Saving Energy in Cellular Networks with National Roaming: A Data-Driven Analysis

Syllas R. C. Magalhães, Suzan Bayhan, Geert Heijenk
University of Twente, The Netherlands, {s.magalhaes, suzan.bayhan, geert.heijenk}@utwente.nl

***Abstract*—Ensuring reliable and ubiquitous mobile network coverage is critical for user experience and competitiveness of mobile network operators (MNOs). To meet varying service demands, MNOs typically over-provision their infrastructure, leading to an increased energy consumption and operational costs. National roaming (NR), a form of cooperation between MNOs, enables infrastructure sharing to reduce redundancy in network infrastructure and improve efficiency by allowing each mobile user to be connected to a base station that provides the *best* connectivity irrespective of which MNO the user is subscribed to. In this paper, we analyse real-world data from three Dutch municipalities to assess NR's impact on operational power consumption as well as coverage and throughput relative to the current de-facto operation mode where MNOs operate in an isolated way with no cooperation (NC). We propose an energy-aware user association scheme (EA) with polynomial complexity for the studied NR scenario and compare it to a load-balancing, energy-oblivious alternative. Our results show that NR can reduce power consumption in the radio access network by up to 30%, reduce the fraction of disconnected population (FDP) by 95%, and raise the fraction of satisfied population (FSP) by 50%. Our analysis suggests that energy savings are most significant in dense, overprovisioned networks, while FDP and FSP gains are most notable in sparse deployments. Moreover, we observe that EA outperforms our comparison baseline for both NR and NC scenarios, which implies that current MNOs can decrease their power consumption significantly by following an energy-aware user association scheme that facilitates putting some more BSs in sleep mode without compromising FDP and FSP.**

I. Introduction

Maintaining ubiquitous mobile network coverage and reliable connection performance is crucial for MNOs to remain competitive and for users to experience a high-quality service. Additionally, regulatory bodies typically impose minimum requirements for service quality and coverage to ensure adequate network performance. For example, in the Netherlands, each MNO is required to provide an outdoor data rate of at least 8 Mbps currently and 10 Mbps after 2026 in every region and mobile network coverage of at least 98% of the surface area in each province [1]. As a result, MNOs must deploy sufficient infrastructure and often over-provision it to accommodate the diverse demands and varying loads of different applications within the network. As mobile networks expand to meet increasing coverage and performance requirements, their energy consumption has become a significant concern, both from an environmental and an operational cost perspective. Infrastructure sharing or other flavours of cooperation among MNOs can serve as an efficient mechanism to decrease energy consumption by dynamically reallocating traffic across networks and selectively deactivating underutilised resources. Consequently, MNOs can lower operational costs and environmental impact without compromising service quality [2], [3].

Previous research studies highlighted the advantages of infrastructure sharing with respect to reducing energy consumption [4]–[6], enhancing failure resilience [7], and lowering both capital and operational expenditures [8]. Moreover, recent reports such as [9] point to inter-MNO radio access network (RAN) sharing as a promising approach for sustainable future mobile networks with lower environmental impact, e.g., lower energy consumption. In our earlier works [6] and [7], we analysed how collaboration between MNOs could improve network coverage and capacity, leveraging data from national authorities in the Netherlands on urbanization, population density, and existing cellular infrastructure across various regions. Both studies show that infrastructure sharing in the form of *national roaming*, defined as inter-MNO collaboration where all MNOs operate as a single national provider with centrally managed user association, yields significant improvements in terms of the *fraction of disconnected population* (FDP) and the *fraction of satisfied population* (FSP). In [6], in addition to improvements in FSP and FDP, we also provided an initial insight into potential energy savings enabled by NR. However, the user association process in [6] does not consider power consumption while determining user-base station (BS) association. To further explore the energy benefits of NR and the impact of user association in decreasing power consumption, this paper *investigates to what extent infrastructure sharing among MNOs can reduce energy consumption and proposes an energy-aware user association scheme*.

Specifically, we aim to address the following research questions:

- RQ1: How much performance improvement can NR offer compared to no cooperation in terms of energy consumption, FSP, and FDP? If any, how do the benefits vary across MNOs?
- RQ2: To what extent can a user association scheme incorporating power consumption cost of connectivity decrease network power consumption compared to an energy-oblivious scheme?
- RQ3: How does the density of BSs affect the performance of the studied scenarios?

To address these questions, we conduct a data-driven analysis

using real-world datasets to quantify the coverage, capacity, and energy consumption performance of NR in comparison to a no-cooperation baseline. This analysis is performed across three Dutch municipalities of varying sizes, allowing us to explore how the benefits of NR vary with BS density and across different MNOs. Additionally, we propose a simple yet effective energy-aware user association scheme and compare its performance against the approach presented in [7].

The remainder of this paper is organised as follows: Section II introduces the system model and key assumptions considered in the analysis while Section III presents the power consumption model adopted to assess the energy performance of the investigated scenarios. Section IV details the proposed energy-aware user association scheme. Section V provides a case study based on real-world data from Dutch MNOs, including data sources and relevant assumptions. Section VI presents the performance analysis addressing the listed research questions. Finally, Section VII discusses key limitations while Section VIII summarises the main insights of this paper.

II. System Model

We consider a downlink system consisting of O MNOs, where MNO_o has a set $\mathcal{B}_o$ of BSs available to serve its active subscribers, represented by the set $\mathcal{U}_o$. We consider the following two scenarios. The first scenario, referred to as *national roaming* (NR), assumes all MNOs cooperate to serve the aggregate of their subscribers, acting as a single national network. In this case, the set of all available BSs $\mathcal{B} = \bigcup_{o=1}^{O} \mathcal{B}_o$ is used to serve the set of all users $\mathcal{U} = \bigcup_{o=1}^{O} \mathcal{U}_o$. The second scenario, referred to as *no cooperation* (NC), represents the current operation of MNOs where each MNO serves only its own subscribers using its own infrastructure, i.e., MNO_o uses the set of BSs $\mathcal{B} = \mathcal{B}_o$ exclusively to serve the set of users $\mathcal{U} = \mathcal{U}_o$, without access to or interaction with the BSs or users of other operators. To simplify the notation, we will omit the MNO index when it is clear from the context. Accordingly, we generically represent the set of BSs as $\mathcal{B} = \{\text{BS}_1, \text{BS}_2, \ldots\}$ and the set of users as $\mathcal{U} = \{u_1, u_2, \ldots\}$, where these sets may include BSs and subscribers of either all MNOs or of a single MNO, depending on the considered scenario.

Each BS is assumed to have three sectors, each covered by a directional sector antenna. For BS_j, a bandwidth of W_j MHz is allocated according to the frequency planning applied by its MNO, and each sector is allowed to transmit with a maximum power level of P_{max} watts. The transmission power of BS_j in a given sector is denoted by P_j, where $P_j \in [0, P_{\text{max}}]$. We assume that each user u_i in the set of users $\mathcal{U}$ has a minimum rate requirement R_i^{min} bits per second to sustain a satisfying user experience for their application. If there is a link between user u_i and BS_j, we denote this link by ℓ_{ij}. To simplify the notation, the sector index is omitted throughout this paper.

We consider two types of cell sites following the models in 3GPP [10]: rural and urban macrocells. The path loss model follows the 3GPP TR 38.901 specification [10], with separate models for line-of-sight (LOS) and non-line-of-sight (NLOS) conditions. The path loss at a receiver u_i located at r_{ij} meters from BS_j is denoted by $\mathcal{L}(r_{ij})$. The antenna gain model for each sector also follows the 3GPP specifications [11]. Therefore, the horizontal radiation pattern A_H is given by:

$$A_H(\phi)[\text{dB}] = -\min\left\{12\left(\frac{\phi}{\phi_{\text{3dB}}}\right)^2, A_m\right\}, \tag{1}$$

where ϕ is the horizontal misalignment angle, $A_m = 20$ is the maximum attenuation, and $\phi_{\text{3dB}} = 65^\circ$ for three-sector antennas [12]. Thus, the antenna gain for user u_i connecting to BS_j is:

$$G_{ij}[\text{dB}] = G_{\text{max}} + A_H(\phi), \tag{2}$$

where G_{max} is the maximum directional gain of the antenna element. Users are assumed to have omnidirectional antennas with 0 dB gain.

Now, let us denote the total noise power over the transmission band by N_{tot} (in watts), which represents the sum of thermal noise power and receiver noise figure [13]. Then, the signal-to-interference-plus-noise ratio (SINR) γ_{ij} at user u_i for link ℓ_{ij} is:

$$\gamma_{ij} = \frac{P_j G_{ij} \mathcal{L}(r_{ij})^{-1}}{N_{\text{tot}} + I_{ij}}, \tag{3}$$

where I_{ij} is the interference experienced by u_i, defined as:

$$I_{ij} = \sum_{m \in \mathcal{B}^j} P_m G_{im} \mathcal{L}(r_{im})^{-1}, \tag{4}$$

where $\mathcal{B}^j$ is the set of BSs operating on the same frequency as BS_j within a radius r_{max}. We assume that the closest three BSs to BS_j implement interference coordination schemes [7], [14] and do not interfere with BS_j. For a signal to be decodable with the lowest modulation scheme, we assume that the required minimum SINR is γ^{min}.

In this paper, we do not explicitly model physical resource blocks (RBs), but instead abstract resource allocation using a time-sharing strategy. Accordingly, we assume that each BS_j applies time-sharing among its served users and define the fraction of time allocated to u_i as ξ_{ij}. Therefore, the achievable throughput for u_i connected to BS_j can be calculated using Shannon's theoretical capacity formula as:

$$R_{ij} = \xi_{ij} W_j \log_2(1 + \gamma_{ij}), \tag{5}$$

where the minimum time required to meet R_i^{min} is:

$$\xi_{ij}^{\text{min}} = \frac{R_i^{\text{min}}}{W_j \log_2(1 + \gamma_{ij})}. \tag{6}$$

As in [7], if the total fraction of time required to serve all users of BS_j exceeds 1, users are assigned time proportionally to their time minimum requirement ξ_{ij}^{min} as follows:

$$\xi_{ij} = \frac{\xi_{ij}^{\text{min}}}{\sum_{k \in \mathcal{U}^j} \xi_{kj}^{\text{min}}}, \tag{7}$$

where $\mathcal{U}^j$ is the set of users connected to BS_j.

We assume that each BS without any associated users switches to a power-saving mode where the BS consumes lower power compared to its state where the transmission circuitry of the BS is active to serve its associated users.

TABLE I: Power consumption parameters used in our evaluations [15], [16].

Symbol	Description and value
P_{airco}	Air conditioning power consumption: 225 W
P_{link}	Microwave backhaul link power consumption: 80 W
P_{proc}	Digital signal processor power consumption: 100 W
P_{trans}	Transceiver power consumption: 100 W
P_{rect}	Rectifier power per sector: 100 W
η	Power amplifier efficiency: 0.12
P_{max}	Maximum transmit power: 40 W
P_{sleep}	Sleep mode power consumption: 75 W

III. Mobile Network Power Consumption Model

To evaluate network-wide power consumption, we adopt a model inspired by [15] and introduce a sleep mode to represent the power-saving mode of BSs without any associated users. The adopted model accounts for both load-dependent and load-independent power consumption components per site and sector, which are introduced next.

Load-independent components: When a BS is active (i.e., at least one user is connected to one of its sectors), it consumes a fixed amount of power independent of load. This includes site-level components such as air conditioning with power consumption P_{airco} watts and backhaul consuming P_{link} watts, as well as sector-level components like rectifiers with power consumption P_{rect} watts per sector.

Load-dependent components: If a BS is active, each sector's power consumption includes the transmission power scaled by the power amplifier efficiency $\eta \in (0, 1]$, the signal-processing related power consumption P_{proc} (in watts) and the transceiver fixed power P_{trans} (in watts). These components are further scaled by a *load factor* (δ), which we assume as the ratio of actual to maximum transmit power per sector (P_j/P_{max}). This assumption for the load factor is made as a simplification with the aim of capturing the effect of reducing the load in a BS. As stated in [16], for macro BSs, mainly the transmitted power scales with the load.

Sleep mode: When a BS is idle (i.e., no users are connected to any of its sectors), it enters a low-power sleep mode with a constant power consumption P_{sleep}.

Total power calculation: The total network power is computed by summing the contributions of all BSs. For BS_j, the total power consumption can be calculated as follows:

$$P_{\text{BS}_j} = \begin{cases} \underbrace{P_{\text{airco}} + P_{\text{link}}}_{P_{\text{site}}} + n_{\text{sectors}} \cdot P_{\text{sector}}, & \text{if BS}_j \text{ is active} \\ P_{\text{sleep}}, & \text{if BS}_j \text{ is idle} \end{cases} \tag{8}$$

where P_{site} is the site-level fixed power components and the sector power is calculated as:

$$P_{\text{sector}} = \left(\frac{P_j}{\eta} + P_{\text{proc}} + P_{\text{trans}} \right) \delta + P_{\text{rect}}. \tag{9}$$

Table I summarises the parameters used in the power consumption model.

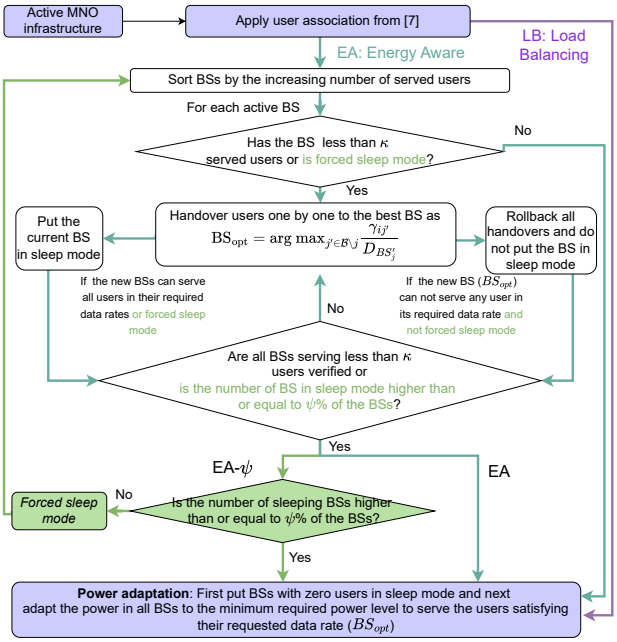

Fig. 1: Flowchart for EA, EA-ψ, and our baseline LB.

IV. Energy-Aware User Association Scheme

As explained in the previous section, there is a fixed and non-negligible power cost of keeping a BS active irrespective of the traffic served by this BS. A well-known approach to decrease power consumption is to consolidate load (hence users) in fewer BSs and switch the remaining BSs without any users to a lower-power sleep mode. To this end, association of users to the BSs considering associated power cost of connectivity becomes paramount. In this section, we present our proposal, referred to as *energy-aware user association* algorithm (EA), whose operation steps are illustrated in Fig. 1. We also present our comparison baseline Load Balancing (LB) approach adapted from [7] and a variant of our proposal EA-ψ where ψ is a design parameter to control the percentage of BSs that will be put in sleep mode, e.g., set by the MNO to cap the number of sleeping BSs or for our controlled experiments for a fair comparison among schemes.

Taking the active mobile infrastructure as its input, EA first determines the user association based on the approach from [7] and later hands over users and switches underutilised BSs to low-power sleep mode while preserving user quality of service. Although putting BSs into sleep mode may lead to a higher transmission power for serving previously-connected users due to the increase in the distance between the new serving BSs and the handed-over users, it reduces fixed power consumption at the sleeping BS. Our goal, therefore, is to design a scheme that identifies which BSs to put in sleep mode and that determines the connectivity configurations that enable net energy savings, while maintaining service quality.

As Fig.1 shows, EA begins by setting all BSs to transmit at full power and connecting each user to the best BS considering

the received signal power and the number of connected users at the corresponding BS. Here, BSs initially broadcast their cell-specific information with power P_{max} and users connect to the least-loaded BS with the highest SINR above γ^{min}. In this scheme, we iteratively connect a randomly-chosen user u_i to BS_{opt} where BS_{opt} is:

$$\text{BS}_{\text{opt}} = \arg\max_{j \in \mathcal{B}} \frac{\gamma_{ij}}{D_{BS_j}}, \tag{10}$$

where D_{BS_j} is the number of users connected to BS_j and $\mathcal{B}$ is the set of all available BSs. This load-aware, SINR-based association scheme provides better user distribution compared to simple $\max$-SINR approaches where each user connects to the BS which offers the highest signal quality. However, it is worth noting that it assumes either centralised control or coordination among BSs to maintain and share current load information (D_{BS_j}), which may require additional signalling overhead. While such centralised association is not typical in standard LTE deployments, similar strategies are employed in self-organizing networks [17].

After user association according to (10), our comparison baseline, referred to as Load Balancing approach (LB), continues with power adaptation whereas our proposal EA seeks further enhancements in user association to save from energy. More specifically, EA first sorts BSs by their number of associated users in ascending order. Consequently, the least-loaded active BSs, those serving fewer than κ users (e.g., 20 users), are considered for deactivation. For each candidate BS, the algorithm attempts to reassign its users to neighbouring active BSs. This is done user by user, selecting the best available BS based on (10). If the new serving BSs can satisfy the rate requirements of all reassigned users, the original BS is put into sleep mode. Otherwise, the handovers are not performed, and the BS remains active. This reassignment and validation process involves scanning multiple BS-user combinations and simulating handover outcomes. Consequently, the algorithm has a worst-case time complexity of $\mathcal{O}(|\mathcal{B}|^2 \cdot |\mathcal{U}|)$, where $|\mathcal{B}|$ is the number of BSs and $|\mathcal{U}|$ is the number of users.

This step is followed by power adaptation which is also performed in LB. First, all BSs without any served users are switched to sleep mode. Next, each active BS iteratively reduces power to the minimum necessary to meet its users' rate requirements. In every iteration, the interference is reduced, allowing for further reductions in transmit power. This iterative process continues until interference levels stabilise and no further significant power adjustments are possible[1].

In Fig.1, we also depict EA-ψ which ensures that exactly $\psi\%$ of the BSs are put in sleep mode. Its operation steps are largely the same as that of EA. Only when this scheme is in use, further improvements in user association stops after $\psi\%$ of the BSs are put into sleep mode. There could be also cases where all BSs with κ users are put in sleep mode but the percentage of sleeping BSs is smaller than $\psi\%$. In this case, referred to as *forced sleep mode* in Fig.1, the BSs are switched to sleep mode at the expense of satisfaction and connectivity. This later case happens rarely and under high ψ values. We consider such values only to investigate certain scenarios and for a fair comparison with our baseline LB.

[1]Empirically, we observe that four iterations offer a good trade-off between computational efficiency and power convergence across all tested scenarios, resulting in less than 0.1% difference in power consumption in subsequent iterations.

V. A Case Study on Dutch Cellular Networks

We perform a case study on the Dutch MNOs based on available data about cell towers, population distribution, and urbanity levels. In the following, we first provide an overview of the datasets and simulation parameters before presenting the performance of the considered schemes[2].

Antenna dataset: Our analysis leverages the antenna registration data provided by the Dutch Telecommunication Authority (Rijksinspectie Digitale Infrastructuur) [18]. For each recorded BS, the dataset includes information on the deployed technology (2G, 3G, 4G, 5G), geographic coordinates, center frequency, effective isotropic radiated power (EIRP) per channel, antenna height, and the number of sectors. As presented in Section II, to calculate the gain of these antennas, we use the antenna radiation pattern as given in the 3GPP TR 36.942 specification [11]. From the antenna registration dataset, we extract the main direction in which the antennas transmit and the EIRP. Therefore, we assume that the power already includes G_{max} in (2). Thus, to obtain G_{ij}, we only add (2) to the given EIRP. As in [7], we opted to remove omnidirectional antennas from the dataset and keep only three sector antennas, as the omnidirectional antennas present mainly indoor locations. The data that has been removed accounts for 5.3% of the total 3G, 4G, and 5G BSs. Moreover, we also removed BSs with 2G technology (corresponding to 29% of all BSs).

Mobile network operators: The Netherlands has three primary MNOs: KPN, Vodafone, and Odido (formerly T-Mobile). To distinguish the equipment owned by each MNO, since the antenna dataset does not specify the owner of each BS, we assign each BS to one of the three MNOs based on the frequency bands allocated to them [19] (see Table II). Throughout this paper, we refer to the three operators as MNO_1, MNO_2, and MNO_3, in no particular order. As shown by the operational frequency ranges in Table II, at the time of data collection (March 2023), 5G deployments were limited to the low-band (below 1 GHz) and mid-band (1–2.6 GHz) spectrum. Higher-frequency bands such as 3.5 GHz had not yet been deployed and are left for future investigation. Furthermore, during this period, the number of BSs varied by operator, ranging from about 8,000 for MNO_3 to roughly 13,000 for MNO_2.

Population density per 500×500m grid: To model population distribution, we use data from Statistics Netherlands (Centraal Bureau voor de Statistiek) [20], which provides the number of residents and the *urbanity* classification for each $500 \times 500\,\text{m}$ grid square across the country as of the

[2]The simulation framework used in this study is openly available at https://github.com/syllasrangel/wmnc-saving-energy-with-nr.

TABLE II: Frequency bands per provider for 3G, 4G, and 5G technologies [7].

	Centre frequency and bandwidth in MHz			
	3G	4G	5G	Total (MHz)
MNO_1	942.2 (5), 2152.6 (5)	816 (10), 1474.5 (15), 1815 (20), 2160 (20), 2605 (30), 2660 (10)	773 (10), 2160 (20)	175
MNO_2	957.4 (5), 2137.4 (10)	796 (10), 950 (10), 1487 (10), 1850 (10), 1860 (30), 1865 (20), 2137.5 (15), 2580 (20), 2652 (4), 2572.5 (15), 2672.5 (15), 2675 (20)	783 (10)	204
MNO_3	—	763 (10), 806 (10), 1459.5 (15), 1835 (20), 2117.5 (15), 2120 (20), 2630 (20), 2644.4 (10)	1835 (20)	140

end of 2020. Urbanity is categorised into five levels, with level-1 indicating the highest address density and level-5 the lowest. As in [7], we use the urbanity level of the region where a BS is located to determine whether it operates in an urban macrocell (UMa) or rural macrocell (RMa) scenario. Specifically, urbanity levels 1 to 3 are treated as UMa, while levels 4 and 5 are considered RMa. Additionally, to model active users, we assume that a fixed proportion f_p of the population is using the mobile network at any time, while the remaining $1 - f_p$ rely on alternative access technologies (e.g., Wi-Fi) or are offline. Active users in each $500 \times 500\,\text{m}$ area are distributed using a homogeneous Poisson point process, with intensity proportional to f_p times the local population density. As exact market shares of each MNO are unavailable, we assume a uniform user distribution across the three MNOs.

Performance metrics: The performance of a user can be characterized by two metrics: whether its signal quality is above the minimum signal level $\gamma^{\min}$ for establishing a connection with the network (if not, we refer to this user as *disconnected user*) and whether its satisfaction level (which reflects the user's perceived data rate) is sufficient for its application. If a user's data rate is above the minimum required rate ($R_i^{\min}$), we refer to this user as a *satisfied user*. The performance of the entire network, either per MNO or with all MNOs together, can then be characterised by the *fraction of disconnected population* (FDP), reflecting the number of users that do not have sufficient signal quality, and the *fraction of satisfied population* (FSP), reflecting the number of satisfied users. FDP and FSP reflect *coverage* and *capacity* of an MNO, respectively. Additionally, we report the total power consumption as in (8).

Simulation settings and scenarios: As previously mentioned, we assume that only a fraction of the population in each area is active at a time. As in [7], we assume the fraction of the active population as $f_p = 2\%$. Additionally, users' rate requirements are selected randomly between 8 and 20 Mbps where 8 Mbps is the minimum outdoor data rate to be provided by an MNO according to the regulations asserted by the Dutch regulatory body RDI [21]. We performed simulations to calculate the FDP, FSP, and power consumption for the considered settings for three municipalities of different sizes and populations, namely Amsterdam, Enschede, and Middelburg. To put this into perspective, the total number of BSs per city and per MNO is as follows: Amsterdam has 1683 BSs (586 for MNO_1, 376 for MNO_2, 721 for MNO_3), Enschede has 345 BSs (121, 73, 151), and Middelburg has 119 BSs (41, 36, 42).

For our performance analysis, we consider two scenarios: NR and no cooperation (NC), where NR denotes the scenario in which all MNOs operate as a single unified operator, while NC refers to the scenario where the MNOs operate their networks independently, representing the current reality. For each scenario, we evaluate two variants of the user association scheme: energy-aware (EA) and load balancing (LB), as introduced in Section IV. The results are reported as the average of 10 samples.

VI. Performance Analysis

This section presents our performance analysis based on our simulations to address the three research questions we introduced in Section I.

RQ1: *How much performance improvement can NR offer compared to NC in terms of energy consumption, FSP, and FDP? If any, how do the benefits vary across MNOs?*

To address this question, we compare the performance of individual MNOs and the system as a whole under NR and NC across three cities: Amsterdam, Enschede, and Middelburg. To ensure consistent conditions for comparison and limit the impact of user association, we assume EA-ψ as the adopted user association scheme for all scenarios and set $\psi = 20$ resulting in 20% of BSs being in sleep mode. Later, we relax this setting to EA to show the full potential of NR over NC. Fig. 2 presents the results.

As shown in Figure 2a, power consumption decreases across all MNOs under cooperation, with particularly noticeable reductions for MNO_3. Under NR, the power consumption decreases compared to NC by 16.77% in Amsterdam, 19.42% in Enschede and 18.59% in Middelburg while absolute power saving is, as expected, the highest in Amsterdam due to the denser BS deployment in Amsterdam compared to the other cities. In larger cities, operators have their largest customer base and prefer to overprovision heavily. In smaller cities, the density is lower, yielding less room for BS deactivation. We also observe the promise of NR in all regions in terms of FSP and FDP in Fig. 2b and Fig. 2c. Zooming in Fig. 2b,NR scenario leads to improved user satisfaction for all MNOs, but with different magnitude. This variation is influenced by the geographical distribution of users and infrastructure ownership in each city. In terms of performance metrics, MNO_3 can realise the greatest benefit from NR, while MNO_2 with the densest deployment the least yet still significant. In summary, all MNOs illustrate some level of improvement from 11.40%-26.47% in FSP and 87.33%-94.23% in FDP.

RQ2: *To what extent can a user association scheme incorporating power consumption cost of connectivity decrease network power consumption compared to an energy-oblivious scheme?*

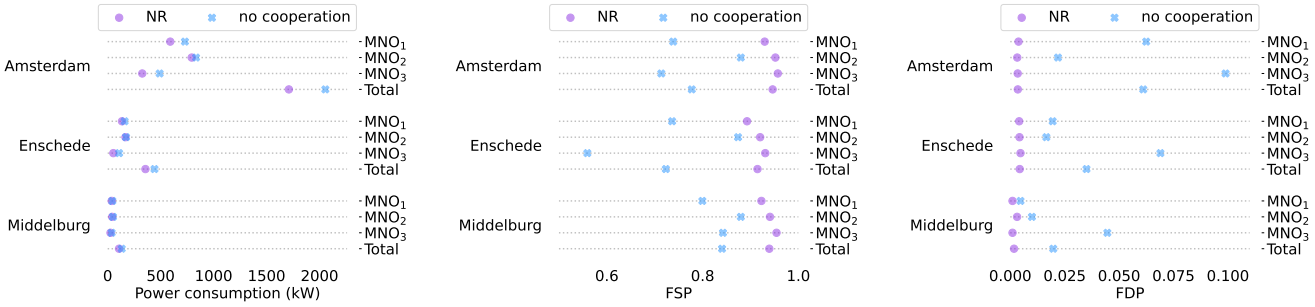

(a) Power consumption. (b) FSP. (c) FDP.

Fig. 2: Performance comparison of each MNO and the overall system in terms of power consumption under the NR and NC scenarios in Amsterdam, Enschede, and Middelburg. In all cases, 20% of BSs are turned off using EA-ψ.

Since our first analysis shows the promise of NR, now, we investigate the role of user association in power consumption of an MNO and the cellular infrastructure. To give also a perspective on how much current networks might benefit from our proposed user association scheme, i.e., EA, we also present results for NC, despite its inferior performance.

To study the effectiveness of EA, we conduct two analyses. First, we compare EA against LB from [7] and later to a baseline where the same number of BSs as EA are put in sleep mode but randomly. This later approach ensures a fair comparison in terms of the number of active BSs, isolating the impact of the intelligent BS deactivation mechanism used in EA. Fig. 3 shows the empirical cumulative distribution function (ECDF) of the resulting number of users per BS in Enschede, a mid-size city, after performing the considered user association. From this figure, we can gain insights into the utilisation of network infrastructure under NR.[3] A key observation from this figure is the significant proportion of BSs that serve no users, especially under EA scenario. These idle BSs represent direct opportunities for energy savings, as they are safely switched to sleep mode without impacting network performance. The variation in the number of users per BS arises from differing user and BS densities across various areas of the city, as well as the user association scheme in effect. Under EA, the ECDF begins at approximately 0.4, indicating that almost 40% of the BSs are switched to sleep mode. Furthermore, the curve rises gradually when the number of users is low, indicating that only a small number of BSs are serving a limited number of users. For example, around 20% of the BSs serve less than 10 users when users are assigned using EA compared to 50% under LB. This occurs because BSs with fewer associated users under EA are prioritised for handing over their users and entering sleep mode. In contrast, LB spreads users across more BSs, reducing the chances of having BSs move into sleep mode. In fact, as we can observe in Fig. 3, EA offers approximately 4$\times$ more opportunities to sleep BSs compared to LB. In this studied scenario, the energy consumption under EA is 27% lower than that of LB.

[3]Since we observe similar trends for both NR and NC scenarios, we report only the results for NR next.

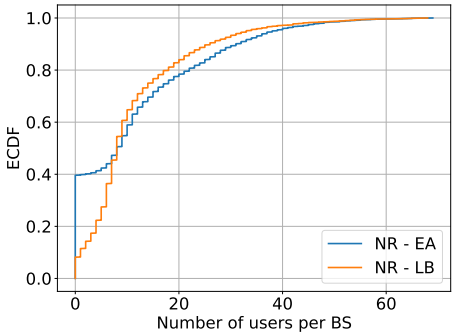

Fig. 3: ECDF of the number of users per BS in Enschede.

Now, we aim to understand how effective EA's BS sleeping approach is, and whether randomly selecting BSs for sleep mode would yield similar power savings and performance benefits. For this analysis, we consider EA-ψ with $\psi = \{0, 10, 20, 30, 40, 50\}$. To mimic our comparison baseline, we consider LB as the association scheme and only construct the active MNO infrastructure in Fig. 1 by randomly putting ψ% of BSs to sleep mode. We refer to this scheme as LB-ψ. Fig. 4 presents the results under increasing ψ values, i.e., the percentage of BSs that are put into sleep mode.

As shown in Fig. 4a, both EA-ψ and LB-ψ achieve reductions in power consumption as more BSs are set to sleep mode. However, EA-ψ consistently delivers this benefit with better user experience. Specifically, Figs. 4b and 4c demonstrate that EA maintains a significantly higher FSP and a lower FDP compared to the random LB-ψ. This is because EA-ψ selectively deactivates BSs that have minimal impact on coverage and user satisfaction, while random deactivation in LB-ψ can inadvertently disable critical BSs, degrading network performance. An interesting observation arises in both the NR scenario at $\psi = 10$ and in the NC scenario for $\psi \geq 10$, where LB-ψ shows slightly lower power consumption compared to EA-ψ. This occurs mainly because, in addition to the initially random selection of sleeping BSs, LB-ψ may include additional BSs that naturally remain unused due to

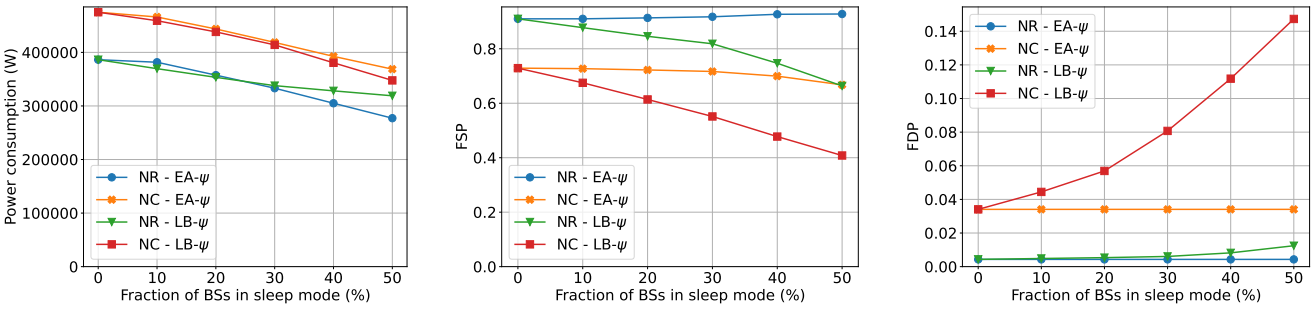

(a) Power consumption. (b) FSP (c) FDP

Fig. 4: Performance of EA-ψ and LB-ψ under both NR and NC. The same number of BSs is deliberately put into sleep mode to enable a fair comparison in terms of active infrastructure.

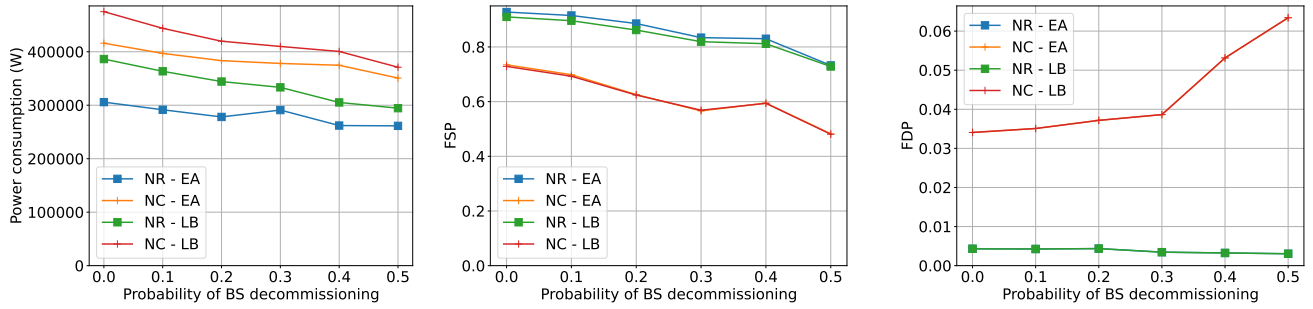

(a) Power consumption. (b) FSP (c) FDP

Fig. 5: Performance of EA and LB under both NR and NC for varying BS densities in Enschede. The x-axis indicates the probability of a BS being decommissioned (removed), reflecting a reduced deployment density.

lack of nearby users, effectively increasing the number of inactive BSs. However, this marginal energy saving comes at a noticeable cost: LB-ψ exhibits lower FSP and higher FDP than EA-ψ, underscoring the importance of controlled and context-aware BS sleeping.

RQ3: *How does the density of BSs affect the performance of the studied scenarios?*

Since our main argument for envisioned benefits of NR and EA is an over-provisioned BS infrastructure, in this section, we aim at understanding the impact of BS density on the observed gains for both NR vs. NC and EA vs. LB. To mimic this setting, we remove the same set of BSs from the active MNO infrastructure in Fig.1. Fig. 5 illustrates the comparative results for Enschede. A clear trend observed in Fig. 5a is that reducing the density of BSs leads to a decrease in overall power consumption. This trend can be logically explained by the power costs associated with each BS, as presented in (8). However, this power consumption reduction comes at a trade-off: as shown in Fig. 5b, the FSP tends to decline with lower BS density. This is expected, as fewer BSs limit the network's ability to serve users effectively in high-demand areas. Interestingly, a counter-intuitive trend is observed in Fig. 5c, where FDP decreases slightly as the BS density is reduced. We attribute this behaviour to the reduced interference in sparser deployments, particularly in overprovisioned networks. With fewer active BSs transmitting simultaneously, signal quality improves in certain areas, allowing previously disconnected users to establish a connection. Across all densities, the NR scenario consistently outperforms the NC configuration in every performance metric, reducing power consumption by up to 30%, decreasing FDP by up to 95%, and increasing FSP by up to 50%. This highlights the benefits of inter-operator cooperation, especially when combined with intelligent user association strategies. Finally, we observe that EA significantly outperforms LB in terms of power consumption and slightly in terms of FSP, while the FDP remains unchanged between EA and LB. This behaviour for FDP is by design, as EA starts from the same user association as LB and only deactivates BSs if doing so does not degrade FSP or not increase FDP. Our observations on the superior performance of EA hold for both NR and NC scenarios, which implies that current MNOs can decrease their power consumption significantly by following an energy-aware user association scheme that facilitates putting some more BSs in sleep mode without compromising their coverage (FDP) and user satisfaction (FSP).

VII. Limitations and Future Work

In this section, we list the key limitations of our work arising from several simplifications and assumptions we had to make for the sake of tractability of our large-scale analysis or due

to the lack of high-quality data on the considered cellular network setting. We believe that future works should address these limitations for a better understanding on the promise of infrastructure sharing and role of user association in reducing power consumption of a mobile network.

First, our user association schemes assume that all BSs initially transmit at full power. This assumption effectively models the worst-case interference scenario. While this conservative approach simplifies computation and ensures robustness, it may underestimate the network's actual capacity. We later adjust the power allocation to the minimum required for communication, reducing interference and potentially allowing more users to connect under the lower interference conditions. However, this potential improvement in user association after power adaptation is not explored in the current study.

Another important caveat is the necessity of making assumptions and simplifications regarding network operations. Specifically, we did not model interference management and coordination techniques in detail, despite their widespread use by operators in real-world deployments. In practice, a variety of mechanisms, ranging from interference cancellation to advanced power allocation, can significantly influence SINR, channel capacity, and ultimately user satisfaction. Hence, our results should be interpreted with these limitations in mind.

Finally, due to limited publicly available data on individual MNOs, we assumed an equal user distribution across all MNOs. While market reports suggest differing user shares, we refrained from using them because the role of mobile virtual network operators and their exact usage of physical infrastructure is not transparent. Additionally, we estimated user counts in each region based on population size. In practice, this may not fully capture reality, as the number of active cellular users can also depend on transient factors such as business or social activity.

VIII. Conclusion

This paper demonstrates the benefits of inter-operator cooperation via national roaming (NR), especially when combined with an energy-aware user association strategy. Using real-world data from three Dutch municipalities with diverse network topologies, we show that NR consistently reduces power consumption while improving user satisfaction and connectivity compared to traditional non-cooperative operation of mobile networks. These gains vary with local infrastructure and user distribution. Our energy-aware user association approach with polynomial complexity, despite being a greedy one, further enhances efficiency by selectively deactivating underutilised base stations (BSs) without degrading service quality, outperforming a load-balancing baseline adopted from the prior work. Overall, our results highlight the value of combining infrastructure sharing with intelligent user association to build more sustainable mobile networks.

Acknowledgements

This research was supported by the National Growth Fund through the Dutch 6G flagship project "Future Network Services".

References

[1] Rijksinspectie Digitale Infrastructuur Ministerie van Economische Zaken, "Dekkings-en snelheidsverplichting)," https://www.rdi.nl/onderwerpen/telecomaanbieders/dekkingseis-en-snelheidsverplichting, accessed on June 2025.

[2] L. Cano, A. Capone, and B. Sansò, "On the evolution of infrastructure sharing in mobile networks: a survey," *ITU Journal on Future and Evolving Technology*, vol. 1, no. 1, p. 21, 2020.

[3] Body of European Regulators for Electronic Communications, "Summary report on the outcomes of mobile infrastructure sharing workshop," https://www.berec.europa.eu, Tech. Rep., Dec. 2020, access on 03/2023.

[4] A. Antonopoulos, E. Kartsakli, A. Bousia, L. Alonso, and C. Verikoukis, "Energy-efficient infrastructure sharing in multi-operator mobile networks," *IEEE Comm. Magazine*, vol. 53, no. 5, pp. 242–249, 2015.

[5] M. A. Marsan and M. Meo, "Network sharing and its energy benefits: A study of european mobile network operators," in *IEEE Global Communications Conference (GLOBECOM)*, 2013, pp. 2561–2567.

[6] L. Weedage, S. R. Magalhaes, and S. Bayhan, "National roaming as a fallback or default?" in *IFIP Networking Conference, Workshop on OT/IT convergence*, 2023, pp. 1–6.

[7] L. Weedage, S. R. Magalhães, C. Stegehuis, and S. Bayhan, "On the resilience of cellular networks: how can national roaming help?" *IEEE Transactions on Network and Service Management*, 2024.

[8] L. Cano, A. Capone, G. Carello, M. Cesana, and M. Passacantando, "On optimal infrastructure sharing strategies in mobile radio networks," *IEEE Trans. on Wireless Comm.*, vol. 16, no. 5, pp. 3003–3016, 2017.

[9] NGMN Alliance, "Green future networks: A roadmap to energy efficient mobile networks," 2025.

[10] 3GPP, "Technical specification group RAN: Study on channel model for frequencies from 0.5 to 100 GHz (release 16)," Technical Specification (TS), 2019, 3GPP TR 38.901 V16.1.0 (2019-12, accessed on Feb.2023), https://www.3gpp.org/ftp/Specs/archive/38_series/38.901/.

[11] E. U. T. R. Access, "Radio frequency (RF) system scenarios. document 3GPP TR 36.942, V. 16.0. 0, 3rd Generation Partnership Project, Jul. 2020," 2020.

[12] M. Rebato, J. Park, P. Popovski, E. De Carvalho, and M. Zorzi, "Stochastic geometric coverage analysis in mmwave cellular networks with realistic channel and antenna radiation models," *IEEE Transactions on Communications*, vol. 67, no. 5, pp. 3736–3752, 2019.

[13] D. C. Halperin, "Simplifying the configuration of 802.11 wireless networks with effective snr," *arXiv preprint arXiv:1301.6644*, 2013.

[14] A. S. Hamza, S. S. Khalifa, H. S. Hamza, and K. Elsayed, "A survey on inter-cell interference coordination techniques in ofdma-based cellular networks," *IEEE Communications Surveys and Tutorials*, vol. 15, no. 4, pp. 1642–1670, 2013.

[15] M. Deruyck, W. Joseph, and L. Martens, "Power consumption model for macrocell and microcell base stations," *Transactions on Emerging Telecommunications Technologies*, vol. 25, no. 3, pp. 320–333, 2014.

[16] G. Auer, V. Giannini, C. Desset, I. Godor, P. Skillermark, M. Olsson, M. A. Imran, D. Sabella, M. J. Gonzalez, O. Blume *et al.*, "How much energy is needed to run a wireless network?" *IEEE wireless communications*, vol. 18, no. 5, pp. 40–49, 2011.

[17] B. Post, S. Borst, and H. van den Berg, "A self-organizing base station sleeping and user association strategy for dense cellular networks," *Wireless networks*, vol. 27, pp. 307–322, 2021.

[18] "Rijksinspectie digitale infrastructuur antenneregister," https://antenneregister.nl, accessed: 2024-09-18.

[19] "Frequentieoverzicht," https://antennekaart.nl/page/frequencies, 2024, accessed: 2024-09-18.

[20] "Kaart van 500 meter bij 500 meter met statistieken (Map of 500 meter by 500 meter with statistics)," https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/kaart-van-500-meter-bij-500-meter-met-statistieken, accessed: 2022-09-15.

[21] "Verlening vergunningen multibandveiling door Agentschap Telecom (Granting of multiband auction licenses by the Telecom Agency)," https://zoek.officielebekendmakingen.nl/stcrt-2020-41318.pdf, accessed on June 2025.