

Application-Level Analysis of LTE Downlink User Throughput Under High Cell Utilization

Imran Kamal Mohammed Kadri

Radio Frequency Optimization and Data Analytics
Eastern Province, Saudi Arabia

Abstract— Downlink user throughput in congested LTE cells is usually diagnosed from cell load and radio indicators, yet cells with similarly high Physical Resource Block (PRB) utilization can deliver very different user throughput. This study examines the top congested L18 cells in Dammam city, selected for consistently high PRB utilization (7-day mean busy-hour DL PRB of 90.3–98.5%), using busy hour data for 6 October 2026 that combine radio KPIs with application-level records. Of 46 cells, 44 with at least 40 average connected users were analysed. Their busy hour DL user throughput ranged from 1.6 to 15.1 Mbps (9.6×) although PRB utilization remained between 90.3% and 99.1%. Average connected users, PRB utilization, MCS, and DL IBLER explained 56% of the variance in log throughput. Application goodput differed far more between applications than between cells: application identity explained 73% of the variance in log goodput, while cell identity explained 11%. Applications were classified as Good (fleet-average goodput of at least 2.0 Mbps, for example YouTube) or Bad (for example Facebook, TikTok, Instagram, and Snapchat). After controlling for users, PRB, MCS, and IBLER, a 10-percentage-point higher Good-application share of DL volume was associated with 9% higher throughput ($p = 0.012$), raising R^2 from 0.56 to 0.63. Five matched cell pairs with similar load and radio conditions illustrate the pattern: the higher-throughput cell carried 18 to 45 percentage points more Good-application traffic. The examples were selected from 14 matched pairs, of which 10 showed this direction, and the association is exploratory rather than causal. Load and radio conditions remain the primary explanation of throughput, while application mix provides an additional explanatory layer.

Keywords— LTE, downlink throughput, PRB utilization, application goodput, application mix, Good and Bad applications, Customer Experience Management, CEM, busy hour data, matched cell pairs, radio optimization.

I. INTRODUCTION

LTE downlink user throughput is one of the most visible indicators of radio network quality because it directly reflects the data rate experienced by active subscribers. In commercial optimization practice, throughput degradation is usually investigated through a combination of cell utilization, radio quality, scheduling, transport, and mobility indicators. DL PRB utilization is a natural first-order indicator of cell loading because 3GPP defines PRB usage as the percentage of available time-frequency resources used during the measurement period [1]. As utilization approaches saturation, fewer resources remain available for each active user and throughput generally decreases.

However, high utilization alone does not uniquely determine the throughput observed by users. Two cells can operate at similar PRB utilization while producing substantially different user throughput because the number of users sharing the resources, the radio conditions, scheduler behavior, device capabilities, server characteristics, and application traffic patterns are different. Application-level measurements have previously shown that different applications can experience different performance characteristics over cellular access, and video streaming and browsing can generate markedly different traffic patterns [2]–[5].

This study was motivated by an operational question: when a group of congested cells is persistently highly utilized, why can the user throughput still vary by several times? The investigation deliberately avoids assuming that one particular

cause must explain all cells. Instead, it follows a diagnostic sequence: first isolate consistently high-PRB cells; second quantify how much of the throughput variation is explained by cell load and radio indicators; third examine whether applications differ systematically in the goodput they achieve; and finally test whether the share of well-performing and poorly performing applications in a cell is associated with its throughput once load and radio conditions are controlled.

The main contribution is methodological rather than a claim that application behavior is the sole determinant of LTE throughput. The study demonstrates a practical diagnostic framework in which busy hour radio and load indicators establish the baseline and application-level records provide an additional explanatory layer for differences that remain after conventional KPI analysis.

II. RELATED WORK AND TECHNICAL BACKGROUND

A. PRB Utilization and Radio Throughput

In LTE, DL PRB utilization represents the fraction of available time-frequency resources used by the cell. High utilization is therefore an indicator of resource pressure, but it does not directly specify how efficiently those resources are converted into user data. CQI influences link adaptation, while MCS reflects the selected modulation and coding level. TA provides an indication of propagation distance and timing alignment. Error-related indicators such as IBLER and RBLER provide additional evidence of retransmission and link quality. These parameters are consequently useful for

distinguishing radio-limited throughput from throughput differences caused by other factors [1].

B. Application-Dependent Performance

Application traffic is not homogeneous. Controlled and measurement-based studies have shown that mobile applications exhibit different throughput requirements and traffic dynamics [2], [3]. Video streaming, for example, commonly uses adaptive and burst-oriented transfer strategies, while web browsing is comparatively bursty and transaction-oriented [4]. These differences mean that a cell-level throughput counter is an aggregate outcome of a heterogeneous set of application sessions rather than a direct measurement of a single application-independent demand.

C. Mobile Traffic Composition

Large-scale traffic studies also demonstrate that a relatively small set of applications can account for a large fraction of mobile data volume and that the traffic composition varies across device populations and network environments [5]–[7]. Consequently, application mix is a plausible dimension for investigating residual cell-level throughput variance after load and radio conditions have been controlled.

III. RESEARCH GAP AND OBJECTIVES

Most operational throughput investigations remain centered on network counters because these measurements are available for every cell and can be monitored continuously. Such analysis is appropriate for identifying load- and radio-limited cells, but it can become inconclusive when cells exhibit similar utilization and comparable radio indicators. The research gap addressed here is the lack of a practical diagnostic sequence that explicitly moves from cell-level counters to application-level behavior when conventional indicators do not explain the observed variance.

The study has five objectives: (1) quantify throughput variance among consistently high-PRB L18 cells; (2) measure how much of that variance is explained by average connected users, PRB utilization, and radio indicators; (3) determine whether applications differ systematically in goodput, independently of the cell in which they are observed; (4) classify applications into Good and Bad categories and test whether the Good-application share is associated with cell throughput after controlling for load and radio conditions; and (5) present matched cell pairs with similar load and radio conditions as concrete examples.

IV. DATASET AND METHODOLOGY

A. Study Population

The top congested L18 cells in Dammam city were considered. These were cells with consistently high PRB utilization: over the 7-day period 30 September to 6 October 2026, the mean busy-hour DL PRB utilization of the selected cells ranged from 90.3% to 98.5% (average 93.2%), and every cell exceeded 90% on at least 4 of the 7 days. The analysis uses busy hour data for 6 October 2026, a day on which all 46

cells had busy-hour PRB utilization above 90%. Only L18 cells were used so that bandwidth and propagation characteristics, including path loss and indoor penetration, are comparable across the population [9]; the selected L18 carriers use 20 MHz. Cell identifiers in this paper have been replaced by anonymized labels (C01 to C46).

B. Busy Hour Data

For each cell, busy hour data were extracted for 6 October 2026. The busy hour of a cell is the hour with the highest DL traffic volume. The radio and load indicators were average number of connected users (L.Traffic.User.Avg), DL PRB utilization, average CQI, average DL MCS, average TA, DL IBLER, DL RBLER, and DL user throughput (RmvLastTTI), taken from network counters. Application-level records were taken from the Customer Experience Management (CEM) platform, which provides for each application in each cell the number of subscribers, DL traffic volume, and average goodput. The top 20 applications per cell were retained and the remainder were aggregated into an Other category.

Two cells with fewer than 40 average connected users (23 and 36 users; throughput 39.6 and 24.8 Mbps) were excluded from the statistical analysis because, with very few users, high PRB utilization reflects a small number of heavy users rather than congestion. The remaining 44 cells form the analysis population; results for all 46 cells are also reported.

Table I. Study design and principal variables

Item	Definition
Population	Top congested L18 cells, Dammam
Selection basis	Consistently high PRB (7-day)
Analysis day	6 October 2026, busy hour data
Cells	46 selected; 44 analysed
Throughput KPI	DL User Throughput (RmvLastTTI)
Load and RF controls	Avg users, PRB, MCS, IBLER
Application metrics	DL volume share, goodput




Fig. 1. Diagnostic workflow used to move from congested-cell selection to application-level analysis.

C. Application Share and Goodput

Application share is the share of the cell's busy hour DL traffic volume carried by the application. Application goodput is the application-layer DL data volume divided by the data-transfer time, summed over all sessions of the application in the cell and expressed in Mbps. Goodput therefore describes the rate achieved while data were actually being transferred and is influenced by the radio conditions, but also by transfer size, application pacing, and the delivery path.

D. Good and Bad Applications

The fleet-average goodput of each application was calculated as its traffic-weighted mean goodput across all 46 cells. An application was classified as Good if this fleet-average goodput was at least 2.0 Mbps and as Bad otherwise. The threshold is a round value that separates the dominant social and short-video applications (1.1 to 1.5 Mbps) from video streaming and bulk-download traffic (2.2 Mbps and above). It was chosen after inspecting the application ranking, and its influence is examined in a sensitivity analysis. The Good-application share of a cell is the sum of the DL volume shares of its Good applications; the Bad share is the remainder.

E. Statistical Model and Matched Pairs

Cell throughput was modeled by ordinary least squares on the logarithm of throughput, with the logarithm of average connected users, DL PRB utilization, average DL MCS, and DL IBLER as controls. The Good-application share was then added, and the improvement was assessed with an incremental F-test. The variance of application goodput was partitioned between application identity and cell identity using a two-way model on log goodput for applications present in at least 35 cells.

Matched cell pairs were defined among the 44 analysed cells as pairs with PRB utilization within 2 percentage points, average connected users within 25%, CQI within 0.6, MCS within 1.3, IBLER within 1.3 percentage points, and a throughput ratio above 1.3. This produced 14 pairs. Five of them are presented as examples in Section V.E.

V. EXPERIMENTAL RESULTS

A. High Utilization Does Not Produce a Single Throughput Level

All 44 analysed cells operated at high PRB utilization on the study day, between 90.3% and 99.1% (mean 94.8%). Nevertheless, busy hour DL user throughput ranged from 1.6 Mbps to 15.1 Mbps, a 9.6-fold spread (Table II). PRB utilization alone therefore does not identify the throughput a user will experience. In simple correlation, throughput decreased with both PRB utilization ($r = -0.50$) and the logarithm of the number of connected users ($r = -0.23$, Fig. 2); within the narrow 90–99% PRB band, the most heavily utilized cells tended to have the lowest throughput.

Table II. Busy hour descriptive statistics of the 44 analysed cells

Variable	Mean	Min	Median	Max
DL user TP (Mbps)	6.9	1.6	6.5	15.1
DL PRB (%)	94.8	90.3	95.1	99.1
Avg users	128	45	116	353
CQI	9.03	7.00	9.14	10.59
DL MCS	12.02	8.37	11.99	15.45
TA (m)	959	284	634	4336
DL IBLER (%)	12.6	10.5	12.4	16.4
Good share (%)	40.9	14.1	39.3	75.4




Fig. 2. Busy hour DL user throughput versus average number of connected users for the selected high-PRB cells.

B. Load and Radio Conditions Explain Much, but Not All, of the Variation

A model with average connected users, PRB utilization, MCS, and IBLER as predictors explained 56% of the variance in log throughput (adjusted $R^2 = 0.51$; Table III, M1). Users, PRB, and IBLER were individually significant, whereas MCS was not once the other variables were included. Doubling the number of connected users was associated with about 25% lower throughput, one additional percentage point of PRB utilization with 5.9% lower throughput, and one additional percentage point of IBLER with 12% lower throughput. This confirms that load and radio conditions are the first layer of diagnosis, but 44% of the variance remains unexplained.

Table III. Regression of log throughput ($n = 44$); M1: load and RF controls, M2: M1 plus Good share

Predictor	M1 coef.	M1 p	M2 coef.	M2 p
$\ln(\text{avg users})$	-0.408	<0.001	-0.287	0.010
DL PRB (%)	-0.061	0.005	-0.056	0.006
DL MCS	0.023	0.465	0.051	0.111
DL IBLER (%)	-0.124	0.004	-0.093	0.024
Good share (pp)	—	—	0.0086	0.012
R^2 (adjusted)	0.559 (0.514)		0.627 (0.578)	

C. Application Goodput Differs by Application, Not by Cell

Application-level goodput was then examined across cells. For the 13 applications present in at least 35 cells (552 application-cell observations), application identity alone explained 73% of the variance in log goodput, whereas cell identity alone explained 11% (both together 83%). The same application tends to achieve a similar goodput wherever it is

observed, while different applications achieve very different goodput inside the same congested cells.

Table IV. Fleet-level goodput of the principal applications

Application	Cells	Fleet gp. (Mbps)	Median (Mbps)	Traffic (%)	Class
YouTube	46	3.32	2.30	13.5	Good
HTTPS	46	2.89	1.50	7.6	Good
Twitter	36	5.47	3.15	2.9	Good
GooglePlay	41	2.20	1.23	3.0	Good
Snapchat	46	1.48	1.35	4.1	Bad
Instagram	46	1.30	1.18	4.9	Bad
Facebook	46	1.21	1.05	16.5	Bad
TikTok	46	1.09	0.99	13.3	Bad
Other	46	0.25	0.24	3.7	Bad
Unknown	44	0.16	0.08	3.9	Bad




Fig. 3. Distribution across cells of application average goodput, with applications classified as Good (dark) or Bad (light) at the 2.0 Mbps threshold.

The separation is clearest between the two classes. YouTube achieved higher goodput than Facebook in 96% and higher goodput than TikTok in 98% of the cells in which both were observed. Within the Bad class the ordering is not stable: Facebook exceeded TikTok in only 58% of cells. The classification is therefore interpreted as a two-class distinction rather than a full ranking. Two applications deserve comment: HTTPS and GooglePlay are Good by their fleet average, but their median per-cell goodput is below 2.0 Mbps (Table IV), so they are borderline; their effect on the results is examined in Section V.D.

D. Good-Application Share and Cell Throughput

Across the 44 analysed cells, the Good-application share of DL volume averaged 41% and ranged from 14% to 75%. It was positively correlated with throughput ($r = 0.39$, Fig. 4). More importantly, adding it to the model with load and radio controls (Table III, M2) gave a coefficient of 0.0086 per percentage point ($p = 0.012$), meaning that a cell with a 10-percentage-point higher Good share had about 9% higher throughput at the same load and radio conditions. R^2 increased from 0.56 to 0.63 (incremental F-test $p = 0.012$). Including the two low-user cells gave a consistent result ($n = 46$; coefficient 0.0078, $p = 0.011$; R^2 0.69 to 0.74).




Fig. 4. Good-application share versus busy hour DL user throughput for the 44 analysed cells (Pearson $r = 0.39$).

Table V shows how the result depends on the definition of Good. The association remains significant for thresholds between 1.5 and 3.0 Mbps, when the analysis is restricted to applications carrying at least 0.5% of traffic, and when HTTPS is not counted as Good. It is not significant under the strictest definition, in which an application is Good only if its per-cell median goodput is at least 2.0 Mbps; this definition retains very few Good applications, and the dependence on the definition is reported as a limitation.

Table V. Sensitivity of the Good-share association ($n = 44$)

Good definition	r	Coef. (per pp)	p	Pairs agree
T = 1.5 Mbps	0.34	0.0075	0.022	9/14
T = 2.0 Mbps (used)	0.39	0.0086	0.012	10/14
T = 2.5 Mbps	0.37	0.0083	0.009	12/14
T = 3.0 Mbps	0.33	0.0078	0.019	12/14
T = 2.0, apps $\geq 0.5\%$ of traffic only	0.31	0.0080	0.030	9/14
T = 2.0, HTTPS not counted as Good	0.35	0.0071	0.037	10/14
Per-cell median goodput ≥ 2.0 Mbps	0.14	0.0051	0.247	9/14

E. Matched Cell Pairs With Similar Load and Radio Conditions

Among the 14 matched pairs, the higher-throughput cell had the larger Good-application share in 10. Five pairs are presented in Tables VI and VII and Fig. 5; they were selected from the 14 and are illustrations of the pattern, not a random sample. In each pair the two cells have PRB utilization within 0.7 percentage points, and CQI and IBLER within the matching limits. Pairs 2 and 3 share the cell C09, compared with two different partners.

Table VI. Load and radio conditions of the five matched pairs (first cell of each pair has higher throughput)

Pair / cell	TP	PRB	Users	CQI	MCS	IBLER	Good %
1 C37 (high)	10.0	95.2	92	9.2	12.8	11.6	73
1 C28 (low)	6.4	95.5	117	9.7	13.5	11.1	29
2 C09 (high)	12.8	94.7	152	9.1	12.2	11.3	66
2 C11 (low)	4.7	95.0	166	9.1	11.6	11.8	34
3 C09 (high)	12.8	94.7	152	9.1	12.2	11.3	66
3 C43 (low)	6.2	95.0	152	9.1	12.8	12.4	38
4 C29 (high)	10.8	95.8	70	8.7	9.9	14.2	44
4 C13 (low)	5.8	96.2	80	8.8	10.8	14.7	26
5 C07 (high)	10.4	94.3	83	9.7	12.0	12.7	67
5 C39 (low)	7.7	94.9	86	9.5	12.0	13.8	44




Fig. 5. Good-application share and throughput of the higher- and lower-throughput cells in the five matched pairs.

Table VII. Application mix in the five matched pairs (% of busy hour DL volume)

Pair / cell	Good	Bad	YouTube	Facebook	TikTok
1 C37 (high)	73	27	26.4	1.6	9.2
1 C28 (low)	29	71	9.1	36.4	18.6
2 C09 (high)	66	34	21.1	5.7	8.7
2 C11 (low)	34	66	11.8	15.3	21.8
3 C09 (high)	66	34	21.1	5.7	8.7
3 C43 (low)	38	62	8.2	6.3	8.5
4 C29 (high)	44	56	8.2	8.2	11.9
4 C13 (low)	26	74	7.8	45.2	11.3
5 C07 (high)	67	33	8.7	2.3	7.0
5 C39 (low)	44	56	4.7	2.9	24.3

In Pair 1, the cells C37 and C28 had PRB utilization of 95.2% and 95.5% and MCS of 12.8 and 13.5, yet throughput was 10.0 and 6.4 Mbps (1.56 $\times$). The faster cell carried 73% Good traffic, with YouTube at 26% and Facebook at 1.6%, whereas the slower cell carried 29% with Facebook at 36%. Pair 2 shows the largest throughput gap (12.8 versus 4.7 Mbps, 2.71 $\times$) between cells with CQI 9.1 and 9.1; the faster cell had 66% Good traffic against 34%, and a TikTok share of 8.7% against 21.8%. Pair 3 is the closest load match, with 152 and 152 connected users; throughput was 12.8 and 6.2 Mbps and the Good share 66% and 38%, with YouTube at 21% and 8%. In Pair 4 the slower cell was dominated by Facebook (45% against 8%), and its Good share was 26% against 44%. Pair 5 is the smallest throughput gap (1.35 $\times$); the slower cell had a TikTok share of 24% against 7%.

Two cautions apply. The higher-throughput cell has fewer connected users in several pairs (for example 92 against 117 in Pair 1), so the pairs are matched within 25% on users rather than identical, and part of the difference may reflect load. In addition, four of the matched pairs did not show the pattern, which is why the pairs are used to illustrate the mechanism and the regression in Section V.D is used to assess its strength.

VI. DISCUSSION

The results support a layered interpretation of LTE user throughput. At the first layer, PRB utilization identifies cells under resource pressure. At the second layer, the number of users sharing the cell and the radio conditions determine how much throughput each user can obtain; in this dataset, average connected users, PRB utilization, MCS, and IBLER together explained more than half of the variation. At the third layer, the composition of the traffic determines the goodput that the sessions in the cell can achieve.

The decomposition of goodput variance is the most distinctive finding. Application identity explained about 73% of the variance in goodput and cell identity only 11%, even though all cells were congested. Goodput differences between applications are therefore not an artifact of which cell they are measured in. The present data cannot separate the reasons. They may include application delivery behavior, such as adaptive streaming versus short bursty transfers, transfer sizes, and content-delivery or server-side characteristics. Prior work shows that streaming and browsing applications use very different transfer strategies [4], and that the delivery path, such as the point of presence a device is attached to, can change end-to-end performance in a commercial 4G network [8]; both are consistent with this explanation.

Given this stable application effect, it is plausible that cells carrying a larger share of high-goodput traffic show higher throughput at equal load and radio conditions, and the cell-level analysis is consistent with this: a 10-point higher Good share was associated with about 9% higher throughput. The effect is moderate. Adding the Good share raised R^2 by 6.8 percentage points, so load and radio conditions remain the dominant explanation. Application mix should be treated as an additional layer for cases in which load and radio indicators do not account for the observed throughput.

VII. PRACTICAL IMPLICATIONS FOR RADIO OPTIMIZATION

The framework can be applied as a troubleshooting sequence. When a high-PRB cell shows poor user throughput, the engineer should first compare average connected users, PRB, IBLER, and the other radio indicators with peer cells. If these clearly differ, load balancing, capacity, or RF optimization remains the primary action.

If a cell has load and radio conditions comparable to a faster peer but lower throughput, the application mix should be examined. A traffic profile dominated by Facebook, TikTok, or other Bad applications indicates that part of the low user

throughput reflects what the users are doing and how those applications deliver data, and not necessarily a radio deficiency. Such cells should not automatically be treated as RF problems.

The approach also suggests a dashboard design in which a throughput card is accompanied by a load and RF indicator and by an application-mix indicator showing the Good and Bad shares. This can reduce cases in which a cell is declared simply “congested” and can direct optimization effort to the correct cause.

VIII. LIMITATIONS

Several limitations should be acknowledged. First, the analysis uses busy hour data for a single city, and 44 cells, so the results are exploratory and should be validated on a larger population. Second, application share was measured as DL volume, which is partly shaped by throughput: adaptive applications can consume more volume in a faster cell. When the Good share was calculated from session counts or subscriber counts instead of volume, no association with throughput was observed. The reported association is therefore specific to the volume-based definition and may partly reflect throughput influencing the traffic mix rather than the reverse.

Third, the Good and Bad classification was defined after inspecting the data. The association weakens and becomes non-significant when Good requires a per-cell median goodput of at least 2.0 Mbps (Table V), and HTTPS and GooglePlay are borderline applications. Finally, the five matched pairs were selected from 14 matched pairs and four of the 14 did not follow the pattern; the pairs also differ slightly in connected users.

IX. FUTURE WORK

The next stage should repeat the analysis on multiple cells in other cities so that the stability of the application effect and of the Good and Bad classification can be assessed. Application shares based on sessions and users, which are less affected by throughput, should be examined in a larger sample.

A within-cell analysis is also promising: because different applications run in the same cell at the same time, comparing their goodput directly controls for radio conditions and load. Combining such an analysis with server and content-delivery information could separate application behavior from server-side effects. Finally, throughput models that include application-mix features together with load and radio indicators could be evaluated with permutation or SHAP importance.

X. CONCLUSION

This study investigated why LTE cells with consistently high PRB utilization can deliver very different user throughput. In busy hour data for 44 top congested L18 cells in Dammam, throughput ranged from 1.6 to 15.1 Mbps although PRB utilization stayed between 90.3% and 99.1%.

Average connected users, PRB utilization, MCS, and IBLER explained 56% of the variance in throughput, confirming that load and radio conditions are the first layer of diagnosis.

Application-level goodput differed much more between applications than between cells (73% versus 11% of the variance). Classifying applications as Good or Bad, a higher Good-application share of DL volume was associated with higher throughput after controlling for users, PRB, MCS, and IBLER (about 9% per 10 percentage points, $p = 0.012$), and five matched cell pairs with similar load and radio conditions illustrated the pattern.

The conclusion is not that application behavior replaces load and RF analysis, but that it adds a moderate explanatory layer that conventional counters do not represent. The association is exploratory, depends on the volume-based definition of application share and on the Good and Bad classification, and should be validated on more cells from different cities. For congested cells whose load and radio conditions do not explain their throughput, examining the Good and Bad application mix gives engineers an additional and practical diagnostic step.

REFERENCES

- [1] 3GPP TS 36.314, “Evolved Universal Terrestrial Radio Access (E-UTRA); Layer 2 Measurements,” Release 19.
- [2] N. Becker, A. Rizk, and M. Fidler, “A measurement study on the application-level performance of LTE,” in Proc. IFIP Networking Conf., Trondheim, Norway, Jun. 2014, pp. 1–9.
- [3] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and V. Bahl, “Anatomizing application performance differences on smartphones,” in Proc. ACM MobiSys, 2010.
- [4] A. Rao, Y.-S. Lim, C. Barakat, A. Legout, D. Towsley, and W. Dabbous, “Network characteristics of video streaming traffic,” in Proc. ACM CoNEXT, Tokyo, Japan, Dec. 2011.
- [5] M. Fiedler, L. Isaksson, and P. Lindberg, “Measurements and analysis of application-perceived throughput via mobile links,” in Network Performance Engineering, Springer, 2011, pp. 37–69.
- [6] Ericsson, Ericsson Mobility Report, Jun 2026.
- [7] ITU-R, “IMT traffic estimates for the years 2020 to 2030,” Report ITU-R M.2370-0, Jul. 2015.
- [8] F. Kaup, F. Michelinakis, N. Bui, J. Widmer, K. Wac, and D. Hausheer, “Assessing the implications of cellular network performance on mobile content access,” IEEE Trans. Netw. Service Manag., vol. 13, no. 2, pp. 168–180, Jun. 2016.
- [9] T. S. Rappaport, Wireless Communications: Principles and Practice, 2nd ed., Prentice Hall, 2002.