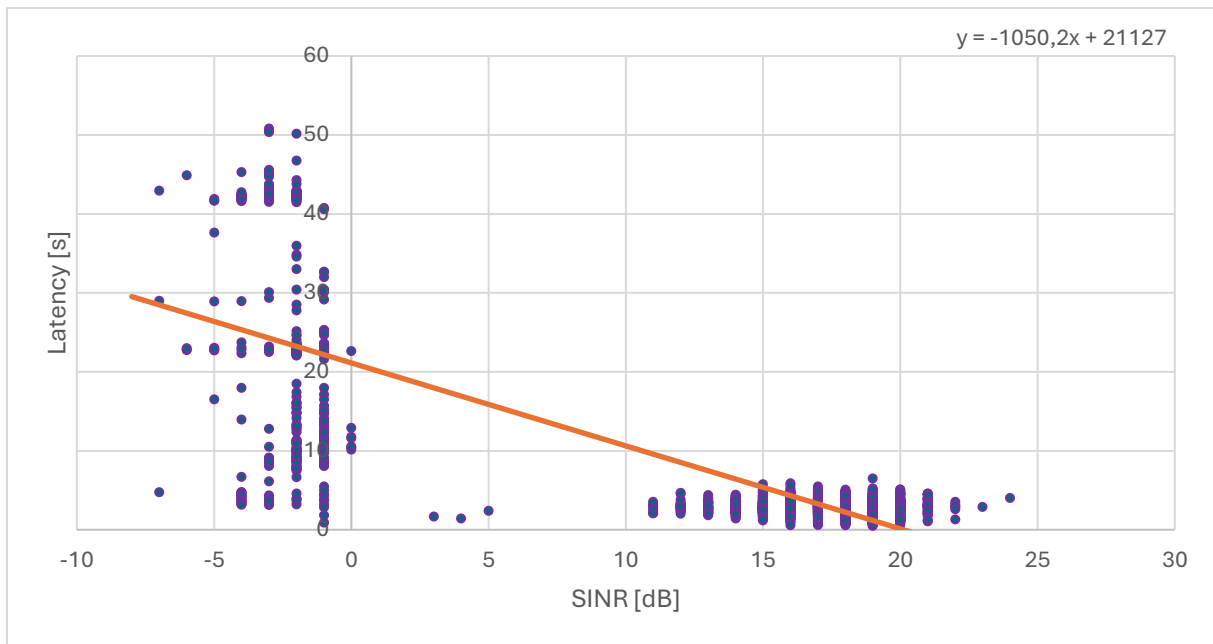


Figure 10: Correlation Between SINR and NB-IoT Message Latency

TN vs. NTN Comparison

Comparing the linear models of both networks highlights a shared behavioral characteristic, albeit with different boundary conditions. Both TN and NTN exhibit a strict negative correlation between signal quality and latency.

However, the terrestrial baseline latency at an SINR of 0 dB sits at approximately 21.1 seconds, which is significantly lower than the 37.0-second baseline of the GEO satellite link. Furthermore, the slope of the TN model (-1050.2 ms/dB) is approximately half that of the NTN model (-2037.4 ms/dB). This data establishes that while both architectures experience massive latency spikes at the edge of their respective link budgets, the terrestrial network degrades at a slower rate than the satellite connection when subjected to equivalent drops in signal quality.

4.1.3 Impact of Meteorological Conditions on NTN Performance

To determine the extent to which atmospheric attenuation affects the satellite link budget, the static rooftop measurements utilizing the standard linearly polarized antenna were cross-referenced with the asynchronous meteorological data captured via the OpenWeatherMap API. Because the GEO NB-NTN service operates in the lower-frequency L-band or S-band spectrum, it theoretically possesses a higher resilience to rain fade compared to higher-frequency Very Small Aperture Terminal (VSAT) systems operating on Ku or Ka-band. To empirically validate this, the SINR and end-to-end latency were categorized across six distinct weather states: Clear, Clouds, Drizzle, Mist, Rain, and Snow.

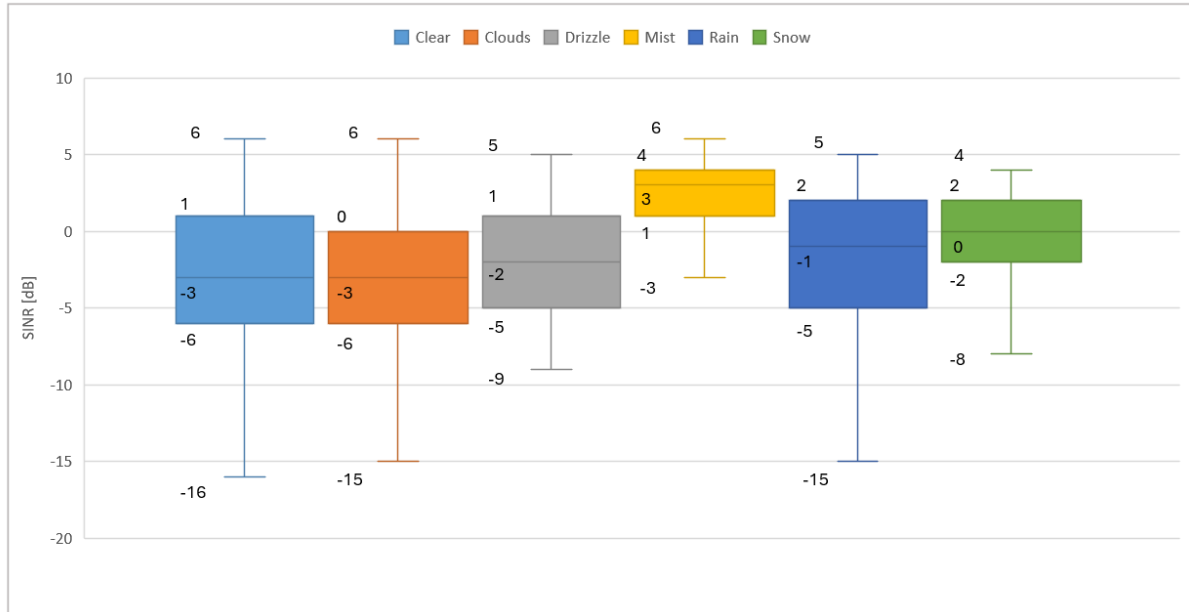


Figure 11: Correlation Between Weather and SINR for NB-NTN

The empirical data confirms the theoretical resilience of the lower-frequency satellite bands against standard precipitation and cloud cover.

- Under clear skies, the linear antenna recorded a median SINR of -3 dB.
- During periods of clouds and active Rain, the median SINR remained highly stable at -3 dB and -1 dB, respectively. The overall distributions across these three states were nearly identical, demonstrating that standard atmospheric moisture and heavy cloud cover do not introduce significant free-space path loss for NB-NTN deployments.
- The Mist category recorded a mathematically higher median SINR of 3 dB. However, with a highly limited sample size of only 39 data points, this is likely a temporal anomaly tied to a highly favorable satellite cell allocation at the time of measurement, rather than a meteorological benefit.




Figure 12: Correlation Between Weather and Latency for NB-NTN

While standard rain and clouds do not degrade the link, the latency metrics reveal a specific vulnerability to snow.

Under baseline conditions (Clear, Clouds, Rain), the median transmission latency remained tightly grouped between 24.7 seconds and 25.1 seconds. However, during periods of snow, the median latency spiked to 33.3 seconds.

More notably, the variance in the snow data is highly compressed, with the lower quartile sitting at 32.34 s and the upper quartile at 36.9 s. Unlike rain, which simply passes through the signal path, snow accumulates directly on the external measurement enclosure. This accumulation potentially alters the near-field electromagnetic environment, such as by inducing localized signal reflections or scattering within the setup, which may account for the observed shift in transmission latency.

4.1.4 Impact of Payload Size on Transmission Latency

To further evaluate the throughput limits of the network architectures, an additional measurement campaign was conducted to analyze the impact of increased message volume on transmission time. The standard 200-Byte payloads were replaced with padded 500-Byte payloads, artificially stressing the constrained link budgets of both TN and NTN.

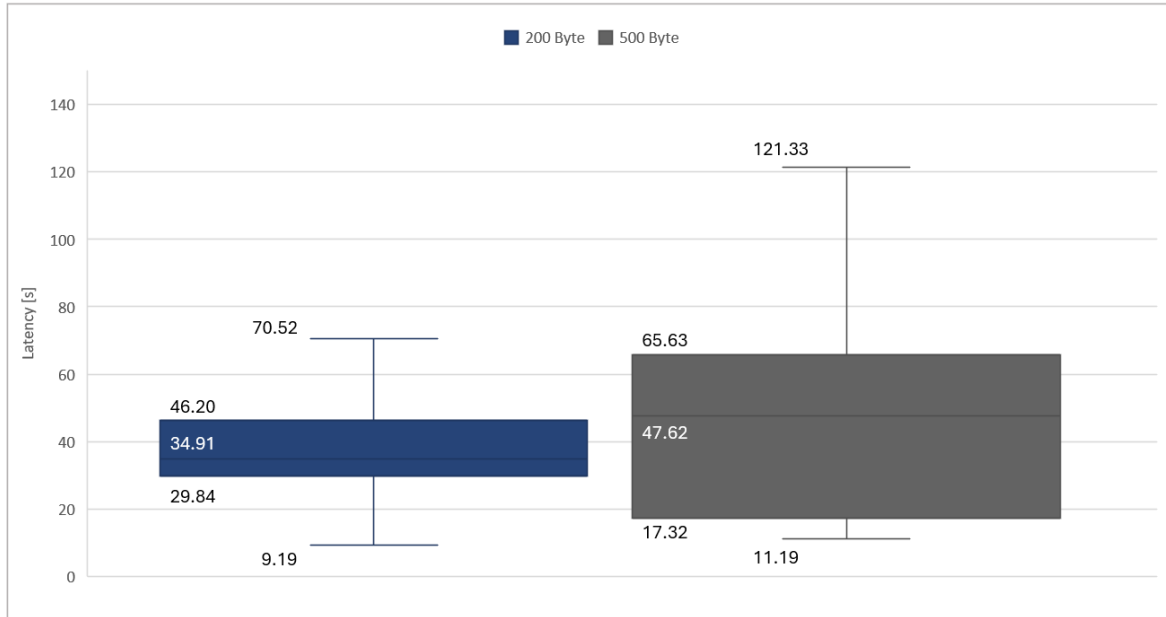


Figure 13: Correlation of Latency and Payload Size for NB-NTN

For the satellite link, increasing the payload size resulted in a substantial deterioration of both median latency and overall delivery predictability. When transmitting the standard 200-Byte payload, the NTN link exhibited a median latency of 34.91 s, with an upper whisker reaching 70.52 s.

However, scaling the payload to 500 Bytes caused a severe expansion in the transmission delay. The median latency increased to 47.62 s, representing a roughly 12.7-second penalty. More critically, the jitter in delivery times expanded drastically: the upper quartile shifted to 65.63 s, and the upper whisker stretched to an extreme 121,33 s. This indicates that while the median delay scales somewhat linearly, the worst-case transmission times for larger payloads on NTN can easily exceed two full minutes.




Figure 14: Correlation of Latency and Payload Size for NB-IoT

In stark contrast, the terrestrial NB-IoT baseline demonstrated extreme resilience to the increased data volume. Under good TN conditions with very high link budgets, the transition from a 200-Byte to a 500-Byte payload induced only a fractional delay.

For the 200-Byte baseline, TN achieved a median latency of 2.76 s, with a tightly bound upper whisker of 3.87 s. When the payload was increased to 500 Bytes, the median latency shifted to only 3.13 s, an increase of just 376 milliseconds. The overall distribution remained highly stable, with the upper whisker barely extending to 4.49 s.

This empirical data illustrates a fundamental divergence in network behavior regarding data throughput. On a terrestrial NB-IoT network with a good connection quality, increasing the payload size within the limits of the standard has a negligible impact on latency, meaning systems engineers can freely aggregate data into larger packets. However, on a GEO NB-NTN link, increasing the payload from 200 to 500 Bytes introduces a disproportionate penalty in both median latency and maximum delay.

4.1.5 Network Reliability and Message Drop Rates

The fundamental reliability of the connection is quantified by the message drop rate (packet loss). To accurately assess the raw physical capability of the TN and NTN links, this analysis isolates the campaigns utilizing UDP. Because UDP is a connectionless, "fire-and-forget" protocol, any messages lost are not recovered by the transport layer, providing a direct metric of link stability.

Table 3: Message Drop Rates Across TN and NTN Measurement Campaigns

Campaign Nr.	01	02	11	12	17	18	19
Sent Messages	1044	144	159	146	456	144	84
Received Messages	1041	131	159	132	456	142	84
Dropped Messages	3	13	0	14	0	2	0
Droprate	0.29%	9.03%	0.00%	9.59%	0.00%	1.39%	0.00%

The terrestrial NB-IoT network demonstrated near-perfect reliability across all tested parameters. During the 7-day baseline reference campaign, the TN successfully transmitted 1,041 out of 1,044 messages, resulting in a negligible drop rate of just 0.29%.

Remarkably, this high reliability remained untouched by extreme physical stress. When the payload was scaled up to 500 Bytes and when the signal was artificially attenuated by -40, the terrestrial network recorded a 0.00% drop rate. This proves that terrestrial NB-IoT, utilizing standard Coverage Enhancement repetitions, are highly capable of masking physical layer degradation from the application layer.

Conversely, the GEO NB-NTN link exhibited significant vulnerability when using the same standard commercial off-the-shelf hardware. In the baseline reference test, the satellite network dropped 13 out of 144 messages, resulting in a severe drop rate of 9.03%. When the linear antenna was further subjected to a -6 dB artificial attenuation to simulate a sub-optimal deployment, the drop rate remained similarly high at 9.59%.

However, the empirical data also demonstrates that the NTN connection is not inherently or permanently unstable. It is subject to high temporal variability. For instance, during Campaign 18, which utilized the exact same linear antenna and network configuration but transmitted larger 500-Byte UDP payloads, the network dropped only 2 out of 142 messages, yielding a vastly improved failure rate of just 1.41%. This sharp contrast between a 9.5% and a 1.4% packet loss rate across different measurement windows highlights the unpredictable nature of the satellite link.

The data clearly illustrates a critical architectural divergence: while terrestrial NB-IoT provides a highly predictable and reliable transport medium even under edge-of-cell conditions, deploying the exact same hardware on a GEO NTN network introduces unpredictable packet loss that can spike to nearly 10%.

4.2 Transmission Protocol Comparison

To evaluate the impact of transport and application-layer protocols on network efficiency, the measurement campaigns systematically tested UDP, CoAP and TCP. The analysis focuses on two primary metrics: the protocol overhead and the resulting message drop rate. These metrics were captured across both the TN and NTN using 200-Byte and 500-Byte payloads.

4.2.1 Protocol Overhead Analysis

The data demonstrates severe differences in data efficiency depending on the selected protocol. Because the device must enter a power-saving state after transmission, connection-oriented protocols inherently require complete session establishment and teardown for each message, directly impacting the overhead.

Table 4: Protocol Performance Comparison in TN and NTN

Network	Protocol	Payload Size	Overhead / Payload	Drop Rate
TN	UDP	200 Bytes	0.14	0.29%
	UDP	500 Bytes	0.06	0.00%
	CoAP	200 Bytes	0.54	0.00%
	CoAP	500 Bytes	0.22	1.35%
	TCP	200 Bytes	1.45	0.00%
	TCP	500 Bytes	0.58	0.00%
NTN	UDP	200 Bytes	0.14	9.03%
	UDP	500 Bytes	0.06	1.39%
	CoAP	200 Bytes	0.54	0.00%
	CoAP	500 Bytes	0.23	3.17%
	TCP	200 Bytes	1.47	7.32%
	TCP	500 Bytes	0.51	1.41%

UDP:

As a connectionless protocol, UDP exhibited the lowest overhead across all test cases. For both TN and NTN, the 200-Byte payload resulted in a static overhead-to-payload ratio of 14%. When the payload was increased to 500 Bytes, the fixed header size was diluted, reducing the ratio to 5.6%.

CoAP:

Operating over UDP, CoAP adds application-layer reliability features. This resulted in a moderate increase in overhead. For a 200-Byte payload, the ratio was measured at 54% on NTN and 53.8% on TN. Scaling to 500 Bytes reduced this overhead to approximately 22.9% on NTN and 21.6% on TN.

TCP:

The three-way handshake and subsequent teardown required by TCP generated the most significant data burden. For a 200-Byte payload, TCP produced an overhead ratio of 147% on NTN and 145% on TN, meaning

the modem transmitted significantly more protocol metadata than actual sensor data. Even when utilizing a 500-Byte payload the overhead ratio remained high at 51.1% on NTN and 57.9% on TN.

4.2.2 Network Reliability and Drop Rates

While protocol overhead directly correlates to active radio time, the message drop rate dictates the functional reliability of the link. The measurement data reveals that the network environment fundamentally alters how these protocols perform.

Terrestrial Network

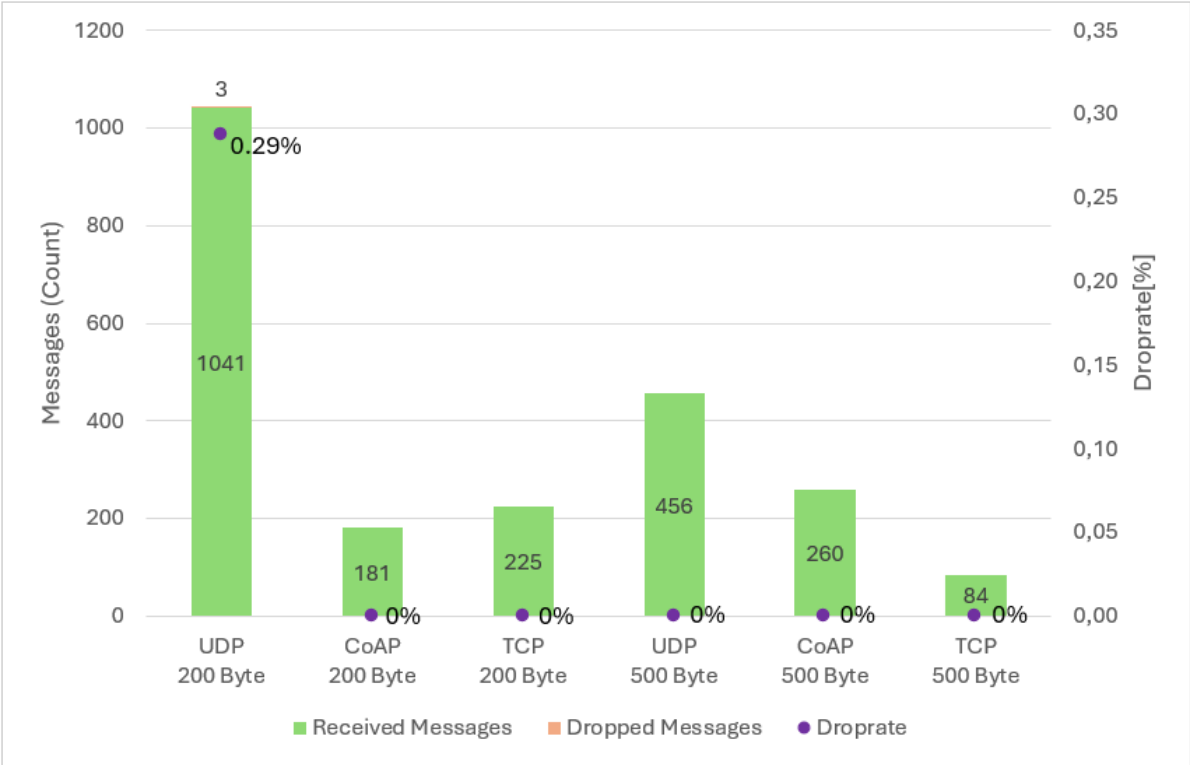


Figure 15: Received vs. Dropped Message by Protocol and Message Size for NB-IoT

On the stable terrestrial NB-IoT network, the physical layer is highly reliable, meaning the choice of transport protocol has minimal impact on the success rate.

TCP achieved a 0.00% drop rate for both 200-Byte and 500-Byte payloads.

CoAP similarly maintained a 0.00% drop rate for 200 Bytes and a low 1.35% drop rate at 500 Bytes.

Even unacknowledged **UDP** shows only a 0.29% drop rate at 200 Bytes and 0.00% at 500 Bytes.

The TN data shows that while TCP guarantees delivery, it introduces over 140% overhead to solve a packet loss problem that does not exist on the terrestrial baseline.

Non-Terrestrial Network

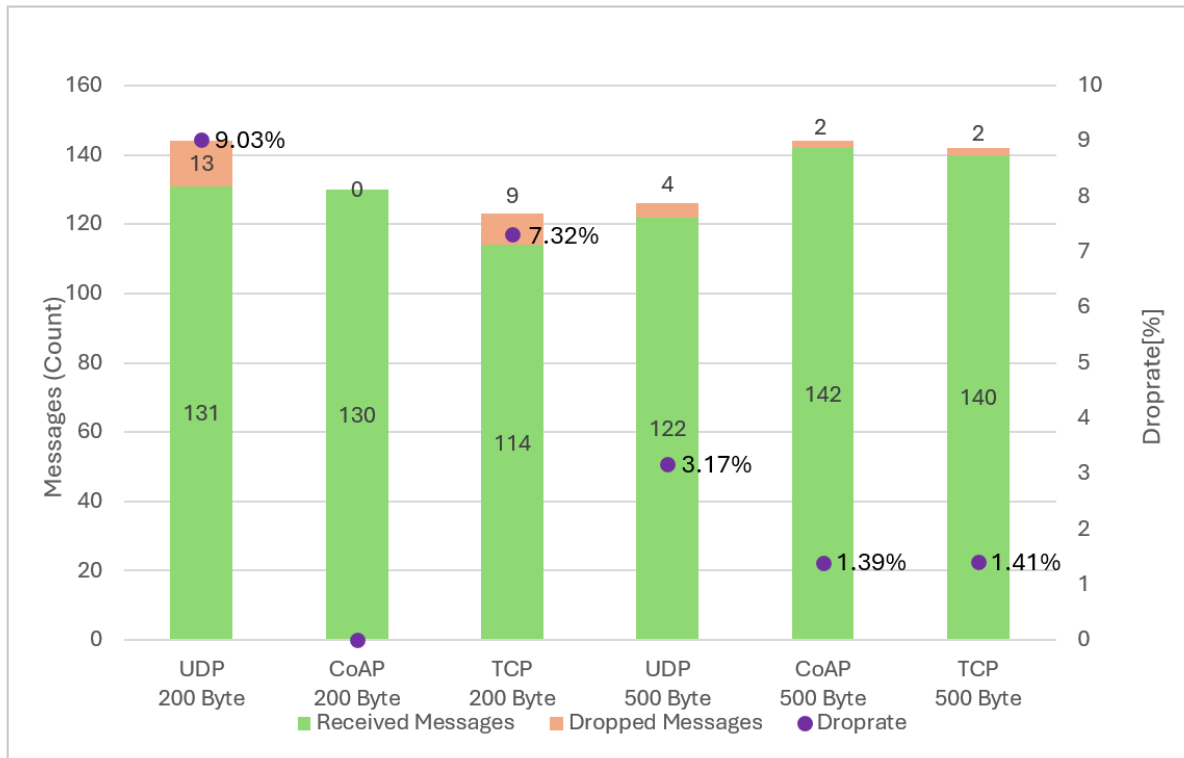


Figure 16: Received vs. Dropped Message by Protocol and Message Size for NB-NTN

The high-latency, high-path-loss environment of the GEO satellite link yielded starkly different reliability results.

UDP: Without transport-layer acknowledgments, the physical degradation of the NTN link resulted in a 9.03% drop rate for 200-Byte UDP payloads. When scaled to 500 Bytes, the drop rate varied, recorded at 1.39%.

TCP: Despite generating 147% overhead, TCP failed to improve reliability over the satellite link. For the 200-Byte payload, TCP experienced a 7.32% drop rate. For the 500-Byte payload, the drop rate was 1.41%. The data indicates that TCP's strict internal timers and connection-oriented nature are highly susceptible to the 37-second baseline latency of the GEO link, resulting in timeouts rather than successful retransmissions.

CoAP: CoAP provided the highest reliability on the NTN link. By utilizing its own application-layer acknowledgment mechanism over connectionless UDP, the 200-Byte CoAP campaign achieved a 0.00% drop rate. When transmitting the larger 500-Byte payloads, CoAP experienced a 3.17% drop rate. This indicates that the CoAP implementation in the device or server might have constraints for higher payloads for NTN environments.

The data shows that in good environments for terrestrial networks, all protocols achieve near-perfect delivery. On the non-terrestrial network, raw UDP suffers from high baseline packet loss (~9%). However, attempting to fix this packet loss with TCP introduces a massive overhead penalty (>140%) while still suffering a comparable drop rate (~7.3%). CoAP emerges in the data as the only tested protocol capable of reducing the NTN drop rate to 0% while maintaining a moderate, scalable overhead profile.

4.3 Power Consumption Measurements

While protocol overhead and latency dictate the temporal efficiency of the network, the ultimate constraint for remote IoT devices is energy consumption. To empirically evaluate the battery drain of the terrestrial and non-terrestrial networks, high-resolution power profiles were captured using a Nordic Semiconductor Power Profiler Kit II. The measurements isolate the electrical charge, measured in milliampere-seconds mAs, required for the initial network attach procedure, as well as the active transmission states across varying payload sizes. Throughout these measurement campaigns, the nRF9151 module was supplied with a constant operating voltage of 3.3V. Consequently, to calculate the absolute total energy consumed by the module, this 3.3V baseline must be applied to the recorded mAs values. The commonly used term ‘power consumption’ will be used throughout the following analysis.

To provide a comprehensive baseline, the terrestrial NB-IoT measurements were executed under two distinct RF conditions: an optimal scenario (RSRP > -100 dBm) and a severely attenuated, edge-of-cell scenario (RSRP < -130 dBm). This methodology allows the inherently constrained link budget of the NTN connection to be directly evaluated against both best-case and worst-case terrestrial deployments.

4.3.1 Network Attach Energy Profile

Before any application data can be transmitted, the cellular modem must exit its deep sleep state, synchronize with the network, and establish an RRC Connected state. This attach procedure represents a significant, fixed energy cost.

Table 5: Network Attach Time and Power Consumption

Network Scenario	PLMN	RSRP Range [dBm]	Mean Attach Time [s]	Max Attach Time [s]	Mean Current Draw [mAs]	Max Current Draw [mAs]	Mean Current Draw [mA]
TN (Optimal)	26201	-53 to -54	3.95	5.2	59.08	129.2	20.94
TN (Moderate Signal)	26202	-84 to -89	34.63	35.3	916.7	950.0	26.48
TN (Edge-of-Cell, Attenuated)	26202	-128 to -134	270.7	445.8	9727.5	14290.0	23.90
NTN	90198	-120 to -126	93.3	120.0	6656.0	11290.0	71.3

The empirical data demonstrates a massive disparity in the energy required to attach depending on the network and physical conditions:

Terrestrial Network (Optimal): Under optimal conditions (-53 to -54 dBm), the TN attach procedure was highly efficient. The average time to attach was just under 4.0 seconds, resulting in a mean power consumption of 59.08 mAs. The current draw per second during this phase averaged 20.94 mAs/s.

Terrestrial Network (Edge-of-Cell): When artificially attenuated to extreme cell-edge limits (-128 to -134 dBm), the modem is forced to utilize maximum Coverage Enhancement repetitions to synchronize with the eNodeB. The attach time skyrocketed to an average of 270.7 seconds, with individual attempts taking over

445 seconds. The mean power consumption surged to 9,727.5 mAs (peaking at 14,290 mAs), with a sustained average current draw of 23.90 mA.

Non-Terrestrial Network: The satellite attach procedure requires negotiation over a high-latency link. Operating at an RSRP between -120 and -126 dBm, the NTN attach procedure took an average of 93.3 seconds. The mean power consumption spiked to 6,656.0 mAs, with individual measurement runs reaching as high as 11,290 mAs.

The data reveals a critical operational insight: attaching to the satellite network is over 110 times more electrically expensive than attaching to a local terrestrial cell under good conditions. However, a terrestrial device buried deep indoors (operating < -130 dBm) can actually consume *more* total energy during an attach procedure than a satellite device with a clear line of sight, due to the extreme length of the TN cell search and repetition cycles.

4.3.2 Active Transmission Power and Payload Scaling

Once attached, the energy required to transmit data and receive the corresponding network acknowledgments was measured. To evaluate how power scales with data volume, payloads from 150 Bytes up to 750 Bytes were transmitted across all three network scenarios.

Power Consumption Averages by Payload Size and Network Condition

Table 6: Energy Metrics for TN vs. NTN Transmissions

Network	Payload	Mean Active Time	Mean Current Draw	Efficiency	Average Current Draw
TN (Optimal)	150 Bytes	~2.4 s	31.2 mAs	0.208 mAs/Byte	3.26 mA
	250 Bytes	~1.8 s	22.4 mAs	0.089 mAs/Byte	11.30 mA
	500 Bytes	~2.8 s	37.1 mAs	0.074 mAs/Byte	16.20 mA
	750 Bytes	~2.7 s	39.1 mAs	0.052 mAs/Byte	17.47 mA
TN (Edge)	150 Bytes	~2.9 s	365.0 mAs	2.433 mAs/Byte	212.86 mA
	250 Bytes	~4.4 s	670.6 mAs	2.682 mAs/Byte	205.90 mA
	500 Bytes	~3.5 s	399.4 mAs	0.798 mAs/Byte	83.70 mA
	750 Bytes	~3.6 s	502.3 mAs	0.669 mAs/Byte	172.40 mA
NTN	150 Bytes	~33.8 s	4,843.0 mAs	32.28 mAs/Byte	154.94 mA
	250 Bytes	~59.9 s	9,083.0 mAs	36.33 mAs/Byte	157.97 mA
	500 Bytes	14.4 s – 120.2 s*	15,111.4 mAs	30.22 mAs/Byte	165.54 mA
	750 Bytes	~105.5 s	15,871.7 mAs	21.16 mAs/Byte	157.39 mA

(*Note: The NTN 500-Byte active time exhibited extreme variance in the raw data, with successful transmissions ranging from as low as 14.4 seconds to as high as 129.2 seconds, reflecting the temporal instability of the satellite link).

Data Analysis of Payload Efficiency

On the optimal terrestrial network, energy efficiency scales positively with payload size. As the payload increases from 150 to 750 Bytes, the cost per byte drops from 0.208 mAs/Byte down to 0.052 mAs/Byte. Because the TN link is highly stable, packing more data into a single transmission requires only marginally more active radio time, making large payloads highly efficient.

When the terrestrial network is severely attenuated, the power consumption jumps drastically. To push data through the poor link, the modem amplifies its average current draw and utilizes CE repetitions, causing a 500-Byte payload to consume 399.4 mAs, roughly 10 times more energy than the optimal TN scenario.

However, despite this degradation, the worst-case terrestrial scenario remains vastly more efficient than the NTN baseline. The NTN link exhibits a drastically different electrical profile. The power cost per byte hovers between 21 and 36 mAs/Byte. At the 750-Byte mark, the extended active time (>100 seconds) fundamentally drives the total power consumption up to an extreme 15,871.7 mAs.

4.3.3 Average Current Draw and Extrapolation

A critical metric derived from this data is the average current draw per second (mAs/s), which indicates the intensity at which the radio module is operating. This metric was calculated by applying a linear approximation through the measured data points, plotting the total power consumption (mAs) against the active transmission duration. The slope of this linear graph yields the average continuous current draw for each respective network state.

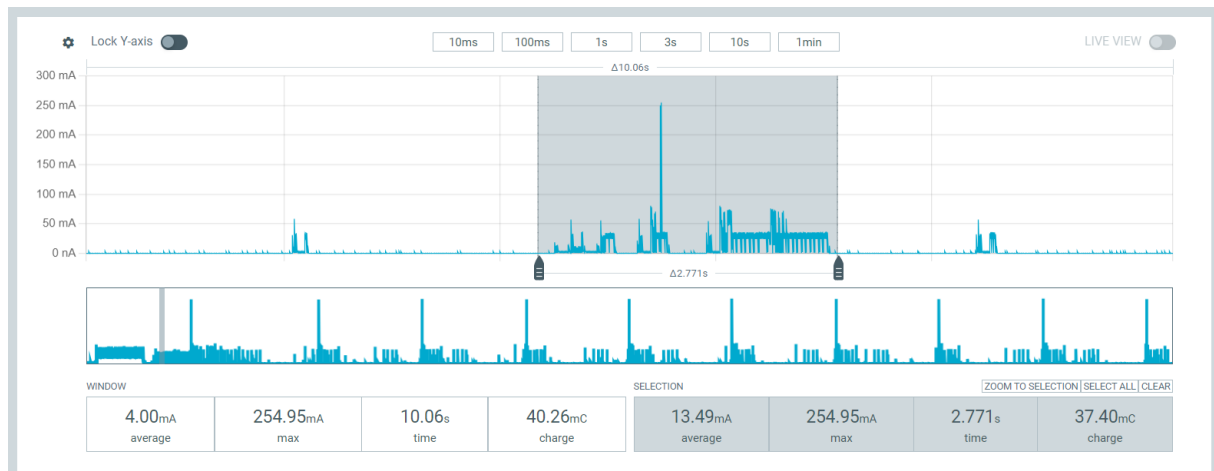


Figure 17: Power Profile of a 1,000 Byte TN Transmission

On the optimal TN link, the average current draw scales dynamically from 3.26 mA up to 17.47 mA, corresponding with conditions where the modem does not need to broadcast at maximum power.

In contrast, both the TN Edge and the NTN link force the modem to transmit at maximum power (typically 23 dBm). The NTN link maintains a consistently extreme current draw across all payloads, ranging from 154.94 mA to 168.5 mA. The TN Edge scenario recorded similar max-power spikes, peaking at 212.86 mA for smaller payloads.

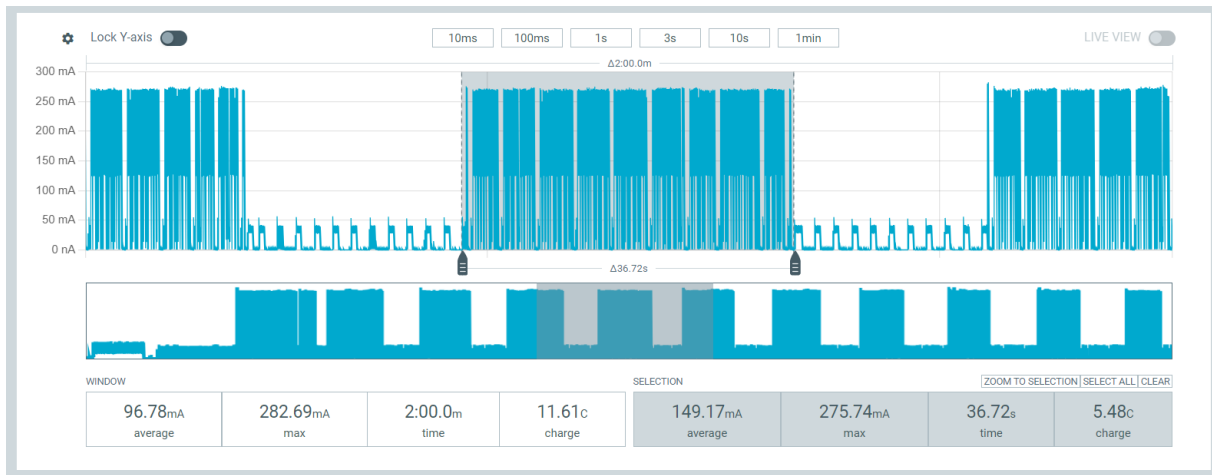


Figure 18: Power Profile of a 150 Byte NTN Transmission

The fundamental difference lies in the temporal domain. While a worst-case TN transmission draws massive instantaneous current, it completes the task in roughly 3 to 5 seconds. The GEO NTN connection draws similarly massive current, but sustains it for typically 10 to 30, but with possible spikes up to over 100 seconds. Therefore, the data conclusively demonstrates that even when a terrestrial NB-IoT device is deployed in absolute worst-case, edge-of-cell conditions requiring maximum repetitions, the GEO NB-NTN link still consumes approximately typically 10 to 20 times more energy to transmit a 250-Byte payload. With spikes of up to 30 times more energy in worse scenarios.

4.3.4 Summary of Energy Expenditure Multipliers

To synthesize the empirical power profiling data, it is highly useful to quantify the GEO NB-NTN energy demands as rough multipliers of the terrestrial baselines. This provides systems engineers with a clear heuristic for battery dimensioning when migrating an IoT device from a terrestrial to a multi-orbit architecture.

The Attach Phase:

Compared to TN Optimal: Attaching to the NTN link is extremely costly, demanding roughly 110 times more energy than a standard TN attach under optimal conditions.

Compared to TN Edge: However, this multiplier flips in extreme cell-edge scenarios. Because the terrestrial network relies on incredibly long repetition cycles at the absolute edge of its link budget (< -130 dBm), the NTN attach procedure consumes approximately the same (0.68x) the energy of the worst-case TN scenario.

While the attach phase shows parity with worst-case terrestrial environments, the active transmission of payload data reveals massive, systemic energy multipliers driven by the extended satellite latency.

Compared to TN Optimal: Transmitting data over NTN demands a catastrophic energy increase compared to a healthy terrestrial connection. For smaller payloads (150 Bytes), NTN draws approximately 150 times more energy. For larger payloads (250 to 750 Bytes), this multiplier scales severely, requiring 400 to over 420 times more energy than the optimal TN baseline.

Compared to TN Edge: Even when comparing NTN to a terrestrial device pushed to its absolute physical limits, the satellite penalty remains severe. Transmitting over NTN requires roughly 13 times more energy for smaller payloads, scaling up to 30 to 38 times more energy for larger payloads (500 to 750 Bytes) compared to the severely attenuated TN baseline.

Ultimately, the power profiling data proves that while the peak RF current draw (mA) is similar across all highly attenuated networks, the temporal expansion caused by the GEO satellite propagation delay

introduces severe energy multipliers. Consequently, devices utilizing NB-NTN require significantly larger battery capacities and highly restrictive payload management to achieve the same operational lifespan as standard terrestrial deployments.

4.4 Antenna Performance

To empirically address Research Question 3 regarding the necessity of specialized hardware for satellite IoT, dedicated measurement campaigns were executed exclusively over the GEO NB-NTN link. Because standard terrestrial NB-IoT devices predominantly utilize Linearly Polarized antennas, they inherently suffer from a theoretical 3 dB polarization mismatch loss when receiving circularly polarized satellite downlinks (as established in Section 2.4).

To quantify the real-world impact of this physical constraint, the standard LP reference antenna was swapped with a specialized, highly directional CP L-band antenna. The performance of both architectures was evaluated across SINR, end-to-end latency, and network reliability using 200-Byte and 500-Byte UDP payloads.

Signal-to-Interference-plus-Noise Ratio

The fundamental metric dictating the performance of the cellular modem is the physical signal quality. The empirical data reveals a drastic disparity in the RF energy captured by the two antenna configurations.

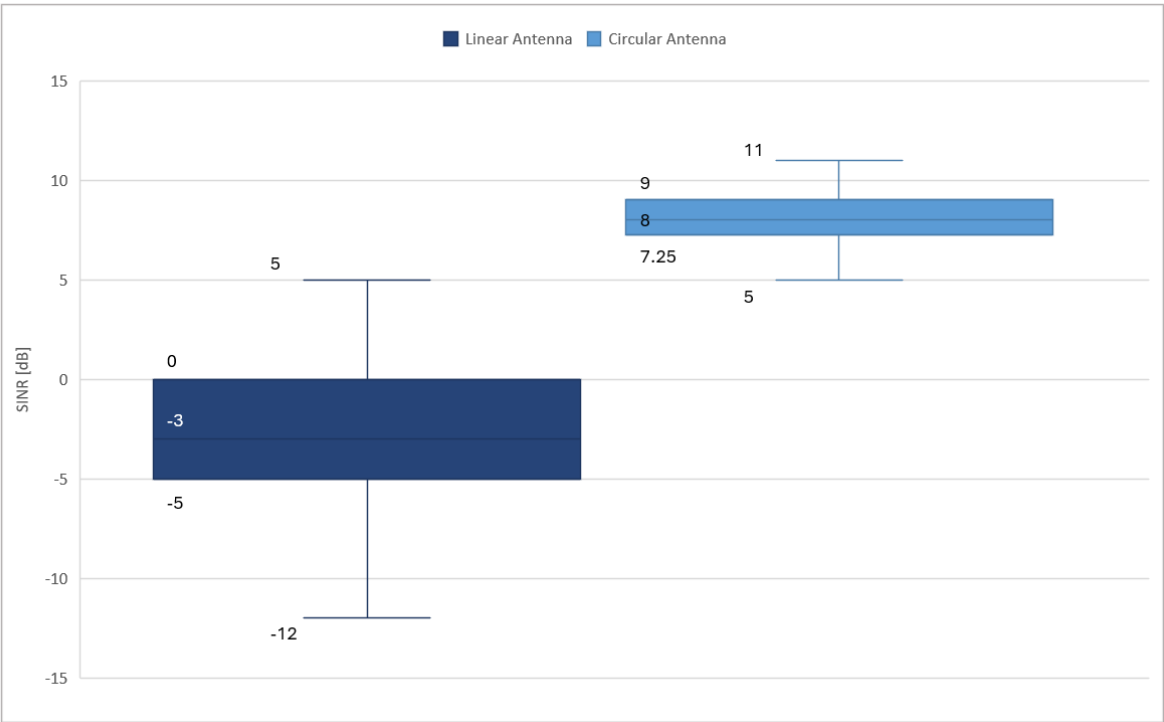


Figure 19: SINR Performance of Directive Circular vs. Omnidirectional Linear Antennas in NB-NTN

Operating at the absolute threshold of the link budget, the standard linear antenna recorded a median SINR of -3 dB, with its lower whisker dropping to a severe -12 dB. In stark contrast, employing a directive circular antenna that matches the satellite downlink's polarization yielded an immediate and massive improvement, raising the median SINR to 8 dB and reaching an upper whisker of 11 dB. This represents an aggregate signal quality improvement of 10 to 11 dB. Ultimately, this gain eliminates the theoretical 3 dB

polarization mismatch penalty and leverages the specialized antenna's higher directional gain, successfully pulling the signal well above the noise floor.

Impact on Transmission Latency and Systemic Jitter

Having a higher SINR in radio networks results in more stable connections and data transfer. Consequently, the 10 dB physical improvement observed with the circular antenna transformed the measurement results of latency and jitter.

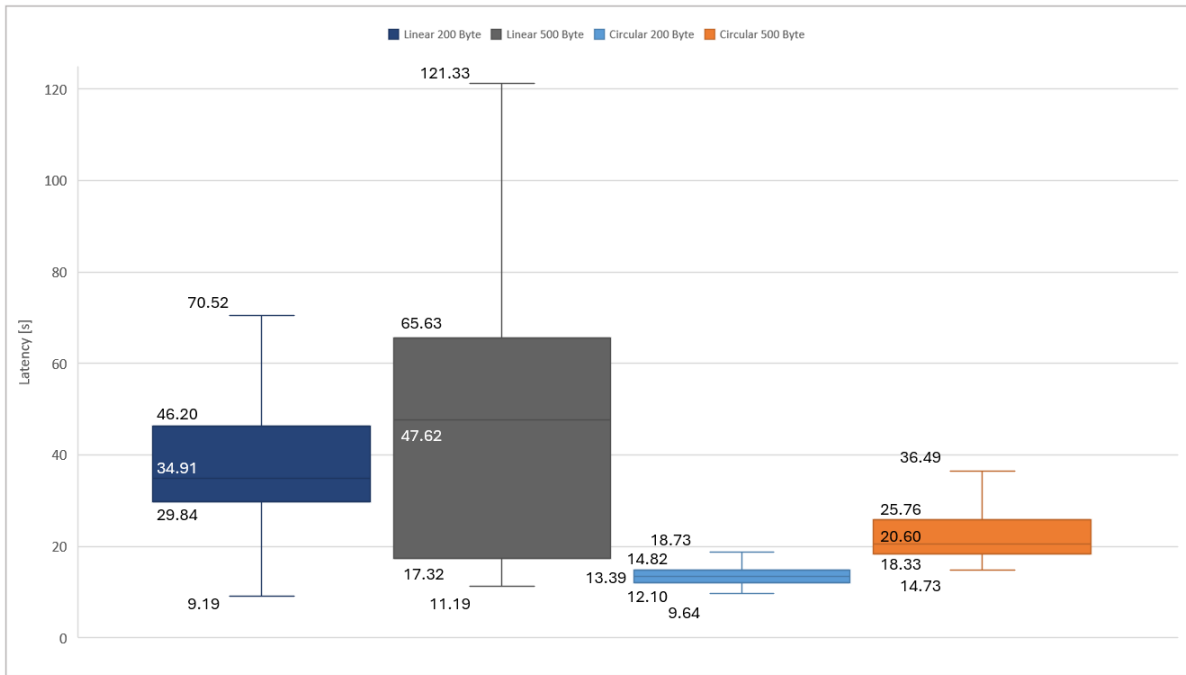


Figure 20: Latency Performance of Directive Circular vs. Omnidirectional Linear Antennas in NB-NTN

200-Byte Payload Latency:

Using the linear antenna, the median transmission latency was 34.91 s. Furthermore, the connection suffered from extreme jitter, with the upper quartile reaching 46.20 s and the upper whisker stretching to a severe 70.52 s.

The circular antenna drastically compressed this timeline. The median latency dropped to 13.39 s indicating a 61% reduction. More importantly, the systemic jitter was virtually eliminated: the upper whisker peaked at only 18.73 s.

500-Byte Payload Latency:

For larger payloads, the linear antenna struggled significantly, yielding a median latency of 47.62 s and extreme worst-case delays of 121.33 s.

The circular antenna handled the 500-Byte payload with high efficiency, recording a median latency of 20.60 s and capping the maximum delay at 36.49 s.

The data proves that the specialized antenna systems not only cuts the median transmission delay by more than half but also completely neutralizes the extreme two-minute latency spikes that plague the standard off-the-shelf hardware.

Network Reliability

To quantify transmission reliability, the message drop rates were analyzed across both the LP and CP antenna configurations using 200-Byte and 500-Byte payloads:

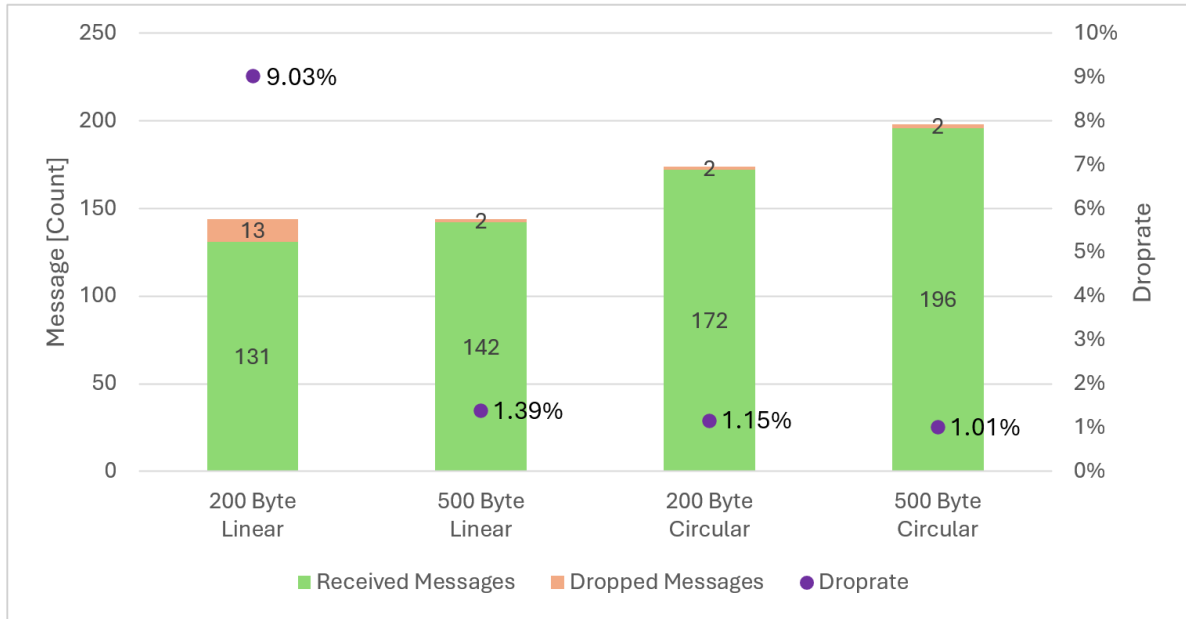


Figure 21: Received vs. Dropped Messages by Antenna and Message Size

Linear Antenna: The LP hardware exhibited high temporal variability and vulnerability to packet loss. During the 200-Byte baseline campaign, it dropped 13 out of 144 messages, resulting in a 9.03% drop rate. While it performed better during the 500-Byte campaign, the inconsistency makes the standard antenna unpredictable for critical deployments.

Circular Antenna: The robust SINR margin provided by the CP antenna stabilized the connection entirely. For both payload sizes, the drop rate remained strictly contained at the baseline threshold. The 200-Byte campaign dropped 2 out of 174 messages and the 500-Byte campaign dropped 2 out of 198 messages.

While these static rooftop measurements prove the undeniable superiority of the CP antenna under optimal Line-of-Sight conditions, its performance in dynamic, reflection-heavy environments will be evaluated in the following chapter.

4.5 Mobility & Automotive Performance

To evaluate the operational viability of GEO NB-NTN for automotive logistics and asset tracking (RQ2), the measurement node was subjected to dynamic mobility testing. As detailed in the methodology, the device was deployed along a standardized test route traversing rural highways, suburban transit corridors, and dense urban canyons in Leipzig. The analysis evaluates two primary variables: the impact of high-speed velocity on the radio link, and the effect of structural RF influence across different in-vehicle placements using both Linear and Circular antennas.

4.5.1 Impact of Velocity

A fundamental concern in satellite communications is the Doppler effect which is the shift in signal frequency caused by the relative motion between the UE and the satellite. To determine if automotive speeds degrade the NB-NTN link, 148 distinct transmission data points were captured at speeds ranging from stationary up to 160 km/h on the Autobahn.

A linear regression analysis was applied to the dataset to approximate the correlation between vehicle speed v (measured in km/h) and network performance.

For transmission latency (measured in milliseconds), the resulting trendline is:

$$Latency(v) = 45.493 * v + 29435$$

For physical signal quality SINR (measured in dB), the trendline is:

$$SINR(v) = 0.0224 * v + 3.8123$$

The mathematical slopes of these functions are statistically negligible. A slope of 0.0224 for SINR indicates that traveling at 100 km/h alters the signal quality by a mere 2.2 dB, well within the margin of standard temporal variance. Similarly, the latency trendline shows that speed has virtually no practical impact on transmission delays. This empirical data clearly demonstrates that the speed of up to 160 km/h of the UE does not negatively affect the radio performance of the NB-NTN link.

4.5.2 Antenna Polarization and Environmental Attenuation

Because the empirical data demonstrates that vehicle speed is not a primary degrading factor, connection failures during the dynamic driving tests must be attributed to a combination of physical obstruction and multipath fading. To systematically evaluate this, the device was rotated through three distinct placements within the vehicle, each representing a progressively harsher RF environment:

1. **Roof Mount**
2. **Behind Windscreen**
3. **Rear Bench**

Throughout the standardized driving route, the campaigns alternated between the standard LP antenna and the highly directional CP antenna. The resulting message drop rates were categorized across the three geographic environments: Rural, Suburban and Urban.

Table 7: Drop Rates by Antenna Type, Placement and Environment

Placement	Antenna	Rural	Suburban	Urban
Roof Mount	Linear	14.3% (1/7)	16.7% (1/6)	0.0% (0/6)
	Circular	0.0% (0/10)	0.0% (0/6)	40.0% (2/5)
Windscreen	Linear	12.5% (1/8)	37.5% (3/8)	16.7% (1/6)
	Circular	0.0% (0/8)	66.7% (4/6)	57.1% (4/7)
Rear Bench	Linear	0.0% (0/9)	0.0% (0/8)	25.0% (1/4)
	Circular	100.0% (9/9)	100.0% (8/8)	100.0% (4/4)

4.5.3 Data Analysis: Polarization Reversal Phenomenon

The empirical data from these mobility campaigns perfectly validates the theoretical physics of electromagnetic reflection as established in Section 2.4.5 and exposes a critical architectural trade-off for NTN IoT deployments. The performance of the link is not universally dictated by antenna gain, but rather by how the chosen polarization interacts with multipath scattering.

The Circular Antenna

The CP antenna is mathematically optimized for the direct satellite downlink. When provided with an unobstructed LOS, such as on the Roof Mount in open Rural and Suburban areas, it performs flawlessly, achieving a perfect 0.0% drop rate across 16 consecutive transmissions.

However, as soon as the physical environment introduces obstructions, forcing the antenna to rely on reflections, its performance collapses.

Urban Shadowing: In the Urban environment on the roof, tall buildings periodically block the direct line-of-sight. The CP drop rate immediately spiked to 40.0%.

Suboptimal Positioning: When moved behind the windscreen, the CP antenna struggled heavily in areas requiring signal penetration, dropping 66.7% in the suburbs and 57.1% in the urban core.

The Faraday Cage: Most severely, when placed on the Rear Bench deep inside the vehicle chassis, the CP antenna suffered a catastrophic 100.0% failure rate across all three environments. The severity of this attenuation was so absolute that the modem was completely unable to establish an initial network attach, continuously failing synchronization procedures for over an hour of active attempts.

This absolute failure on the rear bench is driven by the polarization reversal phenomenon. As the RHCP satellite signal enters the vehicle windows and bounces off the metallic interior of the doors and roof, the phase of the wave inverts, transforming it into a LHCP wave. The specialized CP antenna possesses high cross-polarization isolation, it is physically designed to reject LHCP waves. Consequently, the antenna treats the reflected signals inside the cabin as noise, effectively blinding the modem to the only available RF energy.

The Linear Antenna

Conversely, the standard terrestrial linear antenna demonstrated exceptional resilience in these exact reflection-heavy environments. Because a linear antenna oscillates on a single vector plane, it does not discriminate against the handedness of circular waves. While it inherently loses 3 dB of the direct signal, it is capable of constructively harvesting signals that have scattered, depolarized, or reversed their handedness upon reflection.

This theoretical advantage is proven definitively in the dataset:

Urban Reflection Harvesting: On the Roof Mount in the dense Urban core, the linear antenna achieved a perfect 0.0% drop rate, easily outperforming the CP antenna's 40% failure rate. When the direct signal was blocked by a building, the LP antenna successfully received the satellite waves bouncing off the surrounding concrete facades.

In-Cabin Resilience: Most impressively, when placed on the Rear Bench, the linear antenna maintained a flawless 0.0% drop rate in both Rural and Suburban environments. Even in the deep Urban core, where it faced the compounded attenuation of both the vehicle chassis and city buildings, it dropped only a single message. By simply gathering the depolarized, reflected waves scattering around the interior of the cabin, the standard LTE antenna kept the NTN link alive where the specialized satellite antenna failed entirely.

5. Discussion

The empirical measurement campaigns detailed in Chapter 4 provide a definitive, quantitative baseline for the performance of 3GPP Release 17 NB-NTN over GEO satellites. Moving beyond the raw data, this chapter synthesizes the findings to directly address the core problem statement of this thesis. Section 5.1 interprets these results through the lens of systems engineering, extracting concrete architectural guidelines and practical implications for IoT device manufacturers developing hybrid TN/NTN solutions. Subsequently, Section 5.2 objectively evaluates the inherent operational boundaries of GEO-based NB-NTN, defining the specific use cases where this technology offers the highest commercial and technical value.

5.1 Interpretation of Results and Implications for IoT Device Manufacturers

Designing IoT solutions to operate seamlessly across both conventional Terrestrial Networks and Non-Terrestrial Networks requires good optimization in hardware and software design. The empirical data indicates that directly porting standard terrestrial configurations to a GEO satellite link yields a functional prototype but often results in sub-optimal power efficiency and variable connection stability. To maximize commercial viability and harness the full potential of global satellite coverage, IoT device manufacturers should adopt an "Efficiency-First" engineering paradigm, utilizing the specific physical characteristics of the NTN link to guide the End-to-End architecture.

5.1.1 Redefining the Application Stack

A primary software-level opportunity for manufacturers lies in the strategic transition toward connectionless protocols for constrained devices. As established in Section 4.2, relying on TCP to guarantee message delivery over NTN introduces significant electrical and functional inefficiencies.

TCP was fundamentally designed for low-latency wired and stable cellular networks. The measurement data revealed that TCP experienced a 7.32% message drop rate over NTN, performing comparably to unacknowledged UDP (9.03%). This behavior is a direct consequence of the very high baseline latency of the GEO link. Because the satellite Round-Trip Time consistently exceeds standard TCP retransmission timeout thresholds, the protocol frequently triggers spurious retransmissions. This prematurely reduces the congestion window and can force the modem to drop the socket entirely. Furthermore, the mandatory three-way handshake and subsequent teardown, which is necessary to allow the device to return to power saving mode, generated a substantial 147% protocol overhead for a standard 200-Byte payload.

Conversely, the data identifies the CoAP, operating over UDP, as a highly effective standard for NTN deployments. By shifting the reliability mechanism from the transport layer to the application layer, CoAP eliminates the handshake while completely neutralizing the baseline physical packet loss of the satellite link, achieving a flawless 0.00% drop rate for 200-Byte payloads. Manufacturers can highly optimize their firmware by utilizing CoAP with confirmable messages, specifically tuning the acknowledgment timeout variables to account for a very high delay. This simple architectural adjustment prevents premature backoff failures and guarantees reliable data delivery over satellite.

5.1.2 Payload Sizing

In terrestrial cellular architectures, developers frequently aggregate sensor data into larger, monolithic payloads to maximize protocol efficiency. However, the empirical data from Section 4.1.4 indicates that this strategy introduces non-linear latency scaling over GEO NTN.

While scaling a payload from 200 Bytes to 500 Bytes on the TN link added only 376 milliseconds of latency, the exact same scaling on the NTN link increased the median latency by an additional 12.7 seconds and expanded the worst-case systemic jitter to over 120 seconds. Because GEO NTN disables MAC-layer HARQ, the network relies on blind physical layer repetitions. Larger transport blocks inherently carry a higher probability that a localized signal fade will corrupt a frame during the extended transmission window, requiring the Radio Link Control (RLC) layer to retransmit larger segments of data.

This physical reality presents a strong case for integrating Edge Computing into IoT firmware. Firmware architects should design systems to process, compress, and filter sensor data locally on the device. By transmitting only state changes, compressed binary schemas, or critical alerts rather than verbose IT formats like JSON or XML, devices can maintain payloads under 200 Bytes. This optimization ensures fast, highly predictable transmission times and maximizes the energy efficiency of the satellite link.

5.1.3 Hardware Architecture and Environmental Adaptability

The mobility and antenna performance data provide nuanced insights into hardware requirements for multi-orbit IoT devices. The assumption that specialized satellite antennas are universally required is incomplete. Hardware selection must be context-aware and strictly tailored to the target deployment environment.

Manufacturers can leverage these findings to optimize their hardware designs:

Unobstructed Line-of-Sight: If the device will be deployed with a clear view of the sky for example in maritime monitoring or agricultural sensors, circularly polarized antennas drastically elevate device efficiency. The static data proved that matching the satellite's CP downlink and having high directivity improves the SINR by over 10 dB, cuts latency by 61%, and stabilizes the drop rate to ~1%.

Non-Line-of-Sight and Multipath Environments: Conversely, if the device operates in environments dominated by multipath scattering, such as deep urban canyons or inside vehicle cabins, standard linearly polarized and omnidirectional antennas offer a distinct advantage. As demonstrated in the dynamic driving tests, LP antennas constructively harvest depolarized reflections, maintaining connectivity where CP antennas filter out reflected signals.

Meteorological Resilience: Regarding atmospheric conditions, systems engineers can dimension NTN link budgets without severe fade margins for rain or cloud cover, as the L-band and S-band spectrum demonstrated high resilience to weather.

This data reveals a significant opportunity for hybrid devices designed to operate globally across varied environments. Beyond multipath resilience, standard linearly polarized antennas offer the crucial advantage of broad frequency support across standard terrestrial network cellular bands. This natively enables a converged, single-antenna design capable of seamless TN and NTN roaming. Conversely, because specialized circularly polarized antennas are not optimal for the use with terrestrial linear bands, manufacturers utilizing CP hardware are forced to implement a more complex, dual-antenna architecture to maintain TN capabilities.

For highly cost-constrained or size-limited devices operating in unpredictable environments, a high-efficiency LP antenna provides the ultimate resilient fail-safe. It is uniquely capable of leveraging signal reflections to overcome urban shadowing and vehicle chassis attenuation, while simultaneously supporting NTN roaming without the added cost and footprint of a secondary RF trace.

5.1.4 Optimizing the Power Budget

The power consumption measurements highlight energy management as the primary operational consideration for NTN deployments. The empirical data reveals severe energy multipliers: attaching to the NTN link requires roughly the same amount of energy as a bad-case terrestrial attach. However, active payload transmission is vastly more demanding, requiring over 30 times more energy than an optimal TN link, and approximately 10 times more energy than an edge-of-cell TN deployment.

To unlock multi-year battery lifespans despite these energy consumption adds, IoT manufacturers should implement firmware-level orchestration to mitigate these costs:

Minimize Network Attaches: Because a single satellite attach procedure consumes thousands of milliamperes-seconds, the device architecture should be designed to avoid cold-boot attach procedures whenever feasible. Firmware should utilize eDRX and PSM to maintain the RRC context over long durations. With this “attach once, stay attached” procedure, the device bypasses the high-energy synchronization phases for subsequent messages over extended periods.

GNSS Pre-Provisioning: On top of the consumed NTN attach energy comes the GNSS requirement. For static deployments, for example in agricultural sensors or infrastructure monitoring use cases, manufacturers can pre-provision the device's geographic coordinates into the MCU firmware during installation. This allows the module to bypass the GNSS cold-start entirely, immediately initiating the cellular handshake and saving a significant portion of the initial synchronization energy.

Transmission Throttling: Given that active payload transmission over satellite draws 10 to 30 times more energy than terrestrial equivalents, manufacturers should shift away from frequent reporting. For example, firmware can be event-driven, transmitting data only when specific sensor thresholds are breached or utilizing edge-computing to aggregate daily data into a single, highly compressed transmission.

5.2 Operational Boundaries of GEO-Based NB-NTN

While 3GPP Release 17 successfully bridges the global coverage gap, the empirical data outlines specific operational boundaries inherent to Geostationary NB-NTN. These boundaries are dictated by orbital mechanics, the speed of light, and 3GPP power-saving features. Understanding these parameters allows network architects to accurately align the technology with highly compatible use cases, ensuring successful commercial deployments.

5.2.1 Latency and Half-Duplex Characteristics

A fundamental characteristic of GEO NB-NTN is absolute propagation latency. The empirical baseline established in this thesis demonstrates 10 to 40 seconds of transmission delay under optimal conditions, scaling as payloads increase or SINR drops.

This positions GEO NB-NTN as a good asynchronous messaging system, rather than a real-time control link.

Control Loop Design: Use cases requiring closed-loop remote control are unsuited for a link where a single command-and-acknowledgment cycle may require 60 to 120 seconds. Instead, NTN excels in delay-tolerant reporting, such as daily supply chain updates or environmental telemetry.

Half-Duplex Operations: NB-IoT operates in half-duplex frequency division duplexing, meaning the modem cannot transmit and receive simultaneously. During the extended active transmission phase, the device does not monitor the downlink. Downlink messages dispatched from the cloud are queued at the Evolved Packet Core and delivered once the device transitions to the Rx listening phase. Cloud application logic must be designed to accommodate this asynchronous, queue-based delivery system.

5.2.2 Return-Link Energy Dynamics

While the GEO satellite possesses substantial power to broadcast a downlink signal across the 35,786-kilometer distance, the constrained IoT device faces strict constraints on the uplink. An IoT module operating on battery power is physically limited to a transmit power of 23 dBm (Power Class 3).

To successfully close the ~190 dB Free Space Path Loss, the network utilizes CE repetitions, resulting in continuous maximum-power current draws exceeding 160 mA over extended periods.

This energy dynamic informs the optimal reporting frequency for NTN devices. While terrestrial "Massive IoT" allows for devices to report every few minutes, high-frequency reporting over satellite will accelerate hardware energy depletion. NTN provides the highest value when configured for infrequent, high-value data transmission, easily enabling multi-year lifespans on standard industrial battery packs.

5.2.3 Urban Shadowing

The mobility data (Section 4.5) clearly establishes that GEO NB-NTN is optimally deployed as a complementary technology to terrestrial infrastructure, ensuring continuous coverage as assets move out of TN range.

In the dense urban core, structural shadowing from buildings routinely obstructs the direct satellite downlink. While linear antennas successfully harvest reflections to maintain the link, overall performance is naturally lower than in rural, open-sky environments. Consequently, highly mobile assets frequently transitioning through urban canyons or tunnels should be programmed to prioritize the terrestrial network, utilizing the NTN link primarily when traversing open highways or remote regions to maximize energy efficiency.

5.2.4 Cloud-Side Integration and Protocol Translation

Because edge devices must utilize heavily compressed, connectionless protocols like CoAP or Lightweight Machine-to-Machine (LwM2M) to optimize the NTN link budget, an architectural shift is required within the Cloud Platform Domain.

Standard enterprise IT architectures frequently rely on verbose, text-based protocols like HTTPS or MQTT. Because transmitting these protocols is highly inefficient over satellite, cloud architecture must absorb the complexity of protocol translation. Network architects must deploy ingestion gateways designed to terminate CoAP/LwM2M traffic, unpack optimized binary schemas, and translate the data into RESTful APIs or JSON objects for enterprise consumption.

By moving this computational complexity off the edge device and into the cloud, enterprises can maintain the "plug-and-play" simplicity of their internal IT dashboards while ensuring their remote IoT hardware operates at peak RF efficiency.

6. Conclusion and Outlook

6.1 Conclusion

The standardization of Non-Terrestrial Networks within 3GPP Release 17 represents a monumental shift in cellular telecommunications, theoretically eliminating the 85% global coverage gap that has historically restricted IoT deployments. However, the empirical measurement campaigns conducted in this thesis demonstrate that bridging this gap using Geostationary Earth Orbit satellites introduces profound physical and architectural challenges.

In a direct baseline comparison, terrestrial NB-IoT remains vastly superior in every operational metric: it is highly reliable, functionally superior, and extremely power-efficient. In contrast, GEO NB-NTN is a strictly constrained, asynchronous messaging medium. It fundamentally relies on massive physical layer repetitions and high baseband energy to overcome a ~ 190 dB Free Space Path Loss and a 500 ms round-trip propagation delay. Consequently, while NTN successfully enables global connectivity using standard commercial hardware, treating it as a direct equivalent to a terrestrial network can lead to fatal systemic failures.

To provide actionable parameters for IoT systems engineers, this research systematically addressed four primary Research Questions:

RQ1: Protocol Efficiency and Reliability

The empirical data confirms that transport and application-layer protocols dictate the viability of the satellite link. The connection-oriented TCP is fundamentally incompatible with GEO NTN, introducing an exorbitant 147% protocol overhead for a 200-Byte payload while still suffering a $\sim 7.3\%$ packet drop rate due to latency-induced timeouts. Conversely, CoAP operating over UDP proved to be a good solution. By managing reliability at the application layer, CoAP neutralized the physical packet loss of the satellite link while maintaining a scalable overhead profile. Furthermore, the data proved that payload sizes must be strictly minimized: scaling payloads from 200 to 500 Bytes on NTN induced exponential latency penalties, expanding maximum transmission delays from 30 seconds to over 120 seconds.

RQ2: Environmental and Operational Impact

Dynamic mobility testing on the German Autobahn definitively proved that vehicle velocity (up to 160 km/h) and standard Doppler shifts do not degrade the NB-NTN link, confirming the effectiveness of 3GPP Rel-17 pre-compensation mechanisms. Meteorologically, the L-band satellite link is highly resilient to rain and cloud cover. However, solid precipitation (snow accumulation) acting as a dielectric obstruction was shown to severely compress and elevate median transmission latency by approximately 8.5 seconds. Ultimately, the most significant operational disruptor is urban shadowing. Dense city environments physically block the direct satellite line-of-sight, mandating that highly mobile assets prioritize terrestrial networks and utilize NTN strictly as a rural or maritime fallback.

RQ3: Hardware Architecture and Antenna Polarization

The research explicitly disproves the assumption that specialized satellite antennas are universally required. Hardware selection is strictly context-dependent. Under optimal line-of-sight conditions, a specialized directive circularly polarized antenna is vastly superior, improving the physical SINR by over 10 dB, cutting latency by 61%, and stabilizing the drop rate. However, due to the polarization reversal phenomenon, CP antennas reject reflected signals, suffering a 100% failure rate when placed inside a vehicle cabin. For environments dominated by multipath scattering and non-line-of-sight conditions,

standard linearly polarized antennas are advantageous, as they can constructively harvest depolarized signal reflections to maintain the link.

RQ4: Energy Profiling

The quantitative power profiling establishes that GEO NB-NTN strains the energy budgets of constrained LPWAN devices. While the initial network attach procedure for the satellite link is about 100 times more energy-intensive compared to optimal terrestrial NB-IoT, it is equally electrically expensive as a terrestrial network can be in harsh, edge-of-cell scenarios. However, actively transmitting a standard payload over NTN requires about 10 to 20 times more energy than operating at the absolute cell-edge of a terrestrial network.

In summary, 3GPP Rel-17 GEO NB-NTN is not a high-throughput, real-time network, but rather a highly capable, delay-tolerant telemetry solution. When engineered with a "Optimization-First" paradigm, utilizing CoAP, compressed payloads, and context-aware antenna architectures, it successfully transforms remote, disconnected environments into viable domains for continuous IoT data acquisition.

6.2 Outlook

While the empirical baselines established in this thesis define the operational boundaries of current geostationary deployments, the landscape of non-terrestrial networks is evolving rapidly. Future iterations of both satellite infrastructure and RF hardware promise to resolve many of the most severe constraints identified in this research.

The most significant upcoming architectural shift is the commercialization of NB-IoT over LEO satellite constellations. Operating at altitudes between 500 and 1,200 kilometers, LEO satellites fundamentally alter the physics of the NTN link. The reduced orbital distance will practically eliminate the 500 ms propagation delay that currently breaks terrestrial TCP timers and restricts the network to delay-tolerant applications.

The most significant upcoming architectural shift is the commercialization of NB-IoT over Low Earth Orbit (LEO) satellite constellations [8]. Operating at altitudes between 500 and 1,200 kilometers, LEO satellites fundamentally alter the physics of the NTN link [2]. The reduced orbital distance will practically eliminate the 500 ms propagation delay that currently breaks terrestrial TCP timers and restricts the network to delay-tolerant applications, compressing the physical RTT to approximately 20 to 40 milliseconds [26].

More critically, LEO constellations provide an effective 15 to 20 dB higher link budget compared to GEO satellites [8], [2]. This massive gain in signal strength will allow the network to significantly reduce the extreme coverage enhancement repetitions currently required to close the link. Consequently, future LEO NB-NTN deployments will drastically reduce the active radio time of the IoT device, mitigating the extreme power consumption profiles measured in this thesis and extending battery lifespans to closer match terrestrial expectations [15].

Furthermore, this 15 to 20 dB link budget surplus fundamentally expands the viable deployment environments for satellite IoT. While GEO networks favor unobstructed line-of-sight, the LEO link margin is mathematically sufficient to overcome "light indoor" attenuation, such as the 4 to 8 dB penetration loss typical of wooden shelters, agricultural sheds, or fiberglass enclosures [27]. This allows NTN to support semi-indoor assets [28].

At the hardware level, the polarization dilemma identified in RQ3, which forces engineers to choose between high-efficiency CP antennas for line-of-sight and resilient LP antennas for multipath environments, presents a significant opportunity for RF innovation. One highly promising approach for future IoT hardware architectures is the integration of dual-hand circular antennas [29]. By utilizing

hardware capable of processing both the direct Right-Hand satellite downlink and the Left-Hand reflections bouncing off urban infrastructure or vehicle chassis, systems engineers could theoretically eliminate the polarization mismatch loss without sacrificing multipath resilience [30].

If successfully miniaturized and made cost-effective for constrained LPWAN modules, such dual-polarized architecture would provide a highly viable alternative to standard linear antennas. This approach could enable highly versatile "deploy-anywhere" hardware, allowing a single tracking module to transition seamlessly from the open ocean to the inside of a dense urban environment while maintaining optimal link budgets across both scenarios.

Beyond physical and hardware optimizations, the empirical validation of CoAP over satellite links opens critical pathways for advanced application-layer evolution. Because this research proves that CoAP operates with high efficiency and 100% reliability over constrained GEO NTN links, it establishes the foundational viability for LwM2M in satellite architectures.

LwM2M natively builds upon the CoAP standard to facilitate complex, standardized operations such as Firmware Over-The-Air (FOTA) updates, remote provisioning, and extensive lifecycle device management. Future work and measurement campaigns should focus on empirically validating LwM2M over NTN, specifically profiling the energy consumption and block-wise transfer reliability of FOTA updates, to prove that extensive remote device management is practically feasible for global NTN deployments.

Ultimately, 3GPP Release 17 represents merely the foundational layer of hybrid cellular connectivity. As LEO link budgets mature and specialized hardware evolves the performance gap between terrestrial and satellite IoT will continue to close, paving the way for truly ubiquitous, high-efficiency global sensor networks.

Bibliography

- [1] J. A. Fraire, P. Madoery, R. Burleigh, C. G. Galarza, and J. Finochietto, "Direct-to-satellite IoT—A survey of the state of the art and future research perspectives," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4, pp. 2858-2886, 2020.
- [2] 3GPP, "Study on Narrow-Band Internet of Things (NB-IoT) / enhanced Machine Type Communication (eMTC) support for Non-Terrestrial Networks (NTN) (Release 17)," 3rd Generation Partnership Project (3GPP), Technical Report (TR) 36.763, 2021.
- [3] C. A. Balanis, *Antenna Theory: Analysis and Design*, 4th ed. Hoboken, NJ, USA: John Wiley & Sons, 2016.
- [4] P. Thulasiraman *et al.*, "Performance evaluation of MQTT and CoAP over satellite IoT," in *Proc. IEEE International Conference on Communications (ICC)*, 2022.
- [5] M. Allman, D. Glover, and L. Sanchez, "Enhancing TCP over satellite channels using standard mechanisms," IETF, RFC 2488, 1999.
- [6] D. Richter, "Dear IoT community, I have a question for you! [LinkedIn post]," LinkedIn, Nov. 2025. [Online]. Available: https://www.linkedin.com/posts/dominic-richter_%F0%9D%90%83%F0%9D%90%9E%F0%9D%90%9A%F0%9D%90%AB-%F0%9D%90%88%F0%9D%90%A8%F0%9D%90%93-%F0%9D%90%9C%F0%9D%90%A8%F0%9D%90%A6%F0%9D%90%A6%F0%9D%90%AE%F0%9D%90%A7%F0%9D%90%A2%F0%9D%90%AD%F0%9D%90%B2-%F0%9D%90%88-%F0%9D%90%A1-activity-7390729463925485568-4HHx/ [Accessed: Feb. 22, 2026].
- [7] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation (Release 13)," 3GPP, Technical Specification (TS) 36.211.
- [8] X. Lin *et al.*, "5G non-terrestrial networks: Mobility, link budget, and payload design," *IEEE Communications Magazine*, vol. 59, no. 3, pp. 36-43, 2021.
- [9] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification (Release 17)," 3GPP, TS 36.321.
- [10] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification (Release 17)," 3GPP, TS 36.331.
- [11] GSMA, "IoT device connection efficiency guidelines," GSM Association, Official Document TS.34, Version 8.0.
- [12] B. Martinez *et al.*, "A roaming architecture for low-power wide-area networks," in *Proc. ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2019.
- [13] R. Goyal *et al.*, "Performance of TCP over satellite links," *Computer Networks*, 1999.
- [14] Z. Shelby, K. Hartke, and C. Bormann, "The Constrained Application Protocol (CoAP)," IETF, RFC 7252, 2014.
- [15] L. Zhen *et al.*, "Energy-efficient transmission for NB-IoT over LEO satellite networks," *IEEE Wireless Communications Letters*, vol. 10, no. 5, pp. 933-937, 2021.
- [16] O. Liberg *et al.*, *Cellular Internet of Things: From Massive Deployments to 5G and Beyond*, 2nd ed. Academic Press, 2020.

- [17] J. D. Kraus and R. J. Marhefka, *Antennas for All Applications*, 3rd ed. McGraw-Hill, 2002.
- [18] S. R. Saunders and A. Aragón-Zavala, *Antennas and Propagation for Wireless Communication Systems*, 2nd ed. John Wiley & Sons, 2007.
- [19] IEEE, "IEEE Standard Definitions of Terms for Antennas," IEEE Std 145-2013, 2013.
- [20] Y. Suzuki *et al.*, "Multipath fading characteristics for land mobile satellite systems," *IEEE Transactions on Antennas and Propagation*, 1996.
- [21] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of Things: A survey on enabling technologies, protocols, and applications," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2347-2376, 2015.
- [22] 3GPP, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access," 3GPP, TS 23.401.
- [23] GSMA, "Mobile IoT deployment guide," GSM Association, Official Document, 2023.
- [24] H. Tschofenig *et al.*, "Architectural considerations in smart object networking," IETF, RFC 7452, 2015.
- [25] P. Sethi and S. R. Sarangi, "Internet of Things: Architectures, protocols, and applications," *Journal of Electrical and Computer Engineering*, 2017.
- [26] O. Kodheli *et al.*, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70-109, 2021.
- [27] 3GPP, "Study on New Radio (NR) to support Non-Terrestrial Networks," 3GPP, TR 38.811.
- [28] M. G. Kim and H. S. Jo, "Performance analysis of NB-IoT uplink in low earth orbit non-terrestrial networks," *Sensors*, vol. 22, no. 18, p. 7097, 2022.
- [29] S. Gao, Q. Luo, and F. Zhu, *Circularly Polarized Antennas*. John Wiley & Sons, 2014.
- [30] Z. Zhang *et al.*, "Dual-band dual-circularly polarized antenna for GNSS and satellite communications," *IEEE Transactions on Antennas and Propagation*, vol. 69, no. 8, pp. 4335-4344, 2021.

Appendix

Due to the extensive volume of empirical data collected during the measurement campaigns, the raw data, power profiling logs, and statistical analysis sheets are provided as a digital supplement attached to this thesis (Appendix.zip).

Appendix A: Measurement Data

The digital archive is organized into the following directory structure, containing the Excel files of the raw data and the analysis:

/Analysis/

 /Mobility/

 MASTER_Mobility.xlsx

(Includes mobility campaigns)

 /NTN/

 MASTER_NTN.xlsx

(Includes static campaigns)

 /TN/

 MASTER_TN.xlsx

(Includes static campaigns)

 /Power/

 Power_Consumption.xlsx

/Raw_Data/

 /Main_Measurements/

Contains the individual .xlsx files for all static campaigns

 /Mobility_Measurements/

Contains the individual .xlsx files for all dynamic automotive driving tests

 /Power_Measurements/

Contains the individual .xlsx files for all power measurement tests

Appendix B: Device Source Code

The custom embedded firmware engineered for the Nordic Semiconductor nRF9151 module is provided in the digital supplement under the /Appendix B Device Source Code/ directory.

Appendix C: Server Source Code

The custom receiving server architecture developed to ingest, parse, and analyze the IoT telemetry is provided in the digital supplement under the /Appendix C Server Source Code/ directory.