

Article

Five-Level Adaptive ReportInterval Selection Using a Hysteresis Mechanism for Low-Mobility Devices in 5G NR Networks

Dilmurod Davronbekov ¹, Nurmukhamed Shaudenbaev ¹, Muhammad Sadiq ^{2,*}, Cheng Wen ³, Hua Zheng ⁴ and Kuanishbay Sadatdiynov ⁵

¹ Department of Technologies of Mobile Communications, Tashkent University of Information Technologies Named After Muhammad al-Khwarizmi, Tashkent 100200, Uzbekistan; d.davronbekov@tuit.uz (D.D.); n.shaudenbaev@tuit.uz (N.S.)

² School of International Exchange and Cooperation, Shenzhen University of Information Technology, Shenzhen 518172, China

³ Guangzhou Institute of Technology, Xidian University, Guangzhou 510555, China; wencheng@xidian.edu.cn

⁴ School of Software and Artificial Intelligence, Guangzhou University of Software, Guangzhou 510900, China; zhengh@mail.gzus.edu.cn

⁵ Department of Data Transmission Networks and Systems, Nukus State Technical University, Nukus 230101, Uzbekistan; k.sadatdiynov@nukusstu.uz

* Correspondence: sadiq@suit-sz.edu.cn

Abstract

The expansion of Internet-of-Things (IoT) deployments in 5G New Radio (NR) networks has made periodic measurement reporting a growing burden for low-mobility devices, which benefit little from frequent updates yet must report as often as highly mobile ones. At present, User Equipment (UE) transmits MeasurementReport messages at a fixed ReportInterval—typically 240 ms—regardless of mobility. This continuous transmission needlessly depletes UE battery energy and consumes critical uplink signaling capacity. This paper proposes a five-level adaptive ReportInterval selection scheme driven by the statistical properties of Reference Signal Received Power (RSRP) and Signal-to-Interference-plus-Noise Ratio (SINR). A low-mobility criterion combines four statistical conditions—the variance and gradient of both RSRP and SINR—through a logical AND, while a two-stage hysteresis mechanism (a 3 dB margin and a 2 s holding timer) suppresses unnecessary level transitions. The scheme is slice-agnostic: By relying on observed signal statistics rather than network-slice labels, it serves low-mobility mMTC and stationary eMBB devices while leaving URLLC and high-mobility UEs at their standard configuration. In Monte Carlo simulations over the 3GPP TR 38.901 Urban Micro (UMi) channel model (200 UEs, 300 s, 100 iterations), the algorithm attains a classification accuracy of 91.32% and a sensitivity of 98.77%. Based on the DRX energy model, it yields an estimated 10.87% reduction in average UE power (from 28.15 to 25.09 mW) together with a 51.09% reduction in the network-wide MeasurementReport count. The hysteresis mechanism cuts level transitions by a factor of 31.33 (from 6852.7 to 218.7 per iteration), substantially lowering RRC reconfiguration signaling. Operating at O(n) complexity on the gNodeB and using only conventional MeasConfig signaling, the scheme requires no protocol additions or UE-side modifications, making it directly deployable on existing 3GPP Release 17 infrastructure as a gNB-side software update.

Keywords: 5G NR; ReportInterval; low-mobility criterion; RSRP; SINR; adaptive method; five-level selection; hysteresis; DRX; MeasurementReport; network slicing; IoT; RedCap



Academic Editor: Sarun Duangsuwan

Received: 24 May 2026

Revised: 11 July 2026

Accepted: 17 July 2026

Published: 4 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

5G New Radio (NR) supports a range of services, including high-throughput mobile broadband, ultra-reliable low-latency communication, and massive machine-type communication [1,2]. These services all share one feature, namely their reliance on periodic measurement reporting: Each User Equipment (UE) measures the Reference Signal Received Power (RSRP) and Signal-to-Interference-plus-Noise Ratio (SINR) of the serving cell and sends the findings to the gNodeB (gNB) in a MeasurementReport. These reports inform handover, scheduling, and link-adaptation decisions [3]. The reporting frequency is determined by the ReportInterval parameter, which the operator chooses from a discrete, standardized set specified in 3GPP TS 38.331 [4].

Currently, ReportInterval is statically configured with a single value, usually 240 ms, applied to all devices regardless of their mobility. When the majority of the connected devices were smartphones used by drivers or pedestrians, this consistent structure made sense. The deployment picture has shifted. Ericsson [5] has reported rapid growth of cellular IoT connections, and GSMA Intelligence [6] has estimated that a significant fraction of IoT devices are either static or low-mobility, such as factory sensors mounted on equipment, utility meters mounted on walls and environmental monitors fixed in the field. For such devices, sending a MeasurementReport every 240 ms is wasteful. Each transmission activates the radio, draws current through the power amplifier, and consumes uplink resources that might otherwise be used for active traffic [7,8]. The problem is even more severe for Reduced Capability (RedCap) devices [9,10], since their smaller batteries and reduced circuitry make each millijoule of energy saving important. Various approaches have been explored by researchers to resolve this problem, and thus it is important to understand the pros and cons of each approach.

The 3GPP RRM Relaxation mechanism [11] decreases the frequency of neighbor-cell measurements when the average RSRP exceeds a certain threshold. The concept is simple, and the implementation is lightweight. However, the design has three limitations. It does not consider SINR, so a device in a high-interference area, where RSRP is adequate but SINR is poor, may be incorrectly allowed to ease its operations. It only affects the measurements of neighboring cells and does not alter the MeasurementReport interval of the serving cell. However, it is not able to distinguish between a device with a constant RSRP and a device with an RSRP exhibiting large variations around the same mean, as it is based on the mean instead of the variance or the gradient. The technique only gives a yes/no relaxation decision, thus treating a robust factory sensor the same as a slightly drifting wearable.

Threshold-based approaches [3] work by analyzing the RSRP variance: If the variance is less than a predefined threshold, the device is classified as stationary, resulting in an increase in ReportInterval. This improves RRMRelaxation. However, the decision framework is still binary. The same intervals are used for a device with low variance and a device with moderate variance (though it is much less stable). Without a hysteresis mechanism, a device with variance near the threshold boundary switches between intervals across successive evaluation cycles, leading to RRCReconfiguration signals that increase the signaling overhead rather than decrease it [3,12].

Machine learning approaches—including reinforcement learning [13,14] and low mean square error estimators [15]—have shown promise in improving how networks adapt to changing conditions. That said, these methods typically rely on traffic-related inputs rather than signal-quality measurements, which limits their scope. Reinforcement-learning variants, in particular, come with notable drawbacks: Their computational demand grows quadratically, and they can take considerable time to converge. For devices with limited processing power, such as RedCap, these demands are not just inconvenient—they

are a real barrier to practical deployment [9]. Beyond DRX and reporting optimization, machine-learning methods have also been applied to adjacent 5G problems such as line-of-sight detection [16] and energy-oriented resource allocation [17], underscoring the broad applicability of data-driven techniques in this domain.

Several very recent studies confirm this data-driven trend: Yu et al. [18] employ a random-forest and temporal-convolutional-network ensemble to predict traffic and retime the C-DRX activation window; Kumar and Gautam [19] combine Q-learning with beam-forming to optimize DRX behavior in 5G NR-U; and Khalifa et al. [20] conserve IoT energy through UAV-mounted wake-up radios. Yet across the entire body of work reviewed above—from RRM Relaxation and threshold methods to reinforcement learning, DRX-cycle prediction, and aerial wake-up radios—none adapts the RRC-layer MeasurementReport interval itself, and none exploits the joint variance and gradient of RSRP and SINR for this purpose. The present work fills precisely this gap: It adapts ReportInterval directly from signal statistics already available at the gNB, requiring no traffic-layer input, no training or convergence phase, and no additional hardware.

Despite the variety of approaches reviewed, a common limitation emerges: None of the existing methods jointly exploit both the variance and gradient of RSRP and SINR to dynamically adjust the MeasurementReport interval. Similarly, the integration of a hysteresis mechanism to reduce unnecessary state transitions remains absent from the literature. Table 1 summarizes this comparison.

Table 1. Evaluation of current measurement-optimization methods against the proposed algorithm.

Criterion	RRM Relax [11]	Threshold [3]	Q-Learning [13]	MSE [15]	Proposed
Input parameters	mean RSRP	RSRP var.	traffic stats	traffic stats	RSRP + SINR var. & grad.
Number of levels	2	2	4	2	5
SINR used	No	No	Partial	No	Yes
Variance used	No	RSRP only	No	No	RSRP + SINR
Gradient used	No	No	No	No	RSRP + SINR
Hysteresis	No	No	No	No	Yes
Complexity	O(1)	O(1)	O(S A)	O(1)	O(n)
Standard-compliant	Yes	Partial	No	Partial	Yes

This study addresses that gap. Our contributions are summarized as follows. First, we establish a low-mobility criterion that combines four statistical conditions of RSRP and SINR, including their variance and gradient, via a logical AND operation, resulting in a sensitivity of 98.77%. Secondly, based on the low-mobility criterion, we develop a five-level ReportInterval assignment method (240–5120 ms) using a two-stage hysteresis mechanism, a 3 dB margin and a 2 s holding timer that reduces level transitions by a factor of 31.33. Third, the system is slice-agnostic; it operates on observed signal statistics rather than network-slice labels, thereby accommodating low-mobility devices across mMTC and stationary eMBB categories while automatically excluding URLLC and high-mobility UEs. Fourth, using Monte Carlo simulation based on the 3GPP TR 38.901 UMi channel model

(200 UEs, 300 s, 100 iterations), we show that the algorithm yields an estimated 10.87% reduction in UE energy and a 51.09% reduction in the number of MeasurementReport messages, with no degradation observed in the single-cell scenario considered in this study (inter-cell handover is outside the scope of this work and is identified as future work in Section 5.8). The complete approach operates with $O(n)$ complexity on the gNB and follows conventional MeasConfig signaling, without requiring protocol extensions or modifications on the UE side.

Quantitatively, the proposed scheme achieves the most favorable overall balance among the evaluated methods. As detailed in Section 4.6, it achieves a classification accuracy of 91.32%, exceeding the 3GPP RRM Relaxation baseline by 9.36 percentage points and the Binary Threshold baseline by 45.58 percentage points, while simultaneously attaining the highest energy saving (10.87%) and MeasurementReport reduction (51.09%) among the re-implemented baselines. A key reason for this advantage is the joint use of SINR together with RSRP. Methods that rely on RSRP alone—such as RRM Relaxation, which thresholds only the mean RSRP—cannot distinguish a device with a strong but interference-corrupted link from a genuinely stable one, and consequently reach a specificity of only 30.72%. By additionally exploiting the variance and gradient of SINR, the proposed criterion captures interference-driven instability that RSRP cannot reveal, more than doubling the specificity to 68.96%; the ablation study in Section 4.7 confirms that the SINR variance is one of the two most discriminative of the four conditions.

The remaining sections of this work are structured as follows. Section 2 describes the system model and the proposed algorithm. Section 3 describes the simulation configuration. Section 4 presents the findings. Section 5 examines the effects, the correlation with established 3GPP safety measures, and the limitations. Section 6 concludes this work.

2. System Model and Algorithm

Before describing the algorithm, we state the underlying optimization problem. For each low-mobility UE, the gNB selects the maximum ReportInterval that meets the minimum signal quality and stability requirements.

$$T_{*RI} = \operatorname{argmax}_{T_{RI}} ESR(T_{RI}), \text{ where } T_{RI} \text{ is an element of } T_{set} \quad (1)$$

subject to:

$$\mu_{RSRP} \geq RSRP_{min},$$

$$\sigma_{RSRP}^2 \leq \sigma_{max}^2,$$

$$\mu_{SINR} \geq SINR_{min},$$

$$\sigma_{SINR}^2 \leq \sigma_{max}^2,$$

$$T_{RI} \leq D_{max}.$$

The set T_{set} contains the standardized ReportInterval (T_{RI}) values specified in 3GPP TS 38.331 [4]. The remainder of Section 2 describes the elements that implement this optimization in practice.

The proposed mechanism runs at the gNB, utilizing data natively available in standard MeasurementReport messages. It contains three components: statistical parameter estimation, a low-mobility criterion (C_{low}), and a multi-tier ReportInterval selector with integrated hysteresis. Each element is described below.

2.1. Statistical Parameters of Signals

We examine a 5G NR network functioning in Standalone (SA) mode, wherein N UEs are linked to a single serving gNB. Each UE i measures $RSRP_{i(t)}$ and $SINR_{i(t)}$ at each sample interval Δt , using the Synchronization Signal Block (SSB) reference signal of the serving cell, in compliance with 3GPP TS 38.215 [21] and TS 38.214 [22]. Layer-3 filtering is implemented as described in 3GPP TS 38.331:

$$F_n = (1 - a)F_{n-1} + a M_n, \text{ where } a = \frac{1}{2^{\frac{k}{4}}} \quad (2)$$

where F_n represents the filtered measurement, M_n is the current sample, and k denotes the standard filter coefficient. Within an observation period of duration T_{window} with n samples, the UE calculates the mean, variance, and average gradient for both RSRP and SINR.

$$\mu_X = \left(\frac{1}{n}\right) * \Sigma(X_i) \quad (3)$$

$$\sigma_X^2 = \left(\frac{1}{n-1}\right) * \Sigma(X_i - \mu_X)^2 \quad (4)$$

$$\nabla_X = \left(\frac{1}{n-1}\right) * \Sigma\left(\frac{|X_i - X_{i-1}|}{\Delta t}\right) \quad (5)$$

where X represents RSRP or SINR. The variance measures the extent of fluctuation of the signal around its mean, while the gradient measures the rate of change between successive samples. The two are complementary: A device that drifts slowly across a broad spectrum yields high variance with a low gradient, while a device that remains stationary yet undergoes rapid small-scale fading results in low variance with a higher gradient. Using both, together with the mean, provides the classifier with a more complete understanding than any individual statistic.

Table 2 presents the statistical parameter ranges observed for each UE mobility category. These ranges were obtained empirically from the Monte Carlo simulations described in Section 3, in which UEs of each category were generated under the 3GPP TR 38.901 UMi channel model and their RSRP and SINR statistics were aggregated over 100 runs. They are therefore measured outputs of this study, not normative values specified by 3GPP or taken from prior literature.

Table 2. Empirically observed ranges of RSRP and SINR statistical parameters by UE mobility category, obtained from Monte Carlo simulation under the 3GPP TR 38.901 UMi channel model.

Category (Speed)	σ_{RSRP}^2 (dB ²)	σ_{SINR}^2 (dB ²)	∇_{RSRP} (dB/ms)	∇_{SINR} (dB/ms)
Stationary (0 km/h)	1–4	0.5–2	<0.01	<0.01
Slow (2 km/h)	3–7	1–4	0.01–0.03	0.01–0.03
Pedestrian (4 km/h)	6–12	3–8	0.03–0.08	0.03–0.07
Borderline (10 km/h)	10–20	6–15	0.06–0.15	0.05–0.12
Vehicular (45 km/h)	18–30	12–25	0.12–0.30	0.10–0.25

2.2. The Low-Mobility Criterion

A UE is classified as low mobility only when all four of the following conditions are satisfied simultaneously:

$$C_{low} = \left(\sigma_{RSRP}^2 \leq \gamma_{RSRP} \right) \wedge \left(\nabla_{RSRP} \leq \delta_{RSRP} \right) \wedge \left(\sigma_{SINR}^2 \leq \gamma_{SINR} \right) \wedge \left(\nabla_{SINR} \leq \delta_{SINR} \right) \quad (6)$$

In this context, γ_{RSRP} and γ_{SINR} represent the variance thresholds, while δ_{RSRP} and δ_{SINR} denote the gradient thresholds. The nominal values in this study are $\gamma_{RSRP} = 10 \text{ dB}^2$, $\gamma_{SINR} = 7 \text{ dB}^2$, $\delta_{RSRP} = 0.04 \text{ dB/ms}$, and $\delta_{SINR} = 0.06 \text{ dB/ms}$.

These numbers are not arbitrary; they are situated in relation to the measured parameter ranges presented in Table 2. The RSRP variance threshold of 10 dB^2 is situated above the slow-mobility band (about $3\text{--}7 \text{ dB}^2$) and below the pedestrian band (approximately $6\text{--}12 \text{ dB}^2$), positioning it inside the transitional range between the two, with a margin that supports estimation noise. The SINR variance threshold of 7 dB^2 is established below the RSRP threshold because of the natural noisiness of SINR variance, which is influenced by the interference environment. The gradient thresholds are positioned at the upper limit of the slow-mobility gradient range, ensuring that a truly slow UE is included, while a pedestrian-speed UE, which consistently exceeds this gradient, is excluded. The thresholds were thus calibrated to lie at the empirical boundaries separating adjacent mobility classes in Table 2, rather than being obtained through a formal optimization procedure. To verify that this calibration is robust, a systematic sensitivity analysis—in which each threshold is varied by $\pm 20\%$ while the other three are held fixed—is reported in Section 4.8; it confirms that classification accuracy remains within a narrow band of 90.97% to 93.45% across all configurations, indicating that the operating point is largely insensitive to small calibration errors. A formal learning-based optimization of the threshold set, building on these sensitivity results, is identified as future work in Section 5.8.

The logical AND framework makes the criterion conservative, with conservatism serving as a safety mechanism. If any of the four conditions is unmet, the UE is categorized as mobile and maintains the standard ReportInterval of 240 ms. Since all four conditions must be satisfied simultaneously, the rule is highly selective, making the accidental classification of a mobile UE as stationary extremely unlikely and thus yielding a low false-positive rate. Conversely, a truly stationary UE will typically meet all four criteria, ensuring high sensitivity. The SINR variance threshold is intentionally set lower than that of RSRP, reflecting the inherently noisier nature of SINR due to interference fluctuations. The gradient thresholds, positioned between the ranges characteristic of stationary and pedestrian motion, serve as a secondary verification layer that responds more rapidly to changes in signal behavior.

2.3. Justification for the Four Conditions

Each of the four conditions examines mobility from a distinct perspective, and crucially, none of them is redundant—every criterion contributes information that the others cannot capture. RSRP variance reflects large-scale signal fluctuations, which intensify as a UE traverses regions of varying shadow fading. SINR variance serves a complementary role, capturing degradation in signal quality while also being sensitive to shifts in the interference environment. The two gradient-based terms, by contrast, characterize the rate of signal change rather than its magnitude. This distinction is important: A UE drifting slowly across a wide area will exhibit high variance yet a low gradient, whereas a stationary UE experiencing rapid small-scale fading will present the opposite pattern—low variance alongside a relatively elevated gradient. A classifier relying solely on variance would fail to identify the latter case, just as one based exclusively on gradient would overlook the former. By combining both variance and gradient measures across RSRP and SINR, the classifier gains two complementary views of mobility, each capable of detecting what the

other would otherwise miss. The logical AND connecting them guarantees that a UE is classified as low mobility only when all four perspectives concur, which accounts for the heightened sensitivity discussed in Section 4.

A formal ablation study, in which each condition is sequentially removed and classification performance re-evaluated within the same simulation environment, is presented in Section 4.7; it confirms that the four conditions contribute unequally, with the RSRP gradient and SINR variance emerging as the primary discriminative features.

2.4. Five-Level ReportInterval (T_{RI}) Selection

Upon fulfilling the low-mobility criterion, the algorithm selects one of five ReportInterval values. The five levels are assessed hierarchically, starting with the most aggressive: A UE is assigned to the highest level that meets its mean and variance criteria, and descends to the subsequent level if it does not. The selection rule can be succinctly expressed as a case expression:

$$D_1 : \mu_{RSRP} \geq -80, \sigma_{RSRP}^2 \leq 4, \mu_{SINR} \geq 20, \sigma_{SINR}^2 \leq 2 \rightarrow T_{RI} = 5120 \text{ ms} \quad (7)$$

$$D_2 : \mu_{RSRP} \geq -90, \sigma_{RSRP}^2 \leq 8, \mu_{SINR} \geq 13, \sigma_{SINR}^2 \leq 5 \rightarrow T_{RI} = 2048 \text{ ms} \quad (8)$$

$$D_3 : \mu_{RSRP} \geq -100, \sigma_{RSRP}^2 \leq 12, \mu_{SINR} \geq 6, \sigma_{SINR}^2 \leq 8 \rightarrow T_{RI} = 1024 \text{ ms} \quad (9)$$

$$D_4 : \mu_{RSRP} \geq -110, \sigma_{RSRP}^2 \leq 16, \mu_{SINR} \geq 0 \rightarrow T_{RI} = 480 \text{ ms} \quad (10)$$

$$D_5 : \text{otherwise} \rightarrow T_{RI} = 240 \text{ ms} \quad (11)$$

The per-level threshold sets $\{\Theta_k, \Phi_k, \Psi_k, \Omega_k\}$ are presented in Table 3. The case rule is applied sequentially: The algorithm initially evaluates Level 1. If the conditions for Level 1 are not satisfied, the process advances to Level 2, continuing in this manner for subsequent levels. At Level 4, the SINR variance check is excluded because low SINR values inherently cause signal instability, which makes the check ineffective. If a UE does not meet the condition (C_{low}), the ReportInterval remains at 240 ms, consistent with current system settings. Table 3 provides a summary of the specific conditions for each level and their corresponding T_{RI} values.

Table 3. Five-Level T_{RI} selection rules and typical deployment cases.

Level	$\mu_{RSRP}(\text{dBm})$	$\sigma_{RSRP}^2(\text{dB}^2)$	$\mu_{SINR}(\text{dB})$	$\sigma_{SINR}^2(\text{dB}^2)$	$T_{RI}(\text{ms})$	Typical Scenario
1	≥ -80	<4	≥ 20	<2	5120	Fixed IoT sensor
2	≥ -90	<8	≥ 13	<5	2048	Smart meter, indoor monitor
3	≥ -100	<12	≥ 6	<8	1024	Slow movement indoors
4	≥ -110	<16	≥ 0	(not applied)	480	Near-boundary, mildly drifting
5	otherwise	otherwise	otherwise	otherwise	240	Walking, driving (standard)

All five ReportInterval values are sourced directly from the 3GPP-standardized set, so no proprietary signaling is necessary [4]. The variance thresholds escalate in approximately 4 dB² increments, while ReportInterval approximately doubles at each level, resulting in a proportional trade-off between energy conservation and reporting timeliness. At Level 4, the SINR variance condition is not implemented, as low SINR values naturally show

large fluctuations, making the application of the condition ineffective. A UE that does not satisfy the low-mobility criterion remains at 240 ms—the identical interval used in existing systems.

The advantage over a binary approach is clear. Consider two UEs, one exhibiting an RSRP variance of 4 dB² and the other 1 dB²: A binary system would allocate the same intervals to both, whereas the five-level scheme assigns 2048 ms to the former and 5120 ms to the latter, maximizing energy conservation for the more stable device while exercising prudence with the less stable one. Standard Compliance and UE-Side Considerations. All five ReportInterval values (240, 480, 1024, 2048, and 5120 ms) are valid values from the standardized set defined in 3GPP TS 38.331 [4]. The gNB delivers each updated ReportInterval to a UE through the existing MeasConfig information element carried in the RRCReconfiguration message, which is part of the legacy 5G NR signaling procedure. No new RRC parameters, no protocol extensions, and no firmware updates are required on the UE side. The UE simply applies the updated reporting period as it would for any operator-configured value. The decision logic—statistical estimation, low-mobility criterion evaluation, level selection, and hysteresis—runs entirely at the gNB. Consequently, the algorithm is deployable as a gNB-side software update on existing 3GPP Release 17 infrastructure without modification to deployed UEs.

2.5. Hysteresis Mechanism

In the absence of hysteresis, a UE with variance near a threshold would oscillate between levels during each cycle. This behavior would generate an excessive number of RRCReconfiguration signals. The algorithm addresses this issue using two distinct methods.

The first mechanism introduces a hysteresis margin of $\Delta H = 3$ dB, whereby a transition from level i to level j is permitted only when the variance exceeds the target threshold by at least ΔH . To illustrate, consider a UE currently operating at Level 2 with an RSRP variance of 5.5 dB²—a value that falls short of the threshold by less than ΔH . Under this rule, the UE remains at Level 2 and no transition is triggered until the variance reaches the sum of the threshold and the margin ΔH . The 3 dB margin adopted here is consistent with the hysteresis offsets used for 3GPP A3 measurement events, which are configurable in the 0–15 dB range and are commonly set to 1–3 dB; a systematic optimization of this margin across different interference and mobility scenarios is identified as future work in Section 5.8.

$$\text{upward } (i \rightarrow i-1) : \sigma_{RSRP(t)}^2 \leq \Phi_{i-1} - \Delta H \quad (12)$$

$$\text{downward } (i \rightarrow i+1) : \sigma_{RSRP(t)}^2 \geq \Phi_i + \Delta H \quad (13)$$

The second mechanism employs a hold timer of $T_{hold} = 2$ s. Once the margin check is satisfied, the newly proposed level must remain stable for the entire duration of T_{hold} before the gNB commits to the transition and issues an RRCReconfiguration message. Should the signal revert within this window, the pending transition is immediately cancelled. In this way, the hold timer guards against transient signal dips, while the hysteresis margin addresses measurement noise near level boundaries—the two mechanisms thus serving distinct but complementary protective roles. The design incorporates a deliberate asymmetry: Transitions toward longer ReportInterval values are permitted to occur marginally faster than those toward shorter ones. This ensures that the algorithm promptly rewards sustained signal stability while remaining conservative in the opposite direction, preserving overall system reliability.

2.6. Energy Conservation Model

The power consumption of the UE is represented within the Discontinuous Reception (DRX) framework. A DRX cycle with a total time T_{cycle} consists of an active duration, an idleness phase, many short-sleep sub-cycles, and a deep-sleep period; the UE shows varying power consumption levels in each condition [23]. Consistent with this 3GPP TR 38.840 DRX power-consumption framework and with the DRX analysis of Lin et al. [3], the mean power as a function of T_{RI} is

$$P_{avg}(T_{RI}) = P_{avg,0} + N_{MR(T_{RI})} \cdot \frac{E_{MR}}{T_{cycle}} \quad (14)$$

where $P_{avg,0}$ denotes the baseline average power consumption in the absence of measurement reporting, $N_{MR(T_{RI})} = \left\lfloor \frac{T_{cycle}}{T_{RI}} \right\rfloor$ is the number of MeasurementReport messages transmitted per DRX cycle, and E_{MR} represents the energy consumed by a single report, including radio wake-up and transmission costs. As T_{RI} increases, N_{MR} decreases, thereby reducing the average power. Note: Equation (14) provides a model-based estimate of average UE power. All energy values reported in this study are computed from this model in MATLAB simulation, not derived from real hardware measurement; empirical validation against hardware measurement is identified as future work in Section 5.8. The comparative improvement over the static baseline is known as the Energy Saving Ratio (ESR).

$$ESR = \frac{(P_{avg,static} - P_{avg,adaptive})}{P_{avg,static}} \cdot 100\% \quad (15)$$

The overall average signaling burden across the network is derived by aggregating the contributions from each category, weighted by their respective fractions α_k .

$$S_{net} = \sum_{k=1}^5 \alpha_k \cdot \frac{N_{MR(D_k)}}{T_{cycle}} \quad (16)$$

where $N_{MR(D_k)}$ denotes the number of MeasurementReport messages transmitted per DRX cycle by UEs assigned to level D_k . Equation (16) enables an operator to immediately assess the reduction in network-level signaling based on the mobility-category composition of their deployment.

The practical effect of energy conservation on battery-operated devices can be stated as an increase in device longevity. If $T_{bat,static}$ represents the battery duration in the static 240 ms arrangement, then the adaptive duration is:

$$\Delta T_{bat} = T_{bat,static} \cdot \frac{ESR}{(100 - ESR)} \quad (17)$$

A sensor with a baseline lifespan of one year operating at Level 1 experiences a lifetime extension of several weeks, thereby diminishing the frequency of battery-replacement field visits.

2.7. Energy Model Validation and Scope

The energy values reported throughout this paper are derived from the DRX power consumption model defined in Equation (14). This subsection clarifies the validation status of the model and the intended scope of the reported savings.

Model Basis. The DRX power consumption profile underlying Equation (14) is grounded in 3GPP TR 38.840 [23], which provides standardized power consumption profiles for 5G NR UEs across different DRX states. The reference power values adopted in this study— $P_{ON} = 255.5$ mW for active reception, $P_I = 255.5$ mW for inactive listening, $P_{SS} = 11$ mW for short sleep, and $P_{DS} = 0.7$ mW for deep sleep—are drawn directly from Table 5.1.1.4-1 of 3GPP TR 38.840. The DRX cycle timer values conform to 3GPP TS 38.321 [24].

Estimated Versus Measured Savings. All energy savings reported in this study constitute model-based estimates derived from Equation (14) under the assumed DRX configuration and are not obtained from direct measurements on commercial 5G NR hardware. The reported reduction of 10.87% should therefore be interpreted as a model-level estimate of the power saving expected to be achievable when the proposed adaptive ReportInterval mechanism is deployed under the specified DRX parameters.

Several factors may contribute to a divergence between the model-based estimates and actual hardware measurements. First, chipset implementation differences across vendors—including Qualcomm, MediaTek, and Samsung—typically introduce a variation of ± 20 to 30 percent relative to the 3GPP TR 38.840 reference values. Second, state transition overhead, referring to the energy consumed during wake-up transients between deep sleep and active reception, varies in practice and is heavily dependent on the RF front-end design of the specific UE. Third, background processes—such as modem firmware activity and baseband processing—impose additional power consumption that falls outside the scope of the analytical model and are therefore not accounted for in the reported estimates.

Empirical validation of the reported savings would require direct power measurements conducted on commercial 5G NR testbeds—such as those built on OpenAirInterface, Amarisoft, or equivalent commercial test platforms. This represents an essential direction for future work and is restated accordingly in Section 5.8. Nevertheless, the model-based analysis presented here retains significant value for relative comparison across configurations, which constitutes the central objective of this study, and serves as a principled foundation for designing subsequent empirical investigations.

The overall workflow and operational steps of the proposed adaptive ReportInterval scheme are illustrated in Figure 1. The gNB extracts RSRP and SINR from regular MeasurementReport messages, computes the statistical parameters (μ , σ^2 , ∇), evaluates the low-mobility criterion (Equation (6)), and, if satisfied, selects the appropriate level $D_k \in \{D_1, D_2, D_3, D_4\}$ according to Equations (7)–(11); otherwise, D_5 (240 ms) is applied. A two-stage hysteresis mechanism (Equations (12) and (13)) suppresses ping-pong transitions before the new ReportInterval is communicated to the UE via MeasConfig.




Figure 1. Block diagram of the proposed adaptive ReportInterval algorithm.

3. Simulation Setup

3.1. Channel Model

The simulation uses the 3GPP TR 38.901 Urban Micro (UMi) channel model [25], which is the reference model for dense urban 5G NR deployments in the sub-6 GHz bands. Path loss is calculated using the UMi line-of-sight (LOS) and non-line-of-sight (NLOS) equations. Shadow fading is characterized as log-normal, exhibiting a standard deviation of 4 dB in Line of Sight (LOS) conditions and 7.82 dB in Non-Line of Sight (NLOS) conditions, with

spatial correlation adhering to the Gudmundson model. Fast fading is modeled using a Rician distribution in line-of-sight (LOS) conditions and a Rayleigh distribution in non-line-of-sight (NLOS) conditions. A single dominant interferer is located at the cell edge, generating interference at a level 3 dB below the serving cell's transmitted power. This configuration yields the SINR distribution upon which the mobility classifier operates.

The UMi model was chosen because dense urban micro-cellular deployments at sub-6 GHz are precisely the environments in which the targeted low-mobility IoT devices—utility meters, fixed sensors, and similar fixed-location equipment—are most commonly installed, making it the most representative 3GPP channel model for this study. The single dominant interferer at the cell edge was adopted as a deliberately controlled condition that yields a well-defined, reproducible SINR distribution, isolating the behavior of the proposed classifier from the confounding effects of variable multi-cell geometry; it also constitutes a conservative first-order model, since aggregate inter-cell interference is frequently dominated by the strongest neighboring cell. A comprehensive evaluation under multiple simultaneous interferers with explicit multi-cell handover—which would further allow the impact on handover timing to be quantified—lies beyond the scope of this single-cell study and is identified as a priority direction for future work in Section 5.8.

3.2. UE Mobility Model

A total of 200 UEs are partitioned into five distinct mobility groups, with a distribution chosen to reflect realistic IoT deployment conditions, where the majority of devices are either stationary or operating at low speeds. The initial two categories, stationary and slow, are set as the ground-truth low-mobility class (150 UEs); the subsequent three are classified as the ground-truth mobile class (50 UEs). Sensitivity measures the correct identification of the low-mobility class, whereas specificity measures the correct identification of the mobile class. UE placements are re-randomized at the start of each Monte Carlo simulation using a Random Waypoint model for the mobile categories. Table 4 describes the distribution.

Table 4. Categories and distribution of UE mobility.

Category	Speed (km/h)	Count	Share	Ground-Truth Class
1—Stationary	0	120	60%	Low mobility
2—Slow	2	30	15%	Low mobility
3—Pedestrian	4	20	10%	Mobile
4—Borderline	10	20	10%	Mobile
5—Vehicular	45	10	5%	Mobile

3.3. Specifications

Table 5 consolidates all simulation parameters in a single location. We conducted 100 Monte Carlo iterations, each lasting 300 s of simulated time with a 20 millisecond sampling interval, and present the mean along with the standard deviation over the iterations.

Table 5. Parameters for the simulation.

Parameter	Value	Reference
Channel model	3GPP TR 38.901 UMi	[25]
Carrier frequency	3.5 GHz	5G NR n78 band
Bandwidth	100 MHz	-
gNB transmit power	44 dBm	[25]

Table 5. *Cont.*

Parameter	Value	Reference
gNB height	25 m	UMi
UE height	1.5 m	UMi
Cell radius	500 m	-
Shadow fading std	4 dB (LOS), 7.82 dB (NLOS)	[25]
Noise figure	7 dB	-
Noise spectral density	−174 dBm/Hz	-
Number of UEs	200	-
Simulation duration	300 s	-
Sampling step (Δt)	20 ms	-
Monte Carlo iterations	100	-
$\gamma_{RSRP}, \gamma_{SINR}$	10 dB ² , 7 dB ²	Calibrated
$\delta_{RSRP}, \delta_{SINR}$	0.04, 0.06 dB/ms	Calibrated
$\Delta H, T_{hold}$	3 dB, 2 s	Hysteresis
T_{RI} level set	{240, 480, 1024, 2048, 5120} ms	3GPP TS 38.331 [4]
DRX cycle T_{cycle}	870 ms	[23]
$P_{ON}/P_I/P_{SS}/P_{DS}$	255.5/255.5/11/0.7 mW	[23]
E_{MR}	0.7 mJ	[23]
Simulation tool	MATLAB R2022a	-

4. Results and Analysis

4.1. Classification Performance

Table 6 shows the confusion matrix averaged over 100 Monte Carlo simulations. Of the 150 ground-truth low-mobility UEs, the system successfully detected 148.2 on average, missing only 1.8. Of the 50 ground-truth mobile UEs, 34.5 were successfully classified and 15.5 were labeled as low-mobility.

Table 6. C_{low} confusion matrix (averaged over 100 Monte Carlo runs).

	Actual Low-Mobility (150)	Actual Mobile (50)
Predicted low-mobility	TP = 148.2	FP = 15.5
Predicted mobile	FN = 1.8	TN = 34.5

Table 7 describes the standard metrics. The system achieves a sensitivity of 98.77%, indicating that nearly all authentically low-mobility UEs are identified, with a false-negative rate of just 1.23%. The specificity is 68.96%, with a comparable false-positive rate of 31.04%. The Value column reports the mean $\pm$ 95% confidence interval margin, where the margin is calculated using the t-distribution. The classification performance of the low-mobility classifier (C_{low}) across 100 Monte Carlo iterations is presented in Figure 2. $TP = 148.2$, $FP = 15.5$, $FN = 1.8$, $TN = 34.5$. Overall accuracy = 91.32% $\pm$ 1.35%; sensitivity = 98.77%; specificity = 68.96%.

The elevated sensitivity and small false-negative rate are intentional outcomes of the carefully designed AND structure of the low-mobility criterion, as elaborated in Section 2.2. From an operational perspective, a false negative, where a stationary UE is left for a 240 ms interval, simply results in a missed opportunity for energy savings and poses no risk. A

false positive, where a mobile UE is allocated an extended period, is the more consequential error; however, it is self-correcting: As the UE progresses, its signal variance increases, prompting the hysteresis mechanism to revert it to the 240 ms interval within one to two assessment cycles. The specificity of 68.96% indicates a deliberate trade-off rather than a shortcoming, while the F1-score of 94.47% affirms robust overall classification quality.

Table 7. Classification performance metrics (average over 100 iterations).

Metric	Formula	Value %
Accuracy	$(TP + TN) / Total$	91.32 ± 0.268
Sensitivity (recall)	$TP / (TP + FN)$	98.77 ± 0.173
Specificity	$TN / (TN + FP)$	68.96 ± 0.909
Precision	$TP / (TP + FP)$	90.52 ± 0.253
False-positive rate (FPR)	$FP / (FP + TN)$	31.04 ± 0.909
False-negative rate (FNR)	$FN / (TP + FN)$	1.23 ± 0.173
F1-score	$2TP / (2TP + FP + FN)$	94.47 ± 0.167




Figure 2. Confusion matrix of the low-mobility classifier (C_{low}), averaged across 100 Monte Carlo iterations.

To verify that this performance is not specific to the assumed deployment composition, the low-mobility fraction was varied from 50% to 90%—with the internal sub-proportions of each mobility class held fixed—over 100 Monte Carlo runs. Sensitivity and specificity remained essentially invariant across all mixes (98.7–98.9% and 68–70%, respectively), since the criterion classifies each UE solely from its own signal statistics; only the aggregate accuracy varied predictably with the mix, rising from 84.0% for a 50/50 composition to 95.9% for a 90/10 composition. This confirms that the classifier is robust to the proportion of low- and high-mobility devices in the cell. To evaluate robustness against variations in user-mix composition, Figure 3 shows the classification performance of the low-mobility criterion as the low-mobility UE fraction varies from 50% to 90%. Sensitivity and specificity

remain essentially invariant, while the aggregate accuracy increases with the low-mobility fraction, confirming robustness to the user-mix composition.




Figure 3. Classification performance of the low-mobility criterion versus the low-mobility UE fraction (50–90%), averaged over 100 Monte Carlo runs.

4.2. Impact of Energy and Signaling

Table 8 compares the adaptive and static methods. Based on the DRX energy model in Equation (14), the adaptive algorithm yields an estimated reduction in average UE power consumption from 28.15 mW to 25.09 mW, which corresponds to 10.87%, and reduces the network-wide MeasurementReport count by 51.09%. Handover behavior is not modeled in this single-cell scenario, and the effect on SINR is small.

Table 8. Adaptive versus static ReportInterval: Critical performance metrics (mean $\pm$ standard deviation over 100 iterations).

Metric	Static (240 ms)	Adaptive	Change
Average UE power	28.15 mW	25.09 mW	−10.87%
Energy Saving Ratio (ESR)	0%	10.87% $\pm$ 0.38	+10.87 pp
Power Saving Factor (PSF)	0.9092	0.9156	+0.0064
MeasurementReport reduction	0%	51.09% $\pm$ 1.92	+51.09 pp
RRCReconfiguration per run	6852.7	218.7	−96.8%
Handover behavior	not modeled	not modeled	N/A
SINR impact	-	very small	~0 dB

The magnitude of this energy saving is governed primarily by two factors. The first is the distribution of UEs across the five levels: As Table 9 shows, devices assigned to Level 1 (5120 ms) achieve the largest per-device reduction, so the network-wide saving grows with the proportion of UEs that qualify for the longest ReportInterval values—

which in turn depends on the mobility composition of the deployment. The second is the fraction of baseline UE power attributable to measurement reporting itself, since only this component is reduced by extending the interval, while the remaining baseline consumption is unaffected. It should be emphasized that the reported 10.87% is a model-based estimate; its validation scope and the hardware-dependent deviations (approximately ± 20 – 30% across chipset vendors) are discussed in detail in Sections 2.7 and 5.8.

Table 9. Comprehensive per-level evaluation (average across 100 Monte Carlo simulations, 200 UEs).

Level	T_{RI} (ms)	UEs	Share	P_{avg} (mW)	ESR (%)	MR/UE/300 s
1	5120	29.7	14.8%	25.37	9.88	59
2	2048	38.8	19.4%	25.58	9.15	146
3	1024	36.0	18.0%	25.92	7.93	293
4	480	27.9	13.9%	26.69	5.18	625
5 (standard)	240	67.6	33.8%	28.15	0.00	1250

To complement the standard deviations reported above, 95% confidence intervals (CIs) over the 100 Monte Carlo iterations were computed for the principal metrics. The classification accuracy has a 95% CI of [91.06%, 91.58%], the sensitivity a CI of [98.59%, 98.95%], the specificity a CI of [68.06%, 69.86%], the average energy saving a CI of [10.77%, 10.91%], and the MeasurementReport reduction a CI of [50.71%, 51.47%]. As each interval spans well under one percentage point, the reported results are statistically stable across independent runs. A performance comparison of energy consumption between the static (240 ms) and adaptive (5-level) configurations is presented in Figure 4.




Figure 4. Energy consumption comparison between static (240 ms) and adaptive (5-level) configurations. (a) Average UE power consumption: 28.15 mW (static) versus 25.09 mW (adaptive), a 10.87% reduction. (b) Power Saving Factor (PSF): 0.9092 (static) versus 0.9156 (adaptive), an improvement of +0.64 percentage points.

The distribution of 200 UEs across the five ReportInterval levels, averaged over 100 Monte Carlo runs, is illustrated in Figure 5.




Figure 5. Distribution of 200 UEs across the five ReportInterval levels: D₁ (5120 ms)—29.7 UEs (14.8%); D₂ (2048 ms)—38.8 UEs (19.4%); D₃ (1024 ms)—36.0 UEs (18.0%); D₄ (480 ms)—27.9 UEs (13.9%); D₅ (240 ms)—67.6 UEs (33.8%). Averaged over 100 Monte Carlo runs.

4.3. Energy Analysis by Level

Table 9 summarizes the energy consumption values associated with each of the five levels. The average power per UE was computed using Equation (14), with a reference power of $P_{avg,0} = 25.23$ mW derived from simulation results under the static 240 ms baseline configuration. The per-level Energy Saving Ratio (ESR) is then expressed relative to this baseline. The evidence confirms the proportional link between ReportInterval and ESR as predicted by the energy model.

Multiple observations are noteworthy. Level 1 accounts for the majority of energy savings: The 29.7 UEs allocated there each achieve a 9.88% reduction compared to the 240 ms baseline. This results in the network-wide figure of 10.87%, as given in Section 4.2, when aggregated across levels and weighted by the UE share. The MR reduction column presents a contrasting narrative: A Level-1 UE transmits 59 reports during the 300 s simulation period rather than 1250, indicating a 95.3% reduction per device. The removal of these reports accounts for the 51.09% reduction in the number of MeasurementReport messages across all UEs network-wide.

4.4. Category-Specific Analysis

Table 10 categorizes the results by mobility classification. Stationary and slow devices are identified with high accuracy and contribute substantially to energy conservation, but pedestrian, borderline, and vehicular devices are mostly maintained at the standard interval, as designed.

Table 10. Outcomes categorized by mobility (average across 100 iterations).


Category	$C_{low} = \text{True}$	Avg. Level	Energy Saved	MR Reduced
Stationary	~99%	1.6	~17%	~78%
Slow	~95%	2.4	~13%	~64%
Pedestrian	low	4.7	~1%	~5%
Borderline	very low	5.0	0%	0%
Vehicular	~0%	5.0	0%	0%
Network average	-	3.0	10.87%	51.09%

4.5. Effectiveness of the Hysteresis Mechanism

Section 2.5 presented the two-stage hysteresis mechanism. Table 11 measures the effect. In the absence of hysteresis, the algorithm produced 6852.7 level transitions per run; with hysteresis, this decreased to 218.7, a reduction by a factor of 31.33. Each level transfer corresponds to a single RRCReconfiguration message, resulting in a reduction in reconfiguration signaling. This reduction is among the largest reported for comparable mechanisms. The effect of the proposed hysteresis mechanism on stabilizing level transitions and reducing signaling overhead is illustrated in Figure 6.

Table 11. Analysis of level transitions with and without hysteresis (average over 100 iterations, 200 UEs, 300 s).

Metric	No Hysteresis	With Hysteresis	Ratio
Total level transitions (per run)	6852.7	218.7	31.33x
RRCReconfiguration messages (per run)	6852.7	218.7	31.33x
Transitions per UE (mean)	34.3	1.1	31.33x
Power Saving Factor (static -> adaptive)	0.9092	0.9156	+0.0064

**Figure 6.** Effect of the hysteresis mechanism on level transitions per iteration. Without hysteresis: 6852.7 transitions per iteration; with hysteresis: 218.7 transitions per iteration; reduction factor of 31.33 $\times$, equivalent to 96.8% fewer transitions and RRCReconfiguration messages.

The reduction by a factor of 31.33 represents the single largest effect observed in this study, with two direct consequences. First, each avoided level transition results in one fewer RRCReconfiguration message on the downlink, conserving uplink and downlink signaling capacity and preventing activation of the UE radio. Secondly, each evaded RRCReconfiguration eliminates a brief active interval for the UE, substantially enhancing overall energy conservation. The PSF improves from 0.9092 (static baseline) to 0.9156 (with the hysteresis-integrated adaptive method), indicating that hysteresis is an essential element of the design.

The asymmetric design of the hysteresis, featuring a 3 dB tolerance for variation thresholds and a 2 s holding timer, generates this substantial effect. In the absence of the holding timer, transient fading dips would often initiate downward transitions that later reverse; the 2 s window eliminates rapid fluctuations, while the margin supports consistent oscillations around a threshold boundary. A detailed analysis of upward and downward transitions and their respective contributions to the overall reduction is reserved for future research.

4.6. Quantitative Comparison with Baseline Methods

To validate the proposed method, two baseline approaches were re-implemented under identical simulation conditions, and all three methods were evaluated against one another using 100 Monte Carlo iterations applied to the same channel realizations. The first baseline follows the RRM Relaxation specification in 3GPP TS 38.331 [11], which employs a single RSRP threshold ($\text{RSRP} > -100$ dBm) to identify low-mobility UEs and assigns a single extended ReportInterval. The second baseline is a Binary Threshold method that combines RSRP and SINR conditions ($\text{RSRP} > -100$ dBm AND $\text{SINR} > 10$ dB). The proposed method, by contrast, applies the four-condition low-mobility criterion with hysteresis as formulated in Section 2. Additionally, literature-reported values from a recent AI-based DRX approach [14] are included to provide broader qualitative context.

Table 12 presents a quantitative comparison across five performance metrics: classification accuracy, sensitivity, specificity, energy saving, and MeasurementReport reduction. The values are the mean $\pm$ 95% confidence interval margin, where the margin is calculated using the t-distribution.

Table 12. Quantitative comparison of four classification methods (200 UEs, 100 Monte Carlo iterations, 3GPP TR 38.901 UMi channel); AI-DRX values [14] are shown for qualitative reference only.

Method	Accuracy %	Sensitivity %	Specificity %	Energy Saving %	MR Reduction %
RRM Relaxation [11]	81.96 ± 0.272	99.04 ± 0.163	30.72 ± 0.891	9.96 ± 0.026	45.80 ± 0.117
Binary Threshold	45.74 ± 0.641	32.45 ± 0.739	85.62 ± 1.013	3.04 ± 0.063	13.96 ± 0.288
AI-DRX [14] *	~85	N/R	N/R	~15	N/R
Proposed (this work)	91.32 ± 0.268	98.77 ± 0.173	68.96 ± 0.909	10.87 ± 0.076	51.09 ± 0.381

* Approximate values reported in Lazrek et al. [14] under different simulation conditions (different channel model, traffic profile, and UE mobility mix); direct quantitative comparison requires re-implementation in an identical framework, which is identified as future work in Section 5.8. N/R = not reported in the source.

The results reveal a fundamental limitation inherent to single-criterion approaches. RRM Relaxation achieves high sensitivity (99.04%) but at the cost of severely degraded specificity (30.72%): While it correctly identifies the vast majority of stationary UEs, it simultaneously misclassifies 69.3% of mobile UEs as stationary. In a deployed network, this would result in the ReportInterval being extended for a substantial proportion of moving devices, potentially delaying handover triggers with adverse consequences for mobility management. The Binary Threshold method exhibits the opposite weakness—high speci-

ficity (85.62%) coupled with very low sensitivity (32.45%)—detecting only approximately one third of stationary devices and yielding negligible energy savings of just 3.04%.

Among the three re-implemented methods, the proposed approach achieves the most favorable overall balance across all evaluated metrics. Classification accuracy reaches 91.32%, representing an improvement of 9.36 percentage points over RRM Relaxation and 45.58 percentage points over the Binary Threshold method. Sensitivity remains highly competitive at 98.77%, while specificity of 68.96% more than doubles that of RRM Relaxation. The proposed method also records the highest energy saving (10.87%) and MeasurementReport reduction (51.09%) among the re-implemented baselines. These performance gains are attributable to the joint application of four statistical conditions connected by an AND operator, combined with a two-stage hysteresis mechanism, together enabling the classifier to reliably distinguish genuine stationarity from transient periods of signal quietude in otherwise mobile UEs. Pairwise Welch's *t*-tests over the 100 Monte Carlo runs further confirm that these accuracy differences are statistically significant, with $p < 0.001$ for the proposed method versus both RRM Relaxation and the Binary Threshold method. A quantitative performance comparison between the proposed method, RRM Relaxation, and Binary Threshold across various evaluation metrics is illustrated in Figure 7.




Figure 7. Quantitative comparison of accuracy, sensitivity, specificity, energy saving, and MR reduction for RRM Relaxation, Binary Threshold, and the proposed method under identical simulation conditions (100 Monte Carlo iterations). AI-DRX [14] values are presented separately in Table 12 due to the different simulation framework.

Welch's *t*-tests were performed to compare the proposed method, RRM, and Binary across all five metrics (Accuracy, Sensitivity, Specificity, Energy saving, MR reduction). The Bonferroni-corrected significance threshold was set to $\alpha = 0.0033$, considering the 15 pairwise comparisons (Table 13). The proposed method was significantly better than both RRM and Binary for Accuracy, Energy saving, and MR reduction, and RRM was significantly better than Binary (all $p < 0.001$, $d > 3$), resulting in a consistent ranking of Proposed > RRM > Binary. For Specificity, however, the ranking was inverted, with Binary significantly outperforming Proposed and Proposed significantly outperforming RRM (all $p < 0.001$, $d > 3$). For Sensitivity, both Proposed and RRM substantially outperformed

Binary ($p < 0.001$, $d > 24$), while the difference between Proposed and RRM, though nominally significant under an uncorrected threshold ($p = 0.027$), did not survive Bonferroni correction ($d = -0.31$), indicating no reliable difference between the two methods on this metric. Across comparisons that remained significant after correction, effect sizes were uniformly very large (d ranging from 3.09 to 28.73), substantially exceeding conventional benchmarks for large effects ($d \geq 0.8$), indicating that the observed differences reflect not only statistical significance but also substantial practical magnitude.

Table 13. Pairwise Welch's t -test comparisons across methods.

Metric	Comparison	Mean Diff	$t(df)$	p	Cohen's d	Sig. (Bonferroni)
Accuracy	Proposed vs. RRM	9.36	48.61 (198.0)	<0.001	6.87	Yes
	Proposed vs. Binary	45.58	130.14 (132.6)	<0.001	18.40	Yes
	RRM vs. Binary	36.22	103.17 (133.6)	<0.001	14.59	Yes
Sensitivity	Proposed vs. RRM	−0.27	−2.22 (197.3)	0.027	−0.31	No
	Proposed vs. Binary	66.33	173.39 (109.8)	<0.001	24.52	Yes
	RRM vs. Binary	66.59	174.59 (108.6)	<0.001	24.69	Yes
Specificity	Proposed vs. RRM	38.24	59.62 (197.9)	<0.001	8.43	Yes
	Proposed vs. Binary	−16.66	−24.29 (195.7)	<0.001	−3.44	Yes
	RRM vs. Binary	−54.90	−80.74 (194.8)	<0.001	−11.42	Yes
Energy saving	Proposed vs. RRM	0.885	21.83 (120.9)	<0.001	3.09	Yes
	Proposed vs. Binary	7.806	157.00 (190.8)	<0.001	22.20	Yes
	RRM vs. Binary	6.921	203.16 (131.0)	<0.001	28.73	Yes
MR reduction	Proposed vs. RRM	5.286	26.32 (117.6)	<0.001	3.72	Yes
	Proposed vs. Binary	37.121	154.26 (184.3)	<0.001	21.82	Yes
	RRM vs. Binary	31.835	203.15 (131.0)	<0.001	28.73	Yes

Note. Welch's t -test (unequal variances assumed), $n = 100$ runs per group. Cohen's d computed using the average of the two group standard deviations as the pooling denominator. Bonferroni-corrected significance threshold: $\alpha = 0.05/15 = 0.0033$.

For broader context, we also include the AI-based DRX optimization approach reported by Lazrek et al. [14], which applies reinforcement learning to dynamic DRX configuration. Their work reports an accuracy of approximately 85% and energy savings of around 15% under their simulation conditions. However, their evaluation framework differs from ours in several aspects: different channel model, traffic profile, and UE mobility mix. Their method also focuses on DRX cycle adaptation rather than ReportInterval selection. A direct comparison would require independent re-implementation of their reinforcement-learning agent under our simulation framework, which is recognized as essential future work in Section 5.8. The proposed method, in contrast, operates with $O(n)$ complexity without requiring training data or convergence time, making it suitable for direct deployment on existing gNB infrastructure as a software update.


4.7. Ablation Study

To quantify the marginal contribution of each individual condition within the low-mobility criterion, an ablation study was conducted in which each of the four conditions was systematically removed in isolation, and the resulting impact on classification performance was measured. Each variant was evaluated over 100 Monte Carlo iterations under simulation parameters identical to those of the baseline configuration. Table 14 presents the results.

Table 14. Ablation study showing the impact of removing each condition in the low-mobility criterion (100 MC iterations).

Variant	Accuracy	Sensitivity	Specificity	FPR	FNR
All four conditions (baseline)	91.78%	98.93%	70.34%	29.66%	1.07%
Without σ_{RSRP}^2	91.77%	98.93%	70.28%	29.72%	1.07%
Without ∇_{RSRP}	88.92%	98.93%	58.90%	41.10%	1.07%
Without σ_{SINR}^2	90.00%	100.00%	59.98%	40.02%	0.00%
Without ∇_{SINR}	91.78%	98.93%	70.34%	29.66%	1.07%

The ablation results demonstrate that the four conditions do not contribute equally to overall classification performance, see Figure 8. The removal of the RSRP gradient condition produces the most pronounced impact: Accuracy declines from 91.78% to 88.92% (−2.86 percentage points), and specificity falls from 70.34% to 58.90% (−11.44 percentage points), with the false-positive rate rising correspondingly from 29.66% to 41.10%. The removal of the SINR variance condition yields a comparable degradation: Specificity drops to 59.98% and the false-positive rate increases to 40.02%. Notably, sensitivity reaches 100% in this variant, a consequence of the relaxed criterion classifying a broader set of UEs as stationary, thereby capturing all truly stationary devices, but at the cost of increased misclassification of mobile ones.

**Figure 8.** Ablation study results: Impact of removing each of the four mobility conditions on accuracy, sensitivity, specificity, and false-positive rate.

In contrast, removing either the RSRP variance or the SINR gradient condition individually produces no measurable degradation, with accuracy remaining stable at 91.77–91.78% in both cases. This suggests that these two conditions are partially redundant when com-

bined with the remaining conditions through the AND operator. The RSRP gradient and SINR variance thus emerge as the primary discriminative features, while the RSRP variance and SINR gradient function as supplementary safety margins against transient signal fluctuations. It is worth noting that a simplified two-condition variant of the criterion could represent a viable option for resource-constrained deployments, and this possibility is identified as a direction for future work in Section 5.8.

This redundancy finding does not, however, imply that two of the four conditions are dispensable in the full criterion. The ablation removes conditions one at a time; the absence of degradation when either the RSRP variance or the SINR gradient is removed individually indicates that the two pairs provide overlapping protective coverage, not that both can be discarded simultaneously. Retaining all four conditions preserves this redundancy as a safety margin: The supplementary conditions guard against transient fluctuations and boundary cases—such as the longer-duration false-positive episodes of borderline-mobility devices analyzed in Section 4.9—that are not fully reflected in the run-averaged metrics yet may arise under specific propagation or interference conditions. Because each additional condition costs only an $O(1)$ comparison (Section 4.10), the computational penalty of retaining all four is negligible, and the conservative four-condition AND rule remains the recommended default. The simplified two-condition variant is therefore positioned not as a replacement but as an explicit option for resource-constrained deployments.

4.8. Threshold Sensitivity Analysis

To assess the sensitivity of the proposed method to threshold selection, each of the four thresholds in the low-mobility criterion was varied by $\pm 20\%$ in isolation—with the remaining three held at their baseline values—and 100 Monte Carlo iterations were conducted for each resulting configuration, yielding a total of 12 experiments. The corresponding classification accuracy values are reported in Table 15. The results of the threshold sensitivity analysis, evaluating classification accuracy across -20% , baseline, and $+20\%$ variations for each of the four thresholds, are depicted in Figure 9.

Table 15. Classification accuracy under $\pm 20\%$ variation of each threshold (100 MC iterations per configuration).

Threshold	-20%	Baseline	$+20\%$
$\gamma_{RSRP} = 10 \text{ dB}^2$	91.87%	91.78%	91.77%
$\gamma_{SINR} = 7 \text{ dB}^2$	91.83%	91.78%	91.11%
$\delta_{RSRP} = 0.04 \text{ dB/ms}$	93.45%	91.78%	90.97%
$\delta_{SINR} = 0.06 \text{ dB/ms}$	91.83%	91.78%	91.78%

The analysis yields three principal findings. First, the algorithm demonstrates strong robustness to threshold variations: Across all 12 configurations, classification accuracy remains within the narrow range of 90.97% to 93.45%, corresponding to a maximum spread of just 2.48 percentage points. This stability indicates that the operating point is largely insensitive to small calibration errors—a practically important property for deployment across diverse cell environments and propagation conditions.

Second, the RSRP gradient threshold emerges as the most influential parameter. A 20% reduction improves accuracy from 91.78% to 93.45% and specificity from 70.34% to 77.00%, while decreasing the false-positive rate from 29.66% to 23.00%. This suggests that a tighter gradient bound for stationary devices could further enhance discriminative performance, and systematic threshold optimization through learning-based methods is identified as a direction for future work.




Figure 9. Threshold sensitivity analysis: Classification accuracy for each of the four thresholds under -20% , baseline, and $+20\%$ variations.

Third, the SINR variance threshold exhibits a clear sensitivity–specificity trade-off. Reducing it by 20% raises specificity to 75.76% at the cost of a modest decrease in sensitivity to 97.19% , whereas increasing it by 20% produces the opposite effect—sensitivity of 99.47% alongside specificity of 66.00% . This parameter can therefore serve as an application-specific tuning knob, adjusted according to the relative priority of sensitivity versus specificity in a given deployment context.

By contrast, the RSRP variance and SINR gradient thresholds prove highly stable: $\pm 20\%$ variations produce accuracy changes of less than 0.1 percentage points. This is consistent with the findings of the ablation study in Section 4.7, where these same two conditions were identified as partially redundant. Taken together, both analyses indicate that while these thresholds contribute to overall robustness through the AND operator, their individual values can be adjusted across a wide range without any measurable impact on classification performance.

4.9. Temporal Behavior of False-Positive Classifications


The aggregate false-positive rate reported in Section 4.1, while informative, does not address a critical operational question: When a mobile UE is briefly misclassified as stationary, how long does the misclassification persist before the hysteresis mechanism intervenes to correct it? This duration is of direct practical significance, as it determines whether a false positive constitutes a meaningful operational disruption—such as a delayed measurement update during a critical handover window—or remains a transient event with negligible consequence. To investigate this, all false-positive episodes were logged across 30 Monte Carlo iterations, with the start and end times of each misclassification recorded for subsequent analysis. Table 16 summarizes the results, based on 36,539 false-positive episodes.

Table 16. Temporal statistics of false-positive episode durations (36,539 episodes across 30 Monte Carlo iterations).

Statistic	Value
Total FP episodes	36,539
Average per iteration	1218
Minimum duration	1.00 s
Median duration	2.00 s
Mean duration	3.80 s
Standard deviation	6.32 s
95th percentile	12.00 s
99th percentile	25.00 s
Maximum duration	300.00 s
Outliers (>30 s)	220 (0.6%)

The median false-positive duration of 2.00 s corresponds exactly to the hysteresis holding timer ($T_{hold} = 2$ s) defined in Section 2.5, reflecting the expected system behavior: When a mobile UE is briefly classified as stationary, the rising RSRP and SINR variance cross the detection threshold within one hysteresis cycle, prompting the algorithm to revert to the default ReportInterval of 240 ms. The empirical median thus confirms that the hysteresis mechanism operates precisely as designed.

The distribution of false-positive episode durations on a logarithmic scale ($n = 36,539$) is presented in Figure 10, with key statistics highlighted by vertical markers. The mean false-positive duration is 3.80 s (standard deviation: 6.32 s), with the 95th percentile reaching 12.00 s—meaning that 95% of all false-positive episodes are resolved within 12 s. The 99th percentile is 25 s. For typical IoT applications, where measurement reporting periods span from hundreds of milliseconds to several seconds, a characteristic false-positive duration of 2 to 12 s falls well within acceptable operational tolerance.

**Figure 10.** Histogram of false-positive episode durations on a logarithmic scale ($n = 36,539$). Vertical lines mark the median (2.00 s), mean (3.80 s), and 95th percentile (12.00 s).

A small proportion of episodes—220 out of 36,539, representing 0.6% of the total—exceeded 30 s in duration. These outliers correspond predominantly to UEs operating within the transitional speed range of 5 to 15 km/h (mobility category 4, Section 2.2), where instantaneous SINR variance can intermittently fall below the detection threshold during brief periods of signal quietude. The maximum observed duration of 300 s represents an isolated edge case in which the UE sustained near-stationary signal variance throughout the entire simulation window. Adaptive threshold tuning for boundary mobility scenarios is identified as a direction for future work.

4.10. Considerations of Complexity

The per-UE computational burden in a single decision cycle comprises the statistical estimation of four parameters and two gradients, which is $O(n)$ relative to the observation-window sample count n ; the low-mobility criterion evaluation, entailing four comparisons and three AND operations, $O(1)$; the five-level selection, involving a maximum of four hierarchical checks, $O(1)$; and the hysteresis check, consisting of two comparisons and one timer test, $O(1)$. The primary factor is the statistical estimation, resulting in an overall per-UE complexity of $O(n)$. For a typical observation window of $n = 50$ samples—corresponding to one second at 20 ms intervals—the total computational requirement amounts to approximately 134 operations per UE. Scaled to a cell supporting 200 UEs, this translates to roughly 27,000 operations per decision cycle, a negligible burden for a modern gNB processor. Notably, the computational cost scales linearly with the number of UEs, standing in sharp contrast to the quadratic complexity that characterizes reinforcement learning-based approaches.

5. Discussion

The results presented above demonstrate that the proposed approach successfully reduces both energy consumption and signaling overhead without any observed degradation in the single-cell scenario examined in this study, noting that compatibility with inter-cell handover under multi-cell conditions remains an open question identified as a direction for future work in Section 5.8. This section proceeds to examine the underlying factors driving the design's behavior, its alignment with existing 3GPP protocols, and the boundaries of its applicability.

5.1. The Trade-Off Between Sensitivity and Specificity

The proposed method deliberately prioritizes sensitivity (98.77%) over specificity (68.96%), a design choice grounded in the asymmetric consequences of each error type in a real network environment. A false negative—where a stationary UE is left on the standard ReportInterval—carries a relatively mild cost, amounting to a missed opportunity for energy savings. A false positive, by contrast, extends the ReportInterval for a UE that is in fact mobile, potentially delaying a handover check. However, this latter case is self-correcting by nature: A truly mobile UE will exhibit progressively growing signal variance, prompting the algorithm to restore the 240 ms interval within one or two decision cycles. The AND-based classification rule therefore steers errors toward the less harmful outcome, and the moderate specificity is an accepted and deliberate trade-off in exchange for this operational safety.

It is important to delineate precisely which devices are exposed to this error and what its operational cost is. First, latency-critical URLLC traffic is structurally protected: As detailed in Section 5.4, URLLC UEs transmit irregular, bursty traffic that produces elevated signal gradients, causing them to fail the low-mobility criterion automatically and remain at their standard configuration. False positives are therefore confined to the pool of

genuinely low-mobility candidates—stationary eMBB and mMTC devices—and cannot extend the ReportInterval of a URLLC bearer. Second, even when a mobile UE is briefly misclassified, the 3GPP safety mechanisms described in Section 5.3 remain fully active regardless of ReportInterval: Event-triggered measurement reporting (A1–A6 events), Radio Link Monitoring, and the T310/T311 timers continue to operate, so the worst-case reaction delay to an abrupt link degradation is bounded by the DRX wake-up cycle and the RLM procedure—on the order of one second—rather than by the extended ReportInterval. Third, the temporal analysis in Section 4.9 (Table 16) quantifies the duration of these episodes: The median false positive lasts 2.0 s and 95% are resolved within 12 s, confirming that misclassifications are short-lived and self-correcting. Taken together, a false positive is a brief, bounded, and self-correcting event that affects only non-latency-critical devices and does not erode the safety margins guaranteed by the 3GPP framework. For deployments that nonetheless demand higher specificity—for example, networks carrying a large share of borderline-mobility devices—the SINR variance threshold can be tightened to trade sensitivity for specificity, as quantified in the threshold sensitivity analysis of Section 4.8.

5.2. Independence from Application-Layer Traffic

A fundamental aspect of the proposed design is that the algorithm operates exclusively within the RRC control plane, leaving the application layer entirely unaffected. The MeasurementReport is inherently a radio quality assessment mechanism, utilized by the gNB Radio Resource Management (RRM) function to support handover decisions and scheduling—it carries no user data whatsoever. Application traffic, encompassing sensor readings, video streams, alarms, and real-time data, flows independently through the user plane via the PDCP, RLC, and MAC layers and remains entirely decoupled from the adaptive ReportInterval. Consequently, extending the ReportInterval does not delay application traffic. In the event of a time-sensitive occurrence on a UE with an extended ReportInterval, the associated uplink packet is allocated using the standard user-plane request protocol incorporating physical-layer wake-up signaling, without deferring to the subsequent MeasurementReport interval. The method modifies the frequency at which the UE reports radio statistics, rather than the speed of data transmission.

5.3. Coexistence with 3GPP Safety Protocols

The method operates within the established 3GPP measurement framework and does not deactivate any of its safety protocols. Three of these warrant explicit mention as they remain entirely operational regardless of the configured ReportInterval.

First, event-triggered measurement reporting remains active. The 3GPP A1 to A6 measurement events [4,26] are activated by signal quality conditions rather than a periodic timer. If the serving-cell quality falls below a configured threshold—for example, when an A2 event is triggered—the UE reports immediately, without waiting for the next periodic interval. Extending ReportInterval prolongs the periodic reporting interval while maintaining the integrity of event-triggered reporting.

Radio Link Monitoring (RLM). During each DRX wake-up, the UE measures in-sync and out-of-sync conditions relative to the Q_{in} and Q_{out} criteria. This assessment is integral to the DRX protocol [27] and is independent of ReportInterval.

Thirdly, the T310 and T311 timers regulate radio link failure detection and connection re-establishment. These timers constrain the worst-case response to abrupt link degradation, independent of the ReportInterval.

Collectively, these mechanisms indicate that for a UE subject to the algorithm, the maximum reaction delay to an abrupt signal deterioration is dictated by the DRX wake-up cycle and the RLM procedure, approximately one second, rather than by ReportInterval.

Upon observing degradation during the subsequent wake-up, the increasing variance and activated events prompt the algorithm to restore the UE to the 240 ms period. Extending ReportInterval for low-mobility devices does not increase the safety gap established by the 3GPP framework.

5.4. Operation Independent of Slices

An advantageous characteristic of the algorithm is its slice-agnostic nature. It does not use the network-slice identifier (S-NSSAI) [28] to determine its actions; rather, it categorizes each UE based on its observed RSRP and SINR metrics. This results in three effects. Initially, deployment is straightforward: The gNB requires no slice-topology setting to activate the algorithm. The method seamlessly combines with a multi-slice cell. URLLC UEs are generally configured with brief ReportInterval values and transmit irregular traffic that generates elevated gradient values, resulting in an automatic failure of the low-mobility criterion while maintaining their regular configuration; high-mobility eMBB UEs similarly experience low-mobility criterion failure due to their high variance. Stationary eMBB devices, including fixed-wireless-access routers and stationary cameras, along with low-mobility mMTC sensors, fulfill the low-mobility criterion and gain advantages from the adaptive interval. The method shows resilience to temporal variations in slice membership; it adapts on its own to changes in UE usage patterns based on signal statistics rather than labels, eliminating the need for reconfiguration.

5.5. Scalability for Extensive IoT Implementations

The 51.09% reduction in the number of MeasurementReport messages has effects that extend beyond an individual 200-UE cell. In a densely populated cell accommodating a higher UE population with a comparable mobility mix, the relative reduction is dictated by Equation (16) and remains near the per-cell value, as it depends on the category fractions rather than the total UE count. The total decrease in uplink signaling volume is substantial when extended to a city-wide deployment of several cells, achieved only through a gNB-side software modification. The per-UE computational expense, as shown in Section 4.10, increases linearly, ensuring the technique stays efficient even for wide cells.

5.6. Economic Viewpoint

The practical benefit of a 10.87% decrease in average power consumption for battery-operated IoT sensors is an extension of battery longevity and a concomitant decrease in the frequency of battery replacements. In installations needing physical site visits, such as industrial, agricultural, or remote-monitoring contexts, this results in a quantifiable decrease in maintenance expenses. The implementation cost for the operator is small, as the technique is a firmware-level enhancement to the gNB RRM function, needing no new hardware or modifications to existing UEs.

5.7. Significance to Uzbekistan's 5G Implementation

The deployment of 5G in Uzbekistan is advancing through several operators, with coverage enhancement occurring simultaneously alongside the integration of major IoT sectors, including smart metering, industrial automation, and precision agriculture. These sectors are mostly occupied by stationary devices, which is exactly the demographic the algorithm aims to target. The method is ideal for implementation in environments with limited operator budgets and a diverse device fleet, as it needs no new hardware, protocol extensions, or modifications on the UE side. The national 5G standardization efforts, including preliminary technical specifications for 5G network equipment, offer a direct avenue for the proposed models and thresholds to influence deployment protocols.

5.8. Limitations and Future Work

The evaluation presented in Sections 4.6–4.9 confirms that the proposed method achieves robust classification performance across multiple validation dimensions. Several areas, however, warrant further investigation.

First, the energy savings reported in this study are derived from the DRX energy model of Equation (14), which is grounded in 3GPP TR 38.840 [23]. As discussed in Section 2.7, this model represents the industry standard for analytical evaluation of UE power consumption; however, real-hardware power consumption may deviate by approximately ± 20 to 30 percent owing to chipset implementation differences across vendors. Empirical validation through direct power measurements on commercial 5G NR testbeds would substantially strengthen the practical applicability of the reported results and is therefore identified as an essential direction for future work.

Second, the evaluation is conducted within a single-cell scenario incorporating one dominant interferer. Extending the analysis to multi-cell configurations with explicit handover modeling would enable a more comprehensive assessment of the method's behavior under dense deployment conditions, such as urban small-cell networks—a scenario of growing practical relevance. This extension is likewise identified as a priority direction for future work.

Third, the comparison with AI/ML-based methods in Section 4.6 uses literature-reported values from [14]. A complete head-to-head benchmark requires independent re-implementation of reinforcement-learning and Kalman-filter-based approaches [13–15] within the same simulation framework, including their training and convergence phases. This benchmarking study is identified as essential future work.

The threshold sensitivity analysis presented in Section 4.8 demonstrated that the operating point is amenable to further optimization; reducing the RSRP gradient threshold by 20%, for instance, yielded an accuracy improvement to 93.45%. Systematic optimization through learning-based methods, in conjunction with an extension of the framework to the FR2 mm Wave band, represents a promising and practically relevant direction for future research.

Finally, two algorithmic refinements are left for future investigation. As indicated by the ablation results in Section 4.7, a simplified two-condition variant of the low-mobility criterion—retaining only the two most discriminative conditions—may suit resource-constrained gNB deployments and warrants dedicated evaluation. In addition, a systematic optimization of the hysteresis margin (ΔH) and holding timer (T_{hold}) across a range of interference and mobility scenarios, quantifying their joint effect on the signaling-overhead reduction reported in Section 4.5, would further refine the trade-off between transition stability and responsiveness.

6. Conclusions

This research presented a five-level adaptive ReportInterval technique for 5G NR that detects low-mobility devices from the statistical characteristics of RSRP and SINR. A low-mobility criterion combines four conditions—the variances and gradients of RSRP and SINR—through a logical AND operation, and a two-stage hysteresis mechanism, consisting of a 3 dB margin along with a 2 s holding duration, inhibits unwanted level transitions. Under the 3GPP TR 38.901 UMi channel model, with 200 UEs over 100 Monte Carlo runs, the algorithm achieves a classification accuracy of 91.32% and a sensitivity of 98.77%, yields an estimated reduction of 10.87% in average UE power consumption (based on the DRX energy model in Equation (14)), and a 51.09% reduction in the network-wide MeasurementReport count, with no degradation observed in the single-cell scenario considered here (inter-cell handover analysis under multi-cell conditions is identified as

future work). The hysteresis mechanism reduces level transitions by a factor of 31.33, which is among the largest reductions reported for comparable systems.

The technique runs at $O(n)$ cost on the gNB, where n is the small observation-window sample count, operates within the existing 3GPP standard, and is deployable with a gNB-side software update alone. It is slice-agnostic, independent of application-layer traffic, and coexists with the existing 3GPP safety mechanisms, including event-triggered reporting, Radio Link Monitoring, and the T310/T311 timers, all of which remain fully active regardless of the specified ReportInterval. These attributes make the algorithm a viable option for IoT-intensive 5G networks, especially in scenarios where battery replacement logistics are costly [29].

Future endeavors include multi-cell assessment with inter-cell interference management, expansion to the FR2 band, a hysteresis-informed temporal classification analysis, machine-learning-driven threshold optimization, and empirical validation on an operational 5G NR network.

Author Contributions: D.D.: conceptualization, method, supervision, project administration; N.S.: writing, original draft preparation, method, software; M.S.: validation, investigation, resources; C.W.: data curation, visualization; H.Z.: formal analysis, review and editing; K.S.: writing, review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The MATLAB source code and simulation datasets used to generate the results presented in this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. 3GPP. NR; NR and NG-RAN Overall Description; Stage 2; TS 38.300, v17.7.0; 3GPP: Sophia Antipolis, France, 2023.
2. Dahlman, E.; Parkvall, S.; Sköld, J. *5G NR: The Next Generation Wireless Access Technology*, 2nd ed.; Academic Press: London, UK, 2020.
3. Lin, K.-H.; Liu, H.-H.; Hu, K.-H.; Huang, A.; Wei, H.-Y. A Survey on DRX Mechanism: Device Power Saving from LTE and 5G New Radio to 6G Communication Systems. *IEEE Commun. Surv. Tutor.* **2023**, *25*, 156–183. [\[CrossRef\]](#)
4. 3GPP. NR; Radio Resource Control (RRC); Protocol Specification; TS 38.331, v17.5.0; 3GPP: Sophia Antipolis, France, 2023.
5. Ericsson. *Ericsson Mobility Report*; Ericsson: Stockholm, Sweden, 2025.
6. GSMA Intelligence. *The Mobile Economy 2025*; GSMA: London, UK, 2025.
7. Lee, J.; Lee, J. Prediction-Based Energy Saving Mechanism in 3GPP NB-IoT Networks. *Sensors* **2017**, *17*, 2008. [\[CrossRef\]](#) [\[PubMed\]](#)
8. Lopez-Perez, D.; De Domenico, A.; Piovesan, N.; Xinli, G.; Bao, H.; Song, Q.; Debbah, M. A Survey on 5G Radio Access Network Energy Efficiency: Massive MIMO, Lean Carrier Design, Sleep Modes, and Machine Learning. *IEEE Commun. Surv. Tutor.* **2022**, *24*, 653–697. [\[CrossRef\]](#)
9. 3GPP. *Study on Support of Reduced Capability NR Devices*; TR 38.875, v17.0.0; 3GPP: Sophia Antipolis, France, 2022.
10. Kadan Veedu, S.N.; Mozaffari, M.; Høglund, A.; Yavuz, E.A.; Tirronen, T.; Bergman, J.; Wang, Y.-P.E. Toward Smaller and Lower-Cost 5G Devices with Longer Battery Life: An Overview of 3GPP Release 17 RedCap. *arXiv* **2022**, arXiv:2203.05634.
11. Tayyab, M.; Hailu, S.; Jarvela, R.; Kolehmainen, N.; Gursu, H.M. RRM Relaxation in Connected State for Reduced Capability (RedCap) NR UEs. In Proceedings of the 17th International Symposium on Wireless Communication Systems (ISWCS), Berlin, Germany, 6–9 September 2021; pp. 1–6. [\[CrossRef\]](#)
12. Beschastnyi, V.; Ostrikova, D.; Moltchanov, D.; Gaidamaka, Y.; Koucheryavy, Y.; Samouylov, K. Comparison of Energy Conservation Strategies for 5G NR RedCap Service in Industrial Environment. *Comput. Netw.* **2024**, *254*, 110815. [\[CrossRef\]](#)
13. Bruhn, P.; Bassi, G. Machine Learning Based C-DRX Configuration Optimization for 5G. In Proceedings of the Mobile Communication—Technologies and Applications, 25th ITG-Symposium, Osnabrueck, Germany, 9–11 November 2021; pp. 1–6.
14. Lazrek, H.; El Ferindi, H.; Zouiten, M.; Moumen, A. Enhancing Energy Efficiency in 5G Networks through AI-Driven Dynamic Discontinuous Reception. *Discov. Comput.* **2025**, *28*, 245. [\[CrossRef\]](#)

15. Biswas, M.; Chakraborty, A.; Palit, B. A Kalman Filter Based Low Complexity Throughput Prediction Algorithm for 5G Cellular Networks. *IEEE Trans. Veh. Technol.* **2024**, *73*, 7089–7101. [[CrossRef](#)]
16. Jayawardana, P.A.D.N.; Obaid, H.; Yesilyurt, T.; Tan, B.; Lohan, E.S. Machine-Learning-Based LOS Detection for 5G Signals with Applications in Airport Environments. *Sensors* **2023**, *23*, 1470. [[CrossRef](#)] [[PubMed](#)]
17. Yao, X.; Perez Yuste, A. A ML-Based Resource Allocation Scheme for Energy Optimization in 5G NR. *Sensors* **2025**, *25*, 4978. [[CrossRef](#)] [[PubMed](#)]
18. Yu, J.-H.; Choi, Y.-J.; Seo, S.-H.; Choi, S.-G.; Jeong, H.-Y.; Kim, J.-E.; Baek, M.-S.; You, Y.-H.; Song, H.-K. Improved Connected-Mode Discontinuous Reception (C-DRX) Power Saving and Delay Reduction Using Ensemble-Based Traffic Prediction. *Mathematics* **2025**, *13*, 974. [[CrossRef](#)]
19. Kumar, M.; Gautam, A. A Beamforming-Aided DRX Mechanism for Traffic-Aware Energy Optimization in 5G NR-U. *Int. J. Commun. Syst.* **2025**, *38*, e70251. [[CrossRef](#)]
20. Khalifa, O.; Kouzayha, N.; Hussaini, M.A.; ElSawy, H.; Al-Harathi, N.; Elmirghani, J.; Hanif, M.; Al-Naffouri, T.Y. Energy Conservative Data Aggregation for IoT Devices: An Aerial Wake-Up Radio Approach. *IEEE Internet Things J.* **2024**, *11*, 2006–2026. [[CrossRef](#)]
21. 3GPP. NR; *Physical Layer Measurements*; TS 38.215, v17.3.0; 3GPP: Sophia Antipolis, France, 2023.
22. 3GPP. NR; *Physical Layer Procedures for Data*; TS 38.214, v17.6.0; 3GPP: Sophia Antipolis, France, 2023.
23. 3GPP. *Study on UE Power Saving in NR*; TR 38.840, v16.0.0; 3GPP: Sophia Antipolis, France, 2019.
24. 3GPP. NR; *Medium Access Control (MAC) Protocol Specification*; TS 38.321, v17.4.0; 3GPP: Sophia Antipolis, France, 2023.
25. 3GPP. *Study on Channel Model for Frequencies from 0.5 to 100 GHz*; TR 38.901, v17.0.0; 3GPP: Sophia Antipolis, France, 2022.
26. 3GPP. NR; *Requirements for Support of Radio Resource Management*; TS 38.133, v17.7.0; 3GPP: Sophia Antipolis, France, 2023.
27. 3GPP. NR; *Physical Layer Procedures for Control*; TS 38.213, v17.6.0; 3GPP: Sophia Antipolis, France, 2023.
28. 3GPP. *System Architecture for the 5G System (5GS)*; TS 23.501, v17.9.0; 3GPP: Sophia Antipolis, France, 2023.
29. Sorensen, A.; Wang, H.; Remy, M.J.; Kjettrup, N.; Sorensen, R.B.; Nielsen, J.J.; Popovski, P.; Madueno, G.C. Modelling and Experimental Validation for Battery Lifetime Estimation in NB-IoT and LTE-M. *arXiv* **2021**, arXiv:2106.13286.

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