UNICO I+D Project

Grant No. TSI-0630002021-138

6G-SORUS

SORUS-RIS-A2.1-E3

“Final Energy Measurements”

Financiado por la Unión Europea NextGenerationEU

GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES

Plan de Recuperación, Transformación y Resiliencia

UNICO I+D

Document properties

Document number	SORUS-RIS-A2.3-E3
Document title	Final Energy Measurements
Document responsible	Jesus Omaña Iglesias (TID)
Document editor	Jesus Omaña Iglesias
Editorial team	Jesus Omaña Iglesias, Andre Felipe Zanella
Target dissemination level	Public
Status of the document	Final
Version	1.0
Delivery date	19-06-2025
Actual delivery date	19-06-2025

Production properties

Reviewers	Jesus Omaña Iglesias (TID)

Disclaimer

This document has been produced in the context of the SORUS-RIS Project. The research leading to these results has received funding from the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU through the UNICO 5G I+D programme.

All information in this document is provided “as is" and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Content

List of Figures 5
List of Tables 6
Resumen ejecutivo 7
Abstract 8
1. Introduction 9
2. Effect of RATs video streaming application on energy consumption 11
2.1. Impact on power consumption due to the RATs 12
2.2. Effect on smartphone generation 13
3. Methodology for energy measurement at the RAN 14
3.1. Energy-saving strategies and trial set up 15
3.2. Measurement data 18
4. Cells affected during the trial 20
5. Gains in Energy savings during the trial 23
6. Impact on performance during the trials 26
7. Impact on. CO2 emissions 29
Summary and Conclusion 32
References 33

List of Figures

Figure 1 EFFECT OF CACHING ON POWER CONSUMPTION 11

Figure 3 ENERGY CONSUMPTION FOR video streaming APPS ON DIFFERENT RATS. 13

Figure 4 Data consumption (left) power consumption (right) 13

Figure 5 ENERGY CONSUMPTION for two generation of smartphones 14

Figure 6 Example of energy-saving strategy 17

Figure 7 left: Dense region; right: Sparse region. Percentage of capacity cells affected by their energy-saving strategies 21

Figure 8 Hourly time series of the mean cell downtime. Tope: dense region, and bottom sparse region 22

Figure 9 Top - percentage of energy savings; bottom - distribution of energy savings during day and night periods 25

Figure 10 Transmission Time Interval (TTI) utilization and user throughput before a cell enters sleep mode 28

Figure 11 Cells transition percentage for each strategy 28

Figure 12 Estimated carbon intensity in gCO2eq/kWh and energy consumption in the sparse region during the week of the second trial 31

List of Tables

Table 1 Power Consumption Comparison (Using Cache) 11
Table 2 Power Consumption Comparison (Without using Cache) 12
Table 3 Energy saving strategies by the MNO 15
Table 4 Description of the regions under study 18

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Resumen ejecutivo

La eficiencia energética se ha convertido en un objetivo clave para la industria de las telecomunicaciones móviles, especialmente en la Radio Access Network (RAN), responsable de más del 70 % del consumo total. Aunque existen muchas propuestas complejas para apagar portadoras infrautilizadas, en la práctica se usan estrategias simples basadas en umbrales, activadas solo de noche y sin evaluaciones reales de su impacto. En este estudio, analizamos cinco políticas de apagado de celdas basadas en umbrales fijos desplegadas en una red comercial. Los resultados ofrecen una visión inédita sobre el ahorro energético real y el impacto en los usuarios, y muestran que sin afectar la experiencia de usuario, el ahorro tiene un límite claro. Esto respalda la necesidad de enfoques más flexibles para mejorar la sostenibilidad del RAN.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Abstract

Energy efficiency has become a key objective for the mobile telecommunications industry, especially in the Radio Access Network (RAN), which accounts for over 70% of total energy consumption. While many complex solutions have been proposed to switch off underutilized carriers, real-world deployments mostly rely on simple threshold-based strategies, typically activated only at night and lacking proper impact evaluation. In this study, we analyze five fixed threshold-based cell sleep policies deployed in a commercial network. Our results provide unprecedented insights into actual energy savings and user impact, revealing a clear limit to energy reduction when user experience cannot be compromised. This highlights the need for more flexible approaches to improve RAN sustainability.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

1. Introduction

This report provides a comprehensive analysis of energy consumption across the mobile network stack, spanning both user-device and infrastructure levels. It begins by examining how different Radio Access Technologies (RATs) affect smartphone energy usage during video streaming scenarios. This initial analysis also considers how smartphone generation influences energy consumption, offering insights into end-user behavior.

Building on this foundation, the main focus of the report shifts to the mobile network infrastructure—specifically, the Radio Access Network (RAN)—where energy-saving policies are assessed in a live production environment

Mobile telecommunication network operators (MNOs) are currently facing significant challenges related to energy efficiency, driven by escalating energy costs and growing global awareness of the climate emergency (H. Technology, 2020). The Radio Access Network (RAN) is the primary target for energy savings, as it is responsible for over 70% of an operator's total energy expenditure (G. Association, 2024). Consequently, reducing energy consumption within the RAN is a key objective, offering substantial environmental and economic benefits.

Modern RANs possess inherent redundancy, designed to manage significant fluctuations in user demand across both time and geography (GSMA, 2021). This layered architecture presents an opportunity for dynamic management, allowing unneeded equipment, such as individual symbols, radio channels, or entire carriers, to be temporarily switched off to reduce energy costs (S. Mishra, 2022). However, implementing such energy-saving measures in production networks is complex. While the scientific community has proposed numerous sophisticated solutions, real-world deployments largely depend on more basic, threshold-based strategies, typically enabled only overnight. A major impediment is the conflict between sustainability goals and the strict performance requirements of mobile networks, as any degradation in user Quality of Experience (QoE) is largely unacceptable due to potential customer churn and negative feedback.

Given the scarcity of real-world evaluations, this report, which is heavily based on a work we have recently presented at INFOCOM 2025 (Orlando E Martínez-Durive, 2025) , provides critical insights by benchmarking five distinct fixed threshold-based cell sleep policies that were deployed and tested in a large-scale production network across extensive geographical regions. The study rigorously assesses the effectiveness of these policies in terms of actual energy savings and their trade-offs against user performance, using metrics such as cell downtime, energy consumption, and user throughput. The findings underscore that current fixed-threshold strategies face a clear limitation in achieving greater energy savings without some level of user experience degradation, thus supporting the need for more adaptable and dynamic approaches to RAN sustainability.

In addition to the proprietary datasets—which cannot be shared due to confidentiality constraints—we also leveraged recently released public datasets. These public datasets were used to validate some of the values reported in our figures (particularly those related to energy consumption) and to ensure that the performance indicators we relied on align with those commonly used in the literature. While the proprietary data remains unavailable, similar insights can be reproduced using the public datasets available at: https://github.com/tsinghua-fib-lab/NetData , released by (Y. Ma, 2024)

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

2. Effect of RATs video streaming application on energy consumption

Following our examination of the other application types in the previous deliverable, we now delve into the realm of video streaming apps. In this study, we focused on three popular platforms: YouTube Shorts, TikTok, and Instagram Reels. Similar to our previous analysis, we investigated the impact of caching on power consumption to understand its effect overall. Just like with social network apps, enabling caching led to a reduction (although small) in energy usage. Figure 23 presents our findings, while Table 12 provides various statistics on power consumption with caching, and Table 13 displays results without caching.

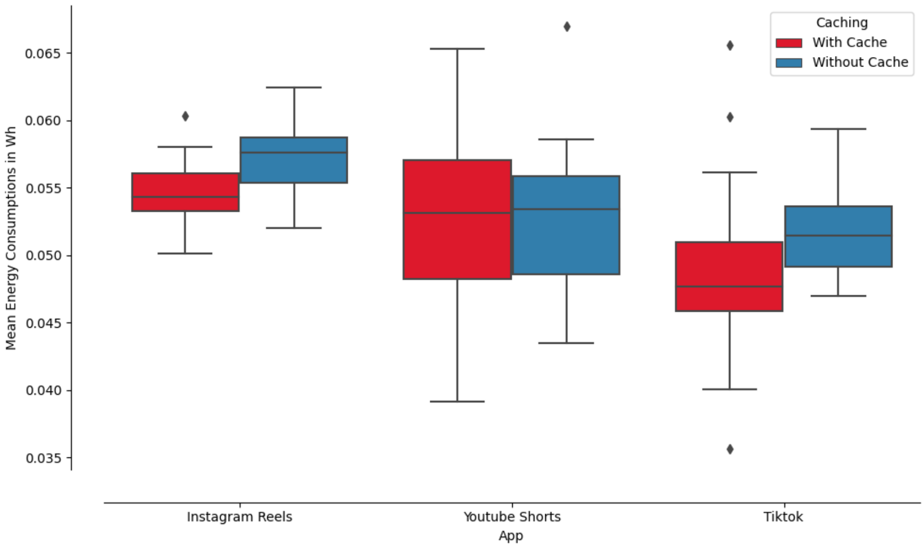

FIGURE 1 EFFECT OF CACHING ON POWER CONSUMPTION

TABLE 1 POWER CONSUMPTION COMPARISON (USING CACHE)

App	Mean (Wh)	Median (Wh)	Standard Deviation	Val. Min. (Wh)	Q1 (Wh)	Q3 (Wh)	Val. Max (Wh)
Instagram Reels	0.0546	0.0543	0.0020	0.0501	0.0532	0.0560	0.0603
YouTube Reels	0.0527	0.0531	0.0069	0.0356	0.0482	0.0570	0.0655
TikTok	0.0484	0.0477	0.0055	0.0391	0.0458	0.0509	0.0652

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

TABLE 2 POWER CONSUMPTION COMPARISON (WITHOUT USING CACHE)

App	Mean (Wh)	Median (Wh)	Standard Deviation	Val. Min. (Wh)	Q1 (Wh)	Q3 (Wh)	Val. Max (Wh)
Instagram Reels	0.0546	0.0543	0.0020	0.0501	0.0532	0.0560	0.0603
YouTube Reels	0.0527	0.0531	0.0069	0.0391	0.0482	0.0570	0.0655
Tiktok	0.0484	0.0477	0.0055	0.0356	0.0458	0.0509	0.0652

2.1. Impact on power consumption due to the RATs

The results for this type of application once again differ from those observed for social networking apps. Specifically, Figure 3 illustrates the findings for the three apps under study. On average, we observe that WiFi is the most energy-efficient RAT, followed by 3G (with an 8.79% increase), then 4G (with a 12.21% increase), and finally, 5G emerges as the most power-hungry RAT, with an 18.32% increase. We are uncertain about the change in order and plan to conduct further experiments to gain more definitive insights. We suspect that this variation may be attributed to the difference in data size; notably, video streaming consumes significantly more data than social networking apps, as depicted in Figure 4. Moreover, in terms of overall power consumption, TikTok appears to be the most energy-efficient app, closely followed by YouTube Shorts (which consumes only 1.54% more). Conversely, Instagram stands out as the most power-hungry app, consuming 10.23% more energy

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

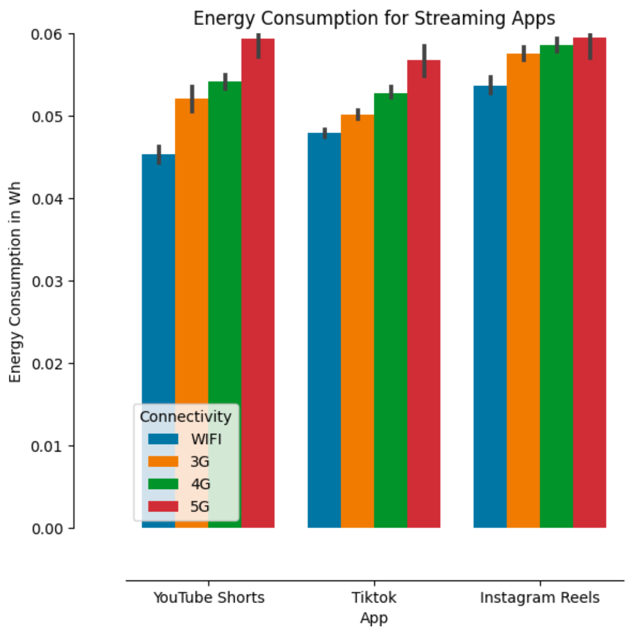

FIGURE 2 ENERGY CONSUMPTION FOR VIDEO STREAMING APPS ON DIFFERENT RATS.

THE ERROR BARS PRESENT THE VARIABILITY OVER 10 INDEPENDENT EXPERIMENTS.

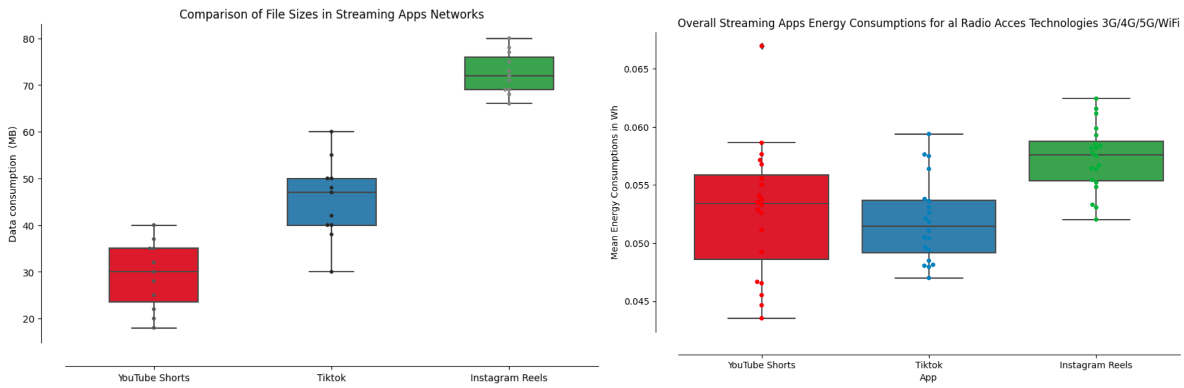

FIGURE 3 DATA CONSUMPTION (LEFT) POWER CONSUMPTION (RIGHT)

2.2. Effect on smartphone generation

Studying the results in Figure 5, no matter the video streaming application used, newer generation phones tend to be more power-hungry than older ones. In this

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

specific instance, the newer Pixel 5 consumes an average of 15.41% more power compared to its predecessor, the Pixel 4.

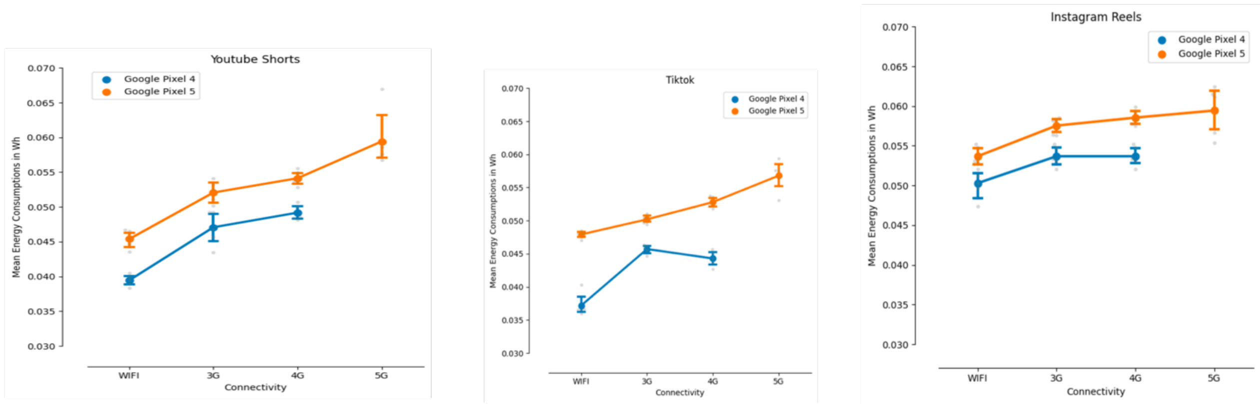

FIGURE 4 ENERGY CONSUMPTION FOR TWO GENERATION OF SMARTPHONES

3. Methodology for energy measurement at the RAN

The previous sections and the two previous deliverables focused on the energy consumption and performance of user equipment (UE) across a variety of applications and radio access technologies (RATs). These measurements provide valuable insights into how device behavior and application demands interact with the network, revealing patterns of energy use and performance at the edge.

To complete the picture, we now shift our focus from end-user devices to the mobile network infrastructure—specifically, the Radio Access Network (RAN). This transition is essential: while smartphones account for the energy consumption perceived by users, RAN components—particularly base stations—are responsible for the majority of energy expenditure incurred by mobile network operators, typically representing around 70% of total consumption.

In this section, we present the methodology used to benchmark five fixed threshold-based cell sleep policies within a large-scale, live production network operated by a top-tier Mobile

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Network Operator (MNO) in Western Europe, serving over 40 million connected devices. The trials were conducted over several weeks in February and March 2023, with each energy-saving strategy deployed for one week (see Table 1). The primary goal was to quantify actual energy savings while assessing any associated impact on user experience.

TABLE 3 ENERGY SAVING STRATEGIES BY THE MNO

Strategy	Deployment time window	Policy ON	Policy OFF
Night-loose	23h - 6h (night)	5%	10%
Night-strict	23h - 6h (night)	10%	20%
Full-loose	24h (all day)	5%	10%
Full-moderate	24h (all day)	7%	12%
Full-strict	24h (all day)	10%	15%

3.1. Energy-saving strategies and trial set up

Figure 6 illustrates the general behavior of the commercial energy-saving solutions tested by the MNO. All solutions rely on continuous monitoring of the Physical Resource Block (PRB) utilization, i.e., the portion of fundamental radio resource units available in a cell that are allocated to users for data transmission or reception. PRB utilization is tracked at the level of each power group, i.e., the set of all cells covering the same geographical area over different frequency bands.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

The tested solutions apply dynamic cell sleeping strategies based on the time evolution of the average PRB utilization in cells of a power group. The energy-saving strategies in this study are defined by three main parameters, which control when and how cells enter or exit sleep mode:

- **Policy Deployment Time Window**: This refers to the specific time period during which the energy-saving policy is enabled (see also the first column in Table 1). For instance, some strategies are active only during night hours (23:00 to 06:00), while others operate throughout the entire 24-hour day.
- **ON/OFF Load Thresholds**: These are the PRB utilization levels that trigger a cell to go into sleep mode (Policy ON) or wake up (Policy OFF). When the average PRB utilization in a power group falls below the "ON" threshold, a cell is selected to enter sleep mode (see column 3 of Table 1). Conversely, if utilization rises above the "OFF" threshold while cells are sleeping, the cell is reactivated (see column 4 of Table 1). Different strategies use varying ON/OFF thresholds, with higher thresholds leading to more aggressive energy saving, potentially impacting user performance.
- **Time to Trigger**: This is a built-in delay designed to prevent frequent and erratic switching of cells due to rapid load fluctuations. For these trials, the delay was set to 10 minutes before a cell enters sleep mode and 5 minutes before it reactivates. This helps ensure a smoother transition for users.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

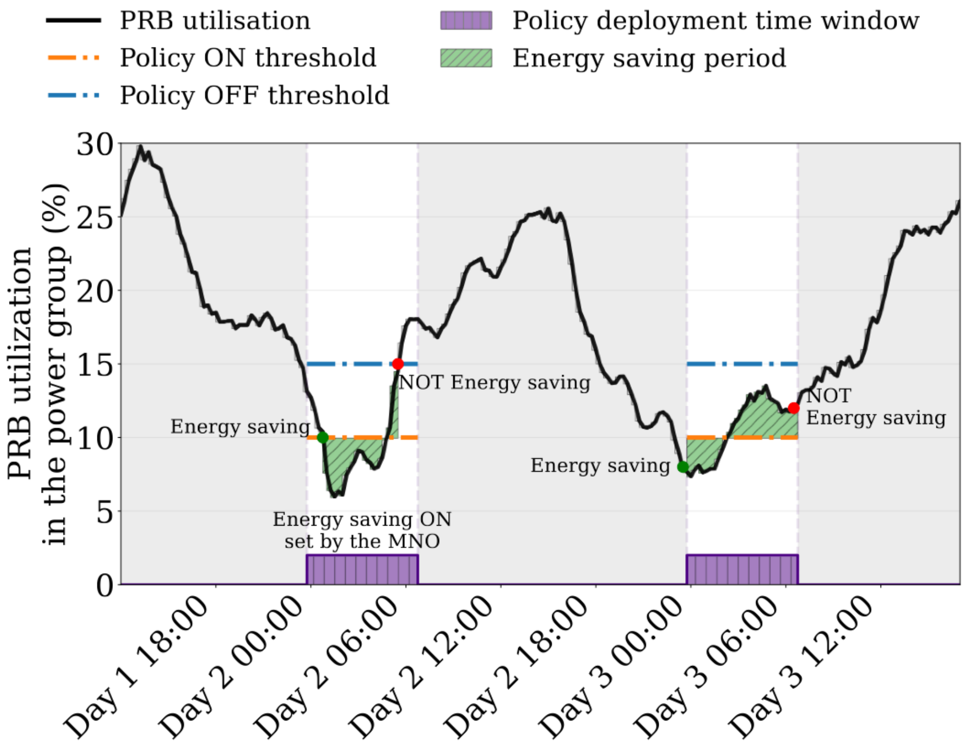

FIGURE 5 EXAMPLE OF ENERGY-SAVING STRATEGY

The core of the methodology involved deploying and evaluating dynamic cell sleeping strategies. These strategies are based on the continuous monitoring of Physical Resource Block (PRB) utilization within each power group, which encompasses all cells covering a specific geographical area across different frequency bands. Five distinct energy-saving policies were rigorously tested (see Table 1).

The trials specifically focused on LTE capacity cells, which account for 23.05% of tested cells in the Dense region and 27.43% in the Sparse region. Coverage cells (e.g., those on low-frequency bands) were kept continuously active to guarantee service availability across the entire geography, while other Radio Access Technologies (RATs) such as 2G, 3G, and 5G were excluded from the trial. The MNO performed these trials across two geographically diverse regions, categorized as "Dense" and "Sparse," characterized by different population and cell deployment densities. The Dense region spanned 189.28 km^2 with 5.41 LTE cells/sq. km, while the Sparse region covered 4986.82 km^2 with 0.21 LTE cells/sq. km. Details of the regions under study can be seen in Table 2.

TABLE 4 DESCRIPTION OF THE REGIONS UNDER STUDY

Region name	Area KM^2	LTE share	LTE density (cells/S1. Km.)	LTE cells/site	Capacity cells (%)
Dense region	189.28	71.51%	5.41	11.0 ± 6.10	23.05%
Sparse region	4986.82	54.14%	0.21	8.5 ± 5.98	27.43%

3.2. Measurement data

For evaluation, cell-level measurements from LTE cells in the two study regions were collected throughout the entire trial period. The data sources included:

- **Radio Access Network (RAN) Deployment Inventory**: This provided daily updated cell-level information, including location (latitude/longitude), eNodeB ID, orientation (azimuth), tilt, antenna manufacturer and model, and carrier frequency channels.
- **Radio Access Network Key Performance Indicators (KPIs)**: The MNO's monitoring infrastructure collected over 160 cell-level KPIs for all deployed RATs (from 2G to 5G NR) every 15 minutes, which were then averaged over hourly intervals. For this study, three main KPIs were used: (i) cell availability, measuring the time a cell was active down to seconds; (ii) cell load, expressed

Financiado por la Unión Europea NextGenerationEU · Gobierno de España · Ministerio de Asuntos Económicos y Transformación Digital · Secretaría de Estado de Telecomunicaciones e Infraestructuras Digitales · Plan de Recuperación, Transformación y Resiliencia · UNICO I+D

as downlink PRB utilization; and (iii) average user throughput, indicating the average data rate served by the cell.

- **Mobility Management Signaling**: Control-plane signaling messages from the 4G Evolved Packet Core's (EPC) Mobility Management Entity (MME) were collected with millisecond timestamps. This data captured all mobility-related events, such as attach, detach, and handover, and was instrumental in analyzing user transitions during cell sleep or wake-up events to identify "fallback cells" – those that absorbed the load from sleeping cells. A data-driven approach with a symmetric time window of 30 minutes (δstart = δend = 30min) was applied to ensure that user transfers were attributed to the activation of the sleeping mode.

Energy consumption was estimated using an operational model (Y. Ma, 2024) ($E_{BS} = E_{BBU} + \sum_i E_{RRU}^i$) that considers both the constant energy consumed by baseband units and the load-dependent energy of radio units (E_{RRU}^i). When a cell enters sleep mode, the E_{RRU}component is considered zero. To ensure a fair comparison across strategies and account for traffic variability, energy savings were reported as a percentage relative to a "no-sleep" benchmark scenario, where capacity cells never enter sleep mode and their radio units always consume power, even at zero load.

To quantify the impact on user performance, the study focused on the average user throughput at the cell level. Fallback cells were analyzed for changes in their Transmission Time Interval (TTI) utilization and user average throughput. Thresholds were established to categorize performance: a TTI utilization of 45% (third quartile of distribution) and an average user throughput below 8 Mbps (based on literature, MNO operations teams, and statistical analysis) were used to identify critical performance degradation. Based on these thresholds, fallback cell transitions were classified into "Critical," "Saturate," "Recovery," and "Normal" states to understand the impact of cell sleep decisions. Details on what the exact definition of these states will be given in section 5.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

4. Cells affected during the trial

During the weekly trial, the different energy-saving strategies affected varying fractions of capacity cells—defined as those that entered sleep mode at least once—and resulted in different levels of average downtime. Nighttime-only strategies (Night-loose and Night-strict) exhibited significantly lower total downtime compared to full-day strategies. However, they impacted a slightly larger number of cells overall. In contrast, full-day strategies (24h) affected more cells, and the average downtime increased with the aggressiveness of the strategy, determined by the ON/OFF thresholds.

Figure 7 illustrates these effects by showing both the percentage of affected cells per strategy and their average downtime, including the standard deviation. As the ON/OFF thresholds become more aggressive, the time cells remain in sleep mode increases. For instance, in the Dense region, the Full-loose strategy (5% ON / 10% OFF load thresholds) leads to an average downtime of 40.6%, while the Full-strict strategy (10% ON / 15% OFF) pushes this value up to 51.5%.

Interestingly, even though Night-loose and Full-loose use the same ON/OFF load thresholds, Night-loose affects a slightly larger fraction of cells. Upon checking with the MNO operations team, this was identified as a symptom of misbehavior in the nighttime energy-saving implementation: some cells were not properly waking up, even when the traffic load exceeded the OFF threshold. Despite this issue, the average downtime across cells still followed the expected trend.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

To complement this analysis, Figure 8 shows the hourly evolution of the average downtime across all strategies. Clear day-night patterns emerge, with all strategies exhibiting peaks in downtime during the night, when traffic demand is low. Full-day strategies consistently show higher downtime across the 24-hour period as they become stricter, with both nighttime peaks and daytime valleys becoming more pronounced. Additionally, the Night-strict strategy shows longer sustained periods of cell sleep during the night, driven by its higher OFF threshold (20% PRB utilization), which causes cells to remain in sleep mode for extended periods while the policy is active.

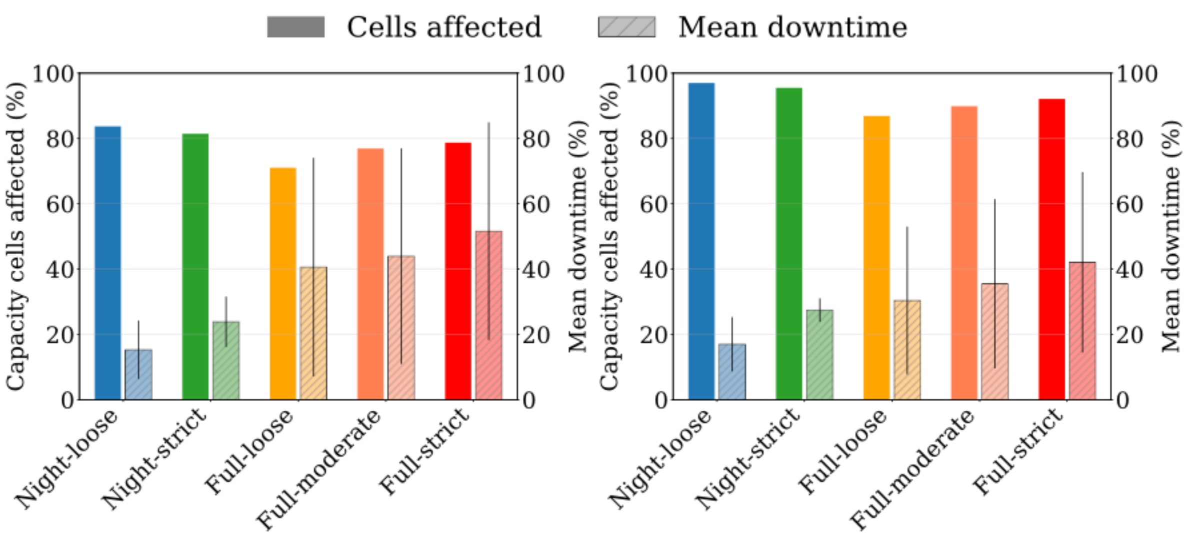

FIGURE 6 LEFT: DENSE REGION; RIGHT: SPARSE REGION. PERCENTAGE OF CAPACITY CELLS AFFECTED BY THEIR ENERGY-SAVING STRATEGIES

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

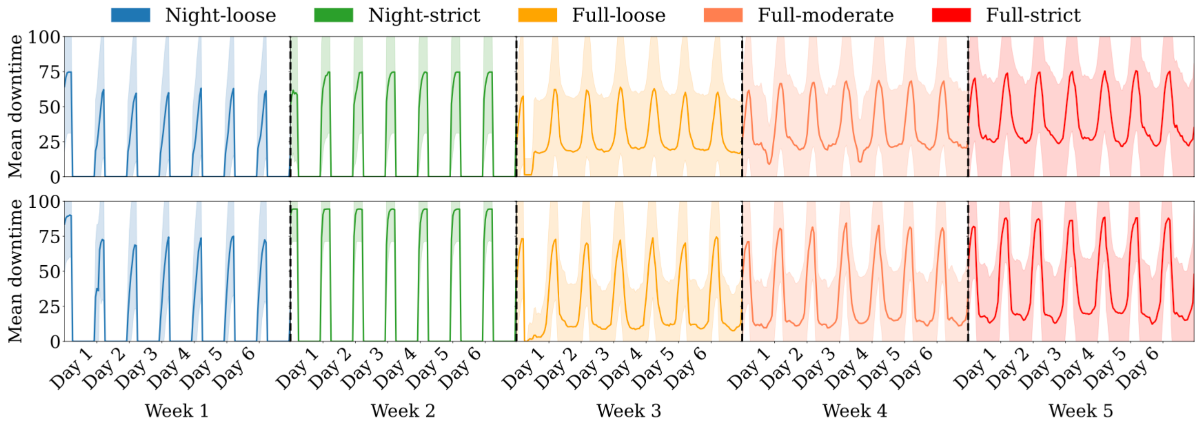

FIGURE 7 HOURLY TIME SERIES OF THE MEAN CELL DOWNTIME. TOPE: DENSE REGION, AND BOTTOM SPARSE REGION

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

5. Gains in Energy savings during the trial

During the weekly trials, we observed some variability in aggregated traffic load across the Dense and Sparse regions, driven by external factors such as local holidays, social events, or sports broadcasts. Despite this, traffic fluctuations remained relatively modest—3.2% in the Dense region and 3.7% in the Sparse region. Since each strategy was tested for an entire week, we mitigated variability caused by differences between weekdays and weekends.

To fairly assess energy savings, we use a baseline model that simulates a no-sleep scenario: all capacity cells remain active throughout the trial, consuming static circuit power (γ) even when carrying zero traffic. In contrast, when a cell enters sleep mode under an energy-saving strategy, its power consumption drops to zero (γ = 0, R_{PRB} = 0), where R_{PRB} is the fraction of PRBs used. It represents the load on a cell, specifically the portion of fundamental radio resource units available within a cell that are allocated for user data transmission or reception. When a cell enters sleep mode, the E_{RRU} (energy consumed by the radio unit) component is considered zero, which implies that both the R_{PRB} and γ (fixed circuit power) are also considered zero. This model allows us to compute the relative energy savings of each strategy, isolating them from external traffic variations.

Figure 9 (top) reports the percentage of energy saved with respect to the no-sleep baseline. All strategies yield noticeable gains, with savings clearly linked to their time windows (night-only vs full-day) and the aggressiveness of their ON/OFF thresholds. Notably, the Full-strict strategy achieves the highest savings—34.5% in the Dense region and 30.2% in the Sparse region.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

To further dissect these gains, we analyze energy savings during daytime (06:00–23:00) and nighttime (23:00–06:00). As expected, most strategies show stronger savings at night, when traffic is lower and both Night-loose and Night-strict are active. Figure 9 (bottom) shows the per-cell distribution of energy savings in both time periods. Median values confirm that most savings occur overnight.

Interestingly, full-day strategies begin to show substantial energy savings during daytime as their aggressiveness increases. While median savings during the day are similar across strategies, the upper quartile of cells (top 25%) achieves significantly higher savings with more aggressive thresholds. For example, during daytime in the Dense region, the top-25% of cells reach savings of 15.8%, 21.8%, and 42.7% under the Full-loose, Full-moderate, and Full-strict strategies, respectively. In the Sparse region, these values are 8.8%, 11.2%, and 17.2%.

These findings highlight that, although nighttime remains the dominant window for energy savings, a non-negligible subset of cells can achieve greater savings during the day, especially under stricter full-day policies. Note that these comparisons depend on our definition of nighttime (7 hours) versus daytime (17 hours).

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

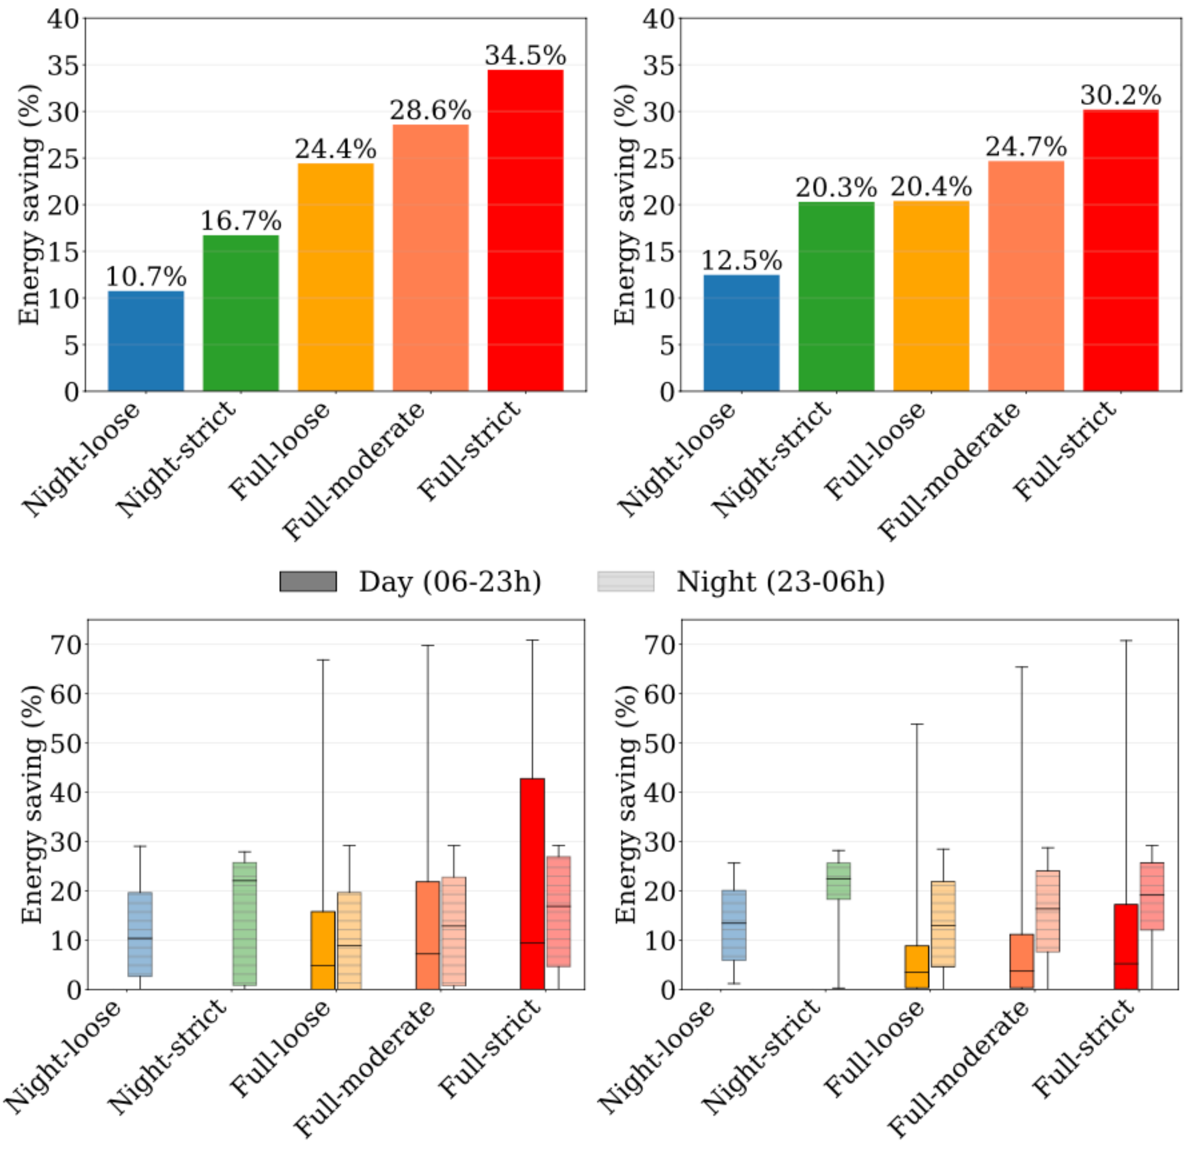

FIGURE 8 TOP - PERCENTAGE OF ENERGY SAVINGS; BOTTOM - DISTRIBUTION OF ENERGY SAVINGS DURING DAY AND NIGHT PERIODS

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

6. Impact on performance during the trials

Our goal in this section is to identify fallback cells that were negatively impacted after nearby capacity cells entered sleep mode. In particular, we focus on cases where the average user throughput in a fallback cell decreased significantly, which could indicate a degraded user experience due to the absorption of traffic from sleeping neighbors.

To isolate performance degradations caused by increased demand (as opposed to low user activity), we jointly analyze average user throughput and cell utilization. Specifically, we use TTI utilization as a proxy for cell load, analogous to PRB utilization. This allows us to distinguish whether a throughput drop is due to congestion or simply due to low traffic activity.

Figure 10 shows the relationship between TTI utilization and average user throughput for fallback cells, measured after a sleep event of a nearby capacity cell in the same site and sector. Each point represents a fallback cell, and we apply two thresholds to classify the state of the cell:

- **TTI utilization threshold**: 45%, chosen as the 75th percentile of the TTI utilization distribution during the trial period.
- **User throughput threshold**: 8 Mbps, derived from a combination of literature, operational team input, and empirical data.

These thresholds do not mark hard failures that would automatically trigger operator intervention but are indicative of significant degradation in user experience. Using these thresholds, we define four quadrants in the TTI-throughput space:

- **Quadrant I**: Low throughput, low utilization

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

- **Quadrant II**: Low throughput, high utilization
- **Quadrant III**: Normal throughput, high utilization
- **Quadrant IV**: Normal throughput, low utilization

We now provide more details on the state transitions of fallback cells before and after capacity cells go to sleep into four main classes:

- **Critical**: Throughput drops below 8 Mbps and utilization rises (e.g., transitions into Quadrant II from I, III, or IV). This indicates a potentially problematic impact from sleep mode.
- **Saturate**: The fallback cell was already in a low-throughput, high-utilization state and remains there (Quadrant II → Quadrant II). Sleep mode didn't cause new degradation but failed to help.
- **Recovery**: The cell improves its average throughput while remaining highly utilized (Quadrant II → I, III, or IV), suggesting adaptive resilience.
- **Normal**: The fallback cell maintains acceptable throughput before and after sleep, showing no performance degradation.

Figure 11 summarizes the percentage of fallback cells falling into each transition category under each policy. Two key insights emerge: (i) all policies lead to a mix of all four outcomes, and (ii) the share of Critical and Saturate transitions does not noticeably increase with more aggressive energy-saving policies. This suggests that performance degradation does not systematically worsen with stricter thresholds.

These findings point to the importance of cell-specific analysis and hint at the potential benefits of tailoring sleep strategies to local conditions rather than applying a uniform policy across the board.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

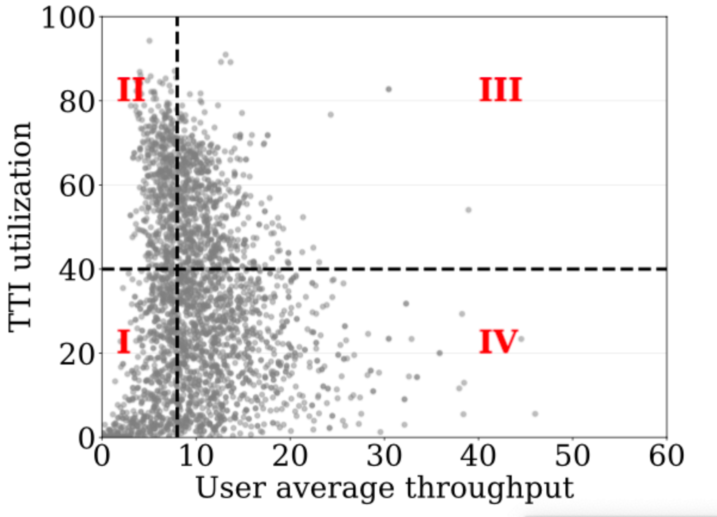

FIGURE 9 TRANSMISSION TIME INTERVAL (TTI) UTILIZATION AND USER THROUGHPUT BEFORE A CELL ENTERS SLEEP MODE

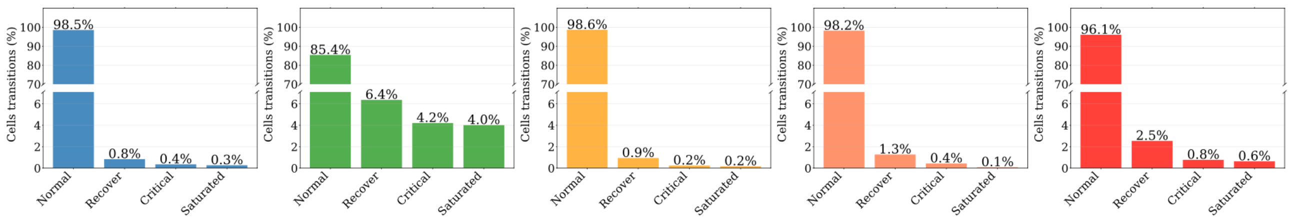

(a) Night-loose (b) Night-strict (c) Full-loose (d) Full-moderate (e) Full-strict

FIGURE 10 CELLS TRANSITION PERCENTAGE FOR EACH STRATEGY

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

7. Impact on. CO2 emissions

Although the primary objective of this work is to evaluate energy savings, we also assess the impact of our strategies on carbon dioxide (CO_2) emissions, with a focus on Scope 2 emissions, as defined by the Greenhouse Gas (GHG) Protocol (GHG Protocol, 2023) .

The GHG Protocol categorizes emissions into three scopes:

- **Scope 1** refers to direct emissions from sources owned or controlled by the organization, such as emissions from company vehicles or on-site fuel combustion.
- **Scope 2** encompasses indirect emissions from purchased energy, primarily electricity consumed by the organization.
- **Scope 3** includes all other indirect emissions, such as those originating from suppliers, logistics, or the end use of products and services.

The threshold-based energy-saving policies evaluated in this study directly impact Scope 2 emissions, which are the most relevant in our context given that mobile networks consume large amounts of electricity. We quantify these emissions in grams of CO_2 equivalent per kilowatt-hour (gCO_2eq/kWh), allowing us to estimate how energy savings translate into emission reductions.

While energy consumption and CO_2 emissions are inherently linked, the relationship between energy saved and emissions reduced is not always linear. One key reason is that energy-saving decisions typically do not account for the carbon intensity of the electricity supply at the time of consumption. That is, a reduction in electricity use is beneficial by default, but its true environmental impact depends on how that electricity was generated—

whether from carbon-intensive sources (e.g., coal or gas) or from low-carbon sources (e.g., wind, solar, or nuclear).

In the country considered for our study, fossil fuels still represent a non-negligible share of the energy mix, meaning that a significant portion of the electricity consumed is associated with high carbon emissions. Emission data were obtained by querying the electricity maps API (Maps, 2025). Therefore, reducing electricity use during periods when the grid is carbon-intensive has a disproportionately higher environmental benefit than reducing it during greener periods.

To illustrate this, Figure 12 reports the carbon intensity of electricity generation in the country during the week in which the Night-strict energy-saving policy was trialed. The data show a clear diurnal pattern, with lower emissions during the night and higher emissions during daylight hours—a result of increased daytime demand and variability in renewable energy availability. However, beyond this expected daily fluctuation, the figure also reveals significant variability across different days of the week.

This variability opens the door to carbon-aware energy-saving strategies that go beyond flat, threshold-based approaches. For instance, policies could be dynamically adjusted to become more aggressive during high-carbon periods (e.g., during daytime hours on days when renewables are scarce) and more relaxed during low-carbon periods. Such alignment would maximize CO_2 reductions per unit of energy saved, making the energy-saving strategy not only power-efficient but also climate-efficient.

In summary, while our current policies already achieve CO_2 reductions proportional to energy savings, there is a clear opportunity to enhance their climate impact by incorporating

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

real-time or forecasted carbon intensity data into the decision-making process. This would enable next-generation energy-aware and carbon-aware network management practices with substantially improved environmental outcomes.

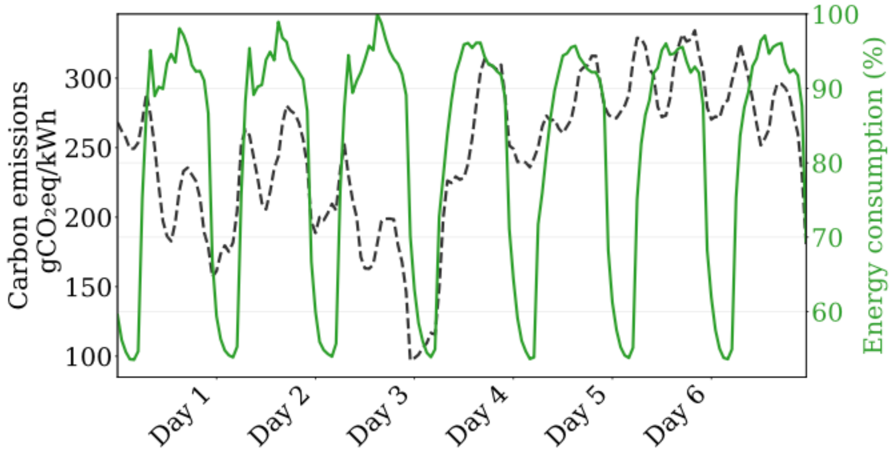

FIGURE 11 ESTIMATED CARBON INTENSITY IN GCO2EQ/KWH AND ENERGY CONSUMPTION IN THE SPARSE REGION DURING THE WEEK OF THE SECOND TRIAL

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Summary and Conclusion

In summary, we analyzed five fixed threshold-based cell sleep energy saving strategies deployed in a production network, examining various dimensions such as downtime, energy consumption, and impact on user performance. The unprecedented visibility into practical solutions for RAN sustainability lets us shed light on the performance and current limitations of these strategies, as well as provide recommendations for the design of new and more effective approaches that can still be deployed in real-world production-grade networks.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

References

H. Technology. (2020). *Green 5g: Building a sustainable world.*

G. Association. (2024). *Mobile net zero, state of the industry on climate action.*

GSMA. (2021). *Going green: Benchmarking the energy efficiency of mobil.* London, UK.

S. Mishra, Z. S. (2022). Second-level digital divide: A longitudinal study of mobile traffic consumption imbalance in france. *ACM Web Conference.*

Orlando E Martínez-Durive, J. S.-V. (2025). An Evaluation of RAN Sustainability Strategies in Production Networks. *IEEE International Conference on Computer Communications (INFOCOM).*

Y. Ma, T. L. (2024). Sustainable con- nections: Exploring energy efficiency in 5g networks. *International Conference on Emerging Networking EXperiments and Technologies, CoNEXT.*

GHG Protocol. (2023). *Greenhouse gas protocol.*

Maps, E. (2025). Obtenido de https://app.electricitymaps.com/map/72h/hourly

Financiado por la Unión Europea NextGenerationEU

GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES

Plan de Recuperación, Transformación y Resiliencia

UNICO I+D