Multiscale Radio Reconfigurations: A Trace-Driven Approach to Estimating Network Performance

Aruna Prem Bianzino, Juan Manuel Montes-Lopez, Pablo Serrano
Universidad Carlos III de Madrid
Spain

Antonio A. de A. Rocha
Fluminense Federal University
Brasil

***Abstract*—As mobile networks become more complex to handle increasing data traffic and a broader range of services, operators must balance the trade-offs between static and dynamic configurations. While traditional static configurations across the entire network are simpler to manage, dynamic adjustments, though more complex to operate, are better suited to adapting to evolving demands. To explore this balance, in this paper, we use real data from a mobile network to evaluate the potential gains in throughput gains, measured by downlink traffic, when dynamically adjusting configurations at both spatial and temporal scales. Our findings show that combining these dynamic adjustments leads to significant performance improvements, with traffic volume gains exceeding 30% when configurations are tailored at the cell level and to the hour scale.**

***Index Terms*—Mobile network, Cell configuration.**

I. Introduction

Mobile networks provide ubiquitous connectivity, not only for users but also for services, devices, and sensors, making them increasingly critical to everyday life and business operations. The supporting infrastructure consists of thousands of access cells, each of which must be configured and maintained by network operators, leading to high operational costs and significant management overhead. The behavior of each cell affects neighboring cells and is influenced by variable network conditions, such as the number of connected users, the type and volume of their traffic, and their mobility patterns. As new services and technologies continue to be integrated [1], operators are challenged to simplify network management while maximizing infrastructure efficiency [2]. This is no easy task, as each additional service adds layers of complexity to the network. Looking ahead, this complexity will only increase, making future networks even more intricate, further complicating the management process [3].

To improve network performance, network slicing has emerged as a solution, enabling more agile networks that can dynamically adjust to varying traffic loads and service requirements. While studies have shown the benefits of such flexibility, they also reveal the added complexity in network management. This evolution toward flexibility has been supported by studies that quantify the benefits of such adaptability using real-world data traces, though it comes at the cost of increased complexity in network management [5]. In line with these developments, we explore the advantages of dynamic network re-configurations. Using real data, we aim to understand how these adjustments could further optimize network throughput, balancing gains in adaptability with the associated management overhead.

In this study, we focus on analyzing how more frequent network reconfigurations could improve performance, particularly in terms of the key performance indicator defined by the network operator: *Downlink Traffic*. We then perform the same analysis considering how the Downlink traffic is affected by allowing different network configurations at different geographical scales, and how the two configuration dimensions (i.e., time and space) affect each others. Using real-world data, we examine different configuration patterns at various scales to explore potential gains.

While this study primarily focuses on LTE technologies, the insights derived from our analysis may extend to Beyond 5G (B5G) networks. The configuration parameters examined, such as cell reselection priorities and signal thresholds, are foundational to mobile network operation and remain relevant in more advanced architectures. As such, the dynamic reconfiguration strategies explored here, along with the performance gains observed, could be replicated in B5G networks.

The rest of this paper is organized as follows: Sec. II analyze the state of the art for the optimization of cell configurations in mobile access networks. Sec. III describes the considered network scenario, as well as the reference dataset. It also details the network configuration parameters, and how they are currently selected by the operator. Sec. IV analyzes the network performance corresponding to the baseline configuration, i.e., the one preferred by the network operator. Sec. V evaluates the network performance under different configurations, as well as the potential gain resulting from their combination over space and time, at different granularities. Finally, Sec. VI presents conclusions and possible future research directions.

II. Related Work

In the defined context, various works describe efforts to optimize the configuration of mobile access networks, including different optimization objectives and methodologies. In particular, [6] analyzes real network traces to identify anomalies, to be used to design mitigation strategies through network reconfiguration of cell clusters. [7] analyzes the usage of neural network to predict the behavior of a network, on the basis of its configuration. This work considers the transmitted traffic as an environmental variable, while we consider it as the main KPI, as suggested by the network operator. [8] proposes the use of Digital Twins to test configurations in Radio Access Networks (RANs), and to train Reinforcement Learning solutions. The solution is then only tested on a small network (i.e., 5 nodes). [1] analyzes the usage of a Machine Learning (ML)-based framework, to find the optimal value of configuration parameters for cells in a mobile access network. The authors use Signal to Interference and Noise Ratio (SINR) as target KPI to optimize the parameters, and optimize only two configuration parameters. Similarly, [9] proposes an ML solution to the cell configuration problem, aiming at maximizing the cell coverage, while reducing the interference (hence not taking traffic nor users into account). At the same time, [10] suggests a solution based on support vector machine, to maximise the user throughput in RANs. Still, authors do not specify the reference scenario used. [11] analyzes a solution to optimize load balancing in 5G mobile cellular networks, through the configuration of cell parameters. The result is then tested on a synthetic network. Finally, [12] tackles the parameter optimization problem from an end-user point of view, using the Quality of Experience as an objective function. The results are then simulated for testing on a 19-cell reference network.

Overall, the solutions present in the reference literature allow to optimize the parameter selection for the cell configuration, taking into account different objectives and using different methodologies. Still, to the best of our knowledge, this work represents the first case of optimization directly involving a real network operator, which provides network data and measurements, as well as insights on the objectives and optimization goals. The findings are evaluated over a real and operating network, showing potential increases in the traffic carried by the network of above 30%, while insights are offered on which configuration changes offer the best increases in terms of network performance.

III. Network Description

A. Deployment sites

We consider the access network of one of the largest mobile network operators in Brazil. The network operates using LTE technology and comprises 1,553 cells, distributed across three municipalities: Itaguaí, São Gonçalo, and Petrópolis, covering a total of 85 neighborhoods with 158 nodes. The spatial distribution of the cells is shown in Fig. 1, with a detailed view of the cells in Itaguaí, shown as a zoom in.

The cells operate across four different frequency bands: 700 MHz (Low Band), 1.8 GHz and 2.1 GHz (Mid Band), and 2.6 GHz (High Band), offering a range of bandwidth configurations to serve varying network demands. There is notable heterogeneity in the network, with each cell using a different combination of frequency bands (42.2% of the cells utilize three frequencies, 27.4% use two frequencies, 17% operate with four frequencies, and 13.4% use just one frequency band).

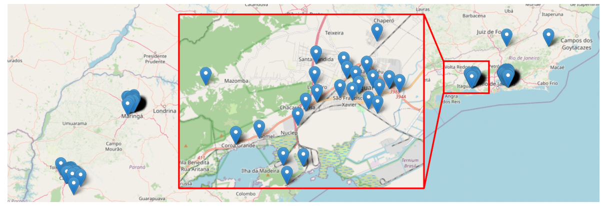

Fig. 1. Location of the cells in the considered network, and in the Itaguaí municipality in detail (zoom in).

Each cell measures different KPIs every hour, including traffic volume, number of users, number of connection drops, etc. Some of the measurements are missing, due to faults in the cell measurements or in the cell communication with the data server. On average, this problem affects about 20% of the measurements. Among the measured KPIs, the network operator indicated the *Downlink Traffic (DL) Volume* as the **most important metric**, with respect to the other ones, as it directly reflects network capacity utilization and the effective delivery of services to end users. Based on this preference, our analysis focuses on traffic volume as the primary metric to evaluate the quality of different network configurations.

B. Network configurations

For each cell, a configuration is selected from a set of predefined parameters, which can vary depending on the equipment, vendor, and version, creating some network heterogeneity. With numerous possible values for each parameter, optimizing configurations becomes complex and costly. To manage this, the operator typically relies on a trial-and-error approach, adjusting configurations based on performance improvements. The parameters are selected from a set of predefined configurations (*configurations 0, 1, 2, 3*), which have proven effective for the operator and help avoid performance issues. While a single configuration is usually applied across the network, adjustments may be made in specific problem areas. However, increasing the granularity of these configurations significantly raises management costs, which operators aim to minimize.

The parameters governing cell configuration play a crucial role in determining how user equipment (UE) transitions between different frequencies and cells. These settings influence decisions like when a UE should switch to a neighboring cell or change frequency based on signal strength, network load, and priority levels. By adjusting these parameters, the network can ensure that UEs are connected to the most appropriate cell, improving overall performance and balancing the distribution of traffic across the network.

The main parameters that influence cell configuration are essential in determining how UEs interact with the network. *Cellreselpriority* establishes the priority of a cell during reselection, meaning that when multiple cells have similar signal strength, UEs will prefer those with higher priority. *QRxLevMin* sets the minimum signal strength required for a UE to access a particular cell, ensuring that only cells with sufficient signal quality are used. The *SNonIntraSearch* parameter triggers the UE to search for cells on other frequencies when the current cell's signal degrades, while *ThrshServLow* focuses on initiating searches within the same frequency when the serving cell's signal becomes weak. Additionally, *q-Offset* applies an offset to the signal strength measurements of neighboring cells on different frequencies, which can influence the reselection process by either boosting or lowering the perceived strength of those cells. *ThreshXHigh* defines the threshold that neighboring cells must meet for reselection during inter-frequency transitions, and finally, *TResLTE* controls the amount of time a UE must measure a stable signal before reselecting, reducing unnecessary switches caused by temporary signal fluctuations.

The network operator has a set of recommended configuration parameters based on experience and performance data. These settings are designed to maintain network stability and optimize performance. However, due to privacy concerns and internal policies, the specific values of these configurations cannot be shared.

IV. Baseline Performance

In the following section, we look at different network statistics under the baseline configuration (`conf0`), which is the one used by the operator because it results in the best throughput in terms of traffic, i.e., the primary KPI. We do this to understand the baseline setup so we can compare it with other configurations and see how they affect network performance.

A. Downlink Traffic

We begin by analyzing the hourly traffic over the course of the week. Fig. 2 shows the traffic in terabytes (TB) for each day, from Monday to Sunday. It can be seen that the amount of traffic generated per hour reveals an average traffic of 3.4 TB. The highest traffic observed during peak hours reaches 5.7 TB, while the minimum traffic during off-peak periods drops to 0.6 KB.

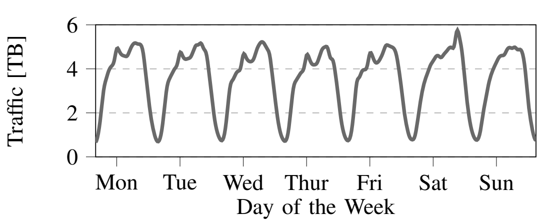

Fig. 2. Hourly traffic across the entire network over a week, using the baseline configuration.

The Traffic measurements in Fig. 2 displays a typical day-night cycle, with consistent patterns throughout both weekdays and weekends, and only slight variations. Interestingly, overall traffic remains stable during the weekends. Additionally, there is a distinct double peak each day, reflecting increased activity during midday and evening hours.

B. Peak hour & frequency distribution

The two peaks that arises in the traffic behaviour from Fig. 2 may be due to different geographical areas experiencing peak hours at different times of the day. Residential neighborhoods typically see peaks in the evening, while commercial areas peak around noon. Fig. 3 shows the distribution of peak hours for different network nodes, with the x-axis representing the time of day and the y-axis showing the percentage of times a given hour was a peak hour during the week. Most peak traffic occurs between 12:00 PM and midnight, with the highest concentration from 6:00 PM to 9:00 PM. There is also a smaller peak in the early afternoon, while very few or none in early morning and late night hours. This suggests that the majority of the network's peak usage happens in the afternoon and evening, aligning with typical user behavior patterns.

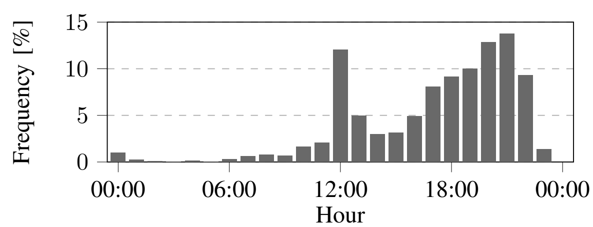

Fig. 3. Distribution of the peak hour for the different nodes in the network.

C. Users over time

We analyze the number of users in the network as it does not depend on the network configuration; rather, it is an external factor that influences the performance of the configurations. This makes it a critical variable to consider when assessing how different configurations impact overall network performance.

Fig. 4a shows the number of users throughout the week.Fig. 4a shows a similar trend to the weekly traffic pattern, with consistent day-night cycles and two peaks each day. However, while the number of users decreases on the weekends, the overall traffic remains steady. This suggests that, on weekends, individual users consume more traffic compared to weekdays, compensating for the lower number of connected users, which is expected.

More importantly, establishing a benchmark is crucial for comparing the different configurations. To this end, in Fig. 4b, we show Δ, the maximum difference in the number of users per hour across the four weeks analyzed. In the figure, the x-axis represents again the days of the week, and the y-axis shows the variation in the number of users, scaled by $\cdot 10^3$. We found that the differences are smaller on Tuesdays, and Wednesdays, indicating more stable conditions on these days, which makes comparisons between different configurations more reliable. In contrast,

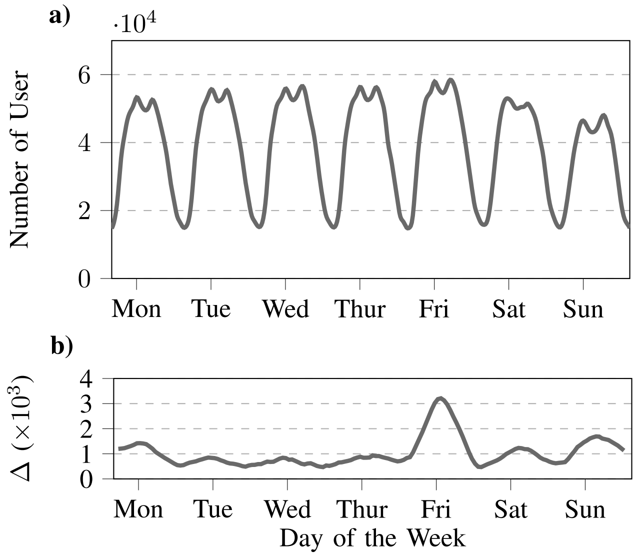

Fig. 4. (a) Average number of connected users per hour during the baseline week. (b) Hourly difference between the maximum and minimum number of connected users, to evaluate the variability across different weeks.

Fridays show greater variability between weeks, making comparisons less consistent.

D. Best static configuration

Finally, we assess the performance of the configurations by analyzing their throughput over two comparable days. As previously mentioned, `conf0` is particularly important, as the operator considers it to be the best configuration based on experience. Since Traffic is the primary KPI, we measure the traffic carried by the network under each configuration during those two days to establish a baseline for comparison and further analysis. This provides a reference point for evaluating the impact of dynamic configuration changes. On the basis of this analysis, `conf0` results in 170.3 TB, `conf1` in 167.1 TB (-1.9% with respect to `conf0`), `conf2` in 169.7 TB (-0.4% with respect to conf0), and `conf3` in 164.9 TB (-3.3% with respect to `conf0`). As we expected, `conf0` achieves the highest throughput. However, the traffic increase compared to other configurations is less than 4%, indicating only a slight improvement. Nonetheless, If we were to choose a single configuration for the whole network and the whole week, `conf0` would still be the optimal choice.

V. Configuration Analysis

To analyses the impact of choosing different configurations over the considered network, we agreed configuration changes with the operator to analyse their effect on the network conditions in order to design a methodology to evaluate their consequences, and be able to properly select the right configuration. In particular, the four indicated presets have been used, sequentially, each one for a consecutive week, over a period without major holidays nor network changes. The obtained configuration sequence has been `conf0` → `conf1` → `conf2` → `conf3`, between August and October 2023.

A. Temporal Analysis

To investigate the performance of each configuration, we analyze how traffic patterns fluctuate throughout the day, to identify which configuration performs best at each hour.

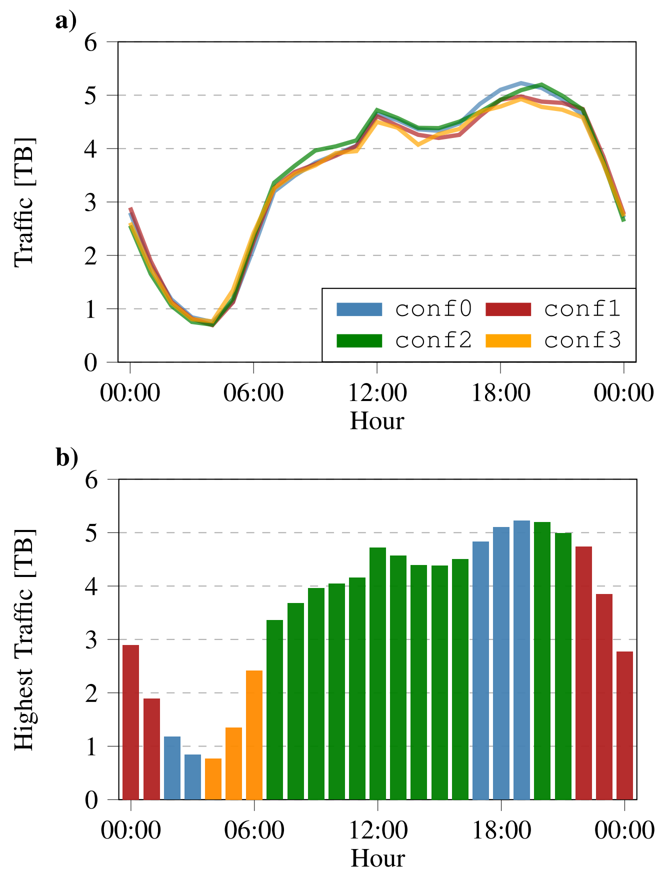

Fig. 5. (a) Total traffic carried by the network, under the different considered configurations, and (b) best performing configuration by hour.

In Fig. 5a four curves are shown, each one representing the amount of traffic observed across the network at each hour of the day, for a different configuration. Interestingly, the curves intersect, suggesting that, at certain times, different configurations result in more traffic than others. This means that, at finer temporal resolutions, `conf0` is not always the most effective choice. This is clearer in Fig. 5b, where the best configuration for each hour is shown, represented by different colors, along with the corresponding amount of traffic generated by that configuration at that time. Notably, `conf2` dominates most of the day, proving to be the optimal choice during the majority of the hours. We observe also that certain hours/conditions are better suited to some specific configurations than other, demonstrating that no single configuration is consistently optimal throughout the entire day. For instance, `conf3` looks the most suitable for off-peak hours, while `conf1` looks the most suitable for decreasing trends.

Since network traffic fluctuates throughout the day, we need to account how much each hourly-optimized configuration performs relative to a static configuration. This can be seen in Fig. 6, which shows the relative traffic gain when using the optimal configuration for each hour compared to `conf0` (i.e., the best configuration when considering only a single configuration over the whole day), and when using the second best configuration (also against `conf0`). In some hours, the gain exceeds 10%, while for most hours, if there is any gain, it remains below 5%. In general, when the best configuration shows no gain (e.g., around 18:00), it means that `conf0` is the best configuration. In these intervals, the second best configuration is performing worst than `conf0`, hence showing a negative gain with respect to that. Similarly, when the second best configuration shows a null gain (e.g., after 12:00), it means that `conf0` is the second best configuration. On average, switching to the optimal configuration each hour results in a 3% traffic increase, which is not significant enough to justify the added complexity of dynamic configuration changes, as this gain is similar to the difference between the best and worst static configurations for the network.

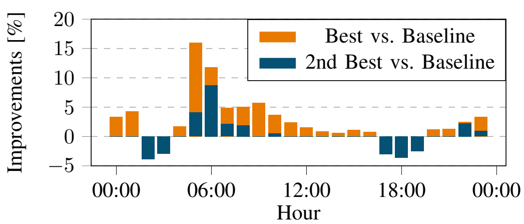

Fig. 6. Traffic gain by hour, when using the best (and second best) configuration for each hour, with respect to using always the `conf0`.

In this subsection, we demonstrated that dynamically adjusting the network configuration on an hourly basis across the entire network leads to some traffic gain. However, as shown in Sec. IVD, the gain is not significant with respect to a static configuration, indicating that the added complexity of dynamic adjustments may not always justify the performance improvements.

B. Spatial Analysis

We now analyze the spatial scale of reconfiguration. Intuitively, one might assume that similar network conditions in nearby areas would allow the same configuration to be effective across cells. To test this, we randomly selected a set of neighboring cells connected to different nodes and analyzed the optimal configuration for each cell, allowing reconfigurations every 6 hours.

We consistently found that the best configuration differed from cell to cell, as shown in Fig. 7, where each color represents a different optimal configuration for a cell over 6-hour periods throughout the reference day. The distribution of optimal configurations is largely uniform, indicating that tailoring configurations to specific geographical areas can significantly enhance traffic performance. This suggests that allowing different configurations for different geographical areas, allows the network to reach a higher gain in terms of carried traffic. In particular, when selecting the best configuration for each one of the 3 municipalities, a traffic increase of 1% is observed, while when selecting the best configuration for each neighborhood, the gain rises to 8.8%, 10.6% for node, reaching 17.5% when selecting the best configuration for each cell.

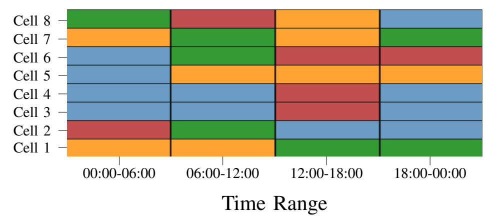

Fig. 7. Best configuration for different neighbour cells, on the reference day (blue: `conf0`, red: `conf1`, green: `conf2`, and yellow: `conf3`).

C. Estimating Spatio-temporal Improvement

Now, we evaluate the potential traffic gain from dynamically adjusting configurations at various spatial and temporal scales. To this extent, we admit, for each dimension of the analysis, different values of the corresponding granularity, and then analyze all the possible combinations, corresponding to different trade-offs. Specifically, we consider time granularities of 1, 2, 4, 8, 12, or 24 hours, alongside spatial granularities at the cell, node, neighborhood, municipality, or region level. For each combination of these granularities, we select the configuration that maximizes traffic over the study period. The resulting values of optimal configurations is presented in Fig. 8, where each curve corresponds to a spatial resolution, the x-axis indicates the time granularity, and the y-axis shows the average percentage of traffic gain for the chosen setting.

The simplest scenario is when a single configuration is applied across the entire network, and it is not allowed to change over the 24h period. This scenario corresponds to the left point of the bottom curve (orange), in Fig. 8. Here the gain is marginal (0.5%), with respect to having a fix configuration always. Remaining on the same curve but moving to the right, we see that selecting a single configuration for the entire region, but allowing hourly adjustments, (i.e., right end of the lower curve - orange - also corresponding to the scenario analysed in Fig. 5.b) yields an average gain of 3.3%. The curve immediately above (light brown) corresponds to the scenario in which different configurations can be selected for different municipalities. In the next scenario, different configurations can be chosen for different municipalities, which leads to a marginal improvement of 1% when changing configurations every 24 hours and 4.9% when changing every hour. In contrast, selecting configurations at a neighborhood or eNodeB level achieves more substantial improvements: when a single configuration is chosen for the entire day, gains reach 8.8% (neighborhood) and 10.4% (eNodeB), while allowing hourly configuration changes increases these gains to around 16.1% and 17.4%, respectively. Finally, allowing configuration changes at cell level results in substantial throughput improvements (dark blue curve). In a less complex scenario, where the configuration is fixed for each cell over 24 hours (left end of the curve), we observe a 17.5% traffic gain. In the most dynamic scenario, with hourly adjustments at the cell level (right end of the curve), traffic volume increases by 30.7%.

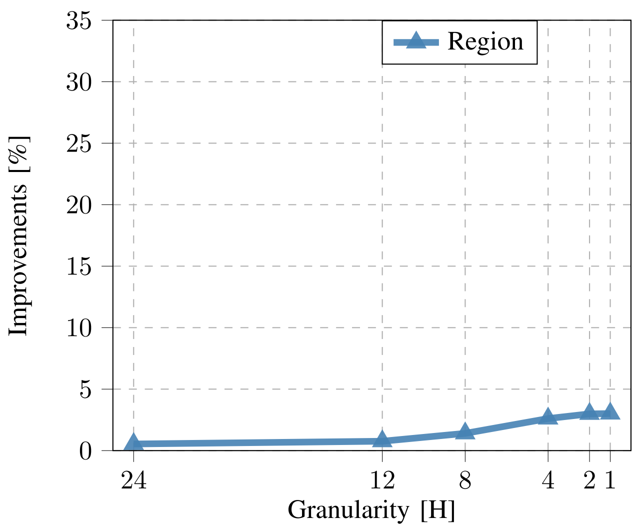

Fig. 8. Average percentage of traffic increase, allowing configuration changes with different time granularities (the x-value indicates the minimum interval over which a configuration cannot change), and at different geographical scales (different curves).

As expected, Increasing both temporal and spatial flexibility yields significant performance gains: allowing different configurations for neighborhoods or cells provides a larger improvement than simply changing configurations more frequently. As time granularity moves from 24-hour to 1-hour intervals, gains increase from about 3% at the region level to nearly 15% at the cell level. In the opposite direction, higher spatial granularity produces gains rising from roughly 17% under 24-hour intervals to about 30% under 1-hour intervals. These results show that each dimension of configuration flexibility not only boosts performance on its own but also enhances the improvements afforded by the other dimension, leading to more-than-linear overall gains.

VI. Conclusions and Future Works

The configuration of cells in mobile access networks is a crucial task for network operators and their business, as it represents the key task regulating network performance and operational and management costs.

In this paper, we first analyzed how the selection of different cell configurations affect the main KPI for the network operator (i.e., traffic volume carried by the network). We then analyzed the potential gain an operator may achieve by introducing flexibility in the selection of the cell configuration, both varying it in time and in spacial granularity, as well as mixing the two approaches. Results show that both flexibilities result in higher carried traffic, but that spatial flexibility in the configuration selection is much more effective, especially when allowing different configurations for different neighborhoods and for different cells (with gains of above 15% and of above 30% respectively). They also show that the two analysed flexibilities boost each other, achieving better network performance when combined.

In future works, we would like to better understand the correlations between network conditions and the corresponding configuration resulting better for them. We will evaluate how a configuration selection affects other KPIs, such as fairness for the traffic distribution among frequencies, or user drops. Finally, we will take into account the operational cost of the configuration changes and of their granularity, in order to evaluate an economic objective function, involving the network operator for its expertise in the field.

Acknowledgment

Thanks to the network operator for providing data and useful insights on the network functioning. This work is part of the TrialsNet project, funded by the European Union's Horizon-JU-SNS-2022 Research and Innovation Programme (Grant No. 101095871). This work is also partly funded by the SNS JU under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101139198, iTrust6G project. Antonio A. de A. Rocha received funding from IARA - Center of Artificial Intelligence in the Remarking of Urban Environments (FAPESP grant 2020/09835-1). We thank Diego Lopez from TID for his valuable feedback analyzing the results obtained.

References

[1] J. Shodamola, U. Masood, M. Manalastas, A. Imran. "A machine learning based framework for KPI maximization in emerging networks using mobility parameters." In 2020 IEEE International Black Sea Conference on Communications and Networking, IEEE, Tbilisi, Georgia, 2020.

[2] S. Almowuena and M. Hefeeda. "Dynamic configuration of single frequency networks in mobile streaming systems." In Proceedings of the 6th ACM Multimedia Systems Conference (MMSys '15), ACM, New York, NY, USA, 153–164, 2015.

[3] A. Yazar, S. Dogan-Tusha, H. Arslan. "6G vision: An ultra-flexible perspective." ITU Journal on Future and Evolving Technologies, 1(1), 121-140, 2020.

[4] GSMA Intelligence, "The Mobile Economy 2024", https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-economy/, last visited September 2024.

[5] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, X. Costa-Perez, "How Should I Slice My Network? A Multi-Service Empirical Evaluation of Resource Sharing Efficiency,", ACM MobiCom, India, 2018.

[6] J. Moysen, F. Ahmed, M. García-Lozano, J. Niemelä. "Big data-driven automated anomaly detection and performance forecasting in mobile networks." In 2020 IEEE Globecom Workshops (pp. 1-5), 2020

[7] M. Panek, I. Jablonski, M. Wozniak. "Modeling configuration-performance relation in a mobile network: a data-driven approach." arXiv preprint arXiv:2407.06702, 2024.

[8] I. Vilà, O. Sallent, J. Pérez-Romero. “On the design of a network digital twin for the radio access network in 5g and beyond.” Sensors, 23(3), 1197, 2023.

[9] Z. Lin, Y. Ouyang, L. Su, W. Lu, Z. Li. “A machine learning assisted method of coverage and capacity optimization (CCO) in 4G LTE self organizing networks (SON).” In Wireless Telecommunications Symposium (WTS) (pp. 1-9), IEEE, United States, 2019.

[10] A. Roshdy, A. Gaber, F. Hantera, M. ElSebai. “Mobility load balancing using machine learning with case study in live network.” In 2018 International Conference on Innovative Trends in Computer Engineering (ITCE) (pp. 145-150), IEEE, Aswan, Egypt, 2018.

[11] A. Asghar, H. Farooq, A. Imran. “Concurrent optimization of coverage, capacity, and load balance in HetNets through soft and hard cell association parameters.” IEEE Transactions on Vehicular Technology, 67(9), 8781-8795, 2018.

[12] R. D. Hegazy, O. A. Nasr. “A user behavior based handover optimization algorithm for LTE networks.” In 2015 IEEE Wireless Communications and Networking Conference (WCNC) (pp. 1255-1260), IEEE, Istanbul, Turkey, 2015.