Different Policies for Different NodeBs: Comparing Downlink Schedulers in Cellular Base Stations

Zesen Zhang[1], Jon Larrea[2], Jarrett Huddleston[3], Haoran Wan[4], Ricky K. P. Mok[1,5], Bradley Huffaker[1,5], KC Claffy[1,5], Kyle Jamieson[4], Alexander Marder[3], and Aaron Schulman[1]

[1] University of California
[2] The University of Edinburgh
[3] Johns Hopkins University
[4] Princeton University
[5] CAIDA

Abstract. Cellular base stations rely on proprietary downlink scheduling algorithms that vendors independently develop to fairly and efficiently schedules traffic to competing users. Schedulers from different vendors can make different scheduling decisions depending on channel conditions, buffer status, fairness, and capability. This work is the first to show the significant scheduling policy differences in a head-to-head comparison of the behavior of downlink schedulers across four base station vendors (Ericsson, Samsung, Nokia and Huawei) running on four cellular providers (AT&T, Verizon, T-Mobile and Vodafone). The evaluation is based on 500Gbytes of downlink transfers across 20 base stations in five cities during semi-controlled network and signal conditions. In particular, we observe different strategies for allocating radio resources, for rate control, and for handling users with asymmetric channel quality. These results challenge the assumptions made about downlink scheduler uniformity in prior cellular performance measurement studies.

1 Introduction

The 3GPP standards [7] provide flexibility for base station vendors to provide their own proprietary downlink scheduling policies. For example, although the 3GPP suggests a certain data rate in a particular channel condition, it does not specify the relationship between these two variables [7]. The 3GPP also only recommends base station scheduler designers to balance between user demands and their equipment's capabilities, while simultaneously maintaining fairness. It does not define a strategy for doing so, as with other mobile protocols [1]. Therefore, each of the four primary base station vendors in the world (Ericsson, Nokia, Samsung, and Huawei) implements their own proprietary scheduling policies. These policies are the primary "secret sauce" that differentiates base station performance across vendors.

Since all LTE base stations are 3GPP standards compliant, and compounded by the lack of visibility into what vendor base station a device is using, there has been an assumption of homogeneous behavior across different vendors' schedulers. However, prior studies have yielded contradictory observations about LTE base station scheduling behavior. Several studies have found bursty patterns in resource allocation across time slots, [4, 5, 27], while others have found equal sharing of resources across UEs in every time slot [25, 26]. This apparent contradiction implies that cellular performance monitoring assessments may require tailoring to a variety of base station scheduling policies.

In this study, we provide the first preliminary evaluation of differences in resource scheduling policies across the top four base station vendors: Ericsson, Samsung, Huawei and Nokia. Ascertaining differences between downlink scheduling policies is challenging because: (1) Cellular modems have no built-in method to determine the vendor of a base station. (2) Observing the scheduling policy of a base station requires a controlled environment, namely an idle base station with two controlled users contending for resources. (3) Base stations may behave differently based on configuration parameters or deployment locations, although operators tend to deploy a single vendor's equipment within a region.

We identified a base station's vendor through a combination of directly viewing the logo on the base station itself, and using a proprietary dataset of vendor fingerprints collected by Revelare Networks [16] We achieved a semi-controlled setting by performing experiments overnight in non-residential areas; we validated that the base stations were idle by checking if one UE could be allocated all resources of the base station. To evaluate a variety of vendors, we collected data from 20 different base stations, with different configurations, in five cities, across two countries, and across deployments of four major mobile carriers (Verizon, AT&T, T-Mobile, Vodafone). We compared scheduler behavior from the four most popular base station vendors, with two competing users, in a range of conditions including differing buffer status, channel quality, and traffic sources (i.e., application and transport protocol). Our contributions are:

1. We found different, but consistent, *radio resource allocation policies* for competing users across base station vendors, providers, and channel conditions.

2. We found differing *rate adaptation* algorithms in use. We found that some base station vendors aggressively pick data rates, while others opt for a linear data rate mapping to channel quality. The aggressive strategy resulted in higher end-to-end throughput in high quality channel conditions.

3. We found that when competing users have different channel qualities, different vendors in different deployments prioritize users differently. Some allocate resources unequally, while others evenly shared radio resources.

2 Background and Related Work

We focus on evaluating LTE base station behavior, as 4G/LTE remains the primary mobile data network in use today. Many 5G services still depend on 4G infrastructure via NSA mode, where 5G radio access is anchored by a 4G core [8].

We confirmed this in our experiments: among the four providers we examined, only T-Mobile operated 5G in SA mode, and only at two base stations.

Because the scheduling algorithm is largely determined by the core network, we configured phones to stay in LTE/4G mode, focusing on 4G scheduling behavior to avoid distortions caused by UEs switching between 5G and 4G bands, which complicates observation of base station scheduling. We discuss 5G results separately to motivate and frame future directions (§5). For context, we describe two key processes not directly defined by the 3GPP— radio resource scheduling (2.1) and link adaptation (2.2)—leaving their implementation and associated network performance implications as vendor-specific characteristics.

2.1 Radio resource scheduling

The scheduler is the process through which the base station distributes available Physical Resource Blocks (PRBs) across the UEs. The base station scheduler manages radio resource allocation for both uplink and downlink directions. Given the dominance of downlink traffic, we focus on downlink scheduling.

During our experiments, we found that a single UE receives all base station resources when it is the only UE downloading traffic. Therefore, we focus on base station scheduling strategies in competing UE scenarios. The scheduler makes a scheduling decision, i.e., distributes available PRBs across contending UEs, every Transmission Time Interval (TTI), which is 1 ms. In general, scheduler decisions consider multiple sources of information from UEs (KPIs), as well as radio measurements collected by the base station, combined with historic allocation data. Relevant KPIs commonly used across schedulers include the number of UEs, UEs' CQI reports, buffer status reports, and QoS rank [2, 22].

The literature defines several theoretical scheduling algorithms, including Round Robin, Maximum CQI, and Proportional Fair [12]. Most commercial base stations implement a custom variant of Proportional Fair that leverages the aforementioned KPIs, among others. This approach balances throughput and fairness by prioritizing UEs that meet custom criteria (e.g., those that have not recently received resources or that have substantial traffic to send) [6, 5].

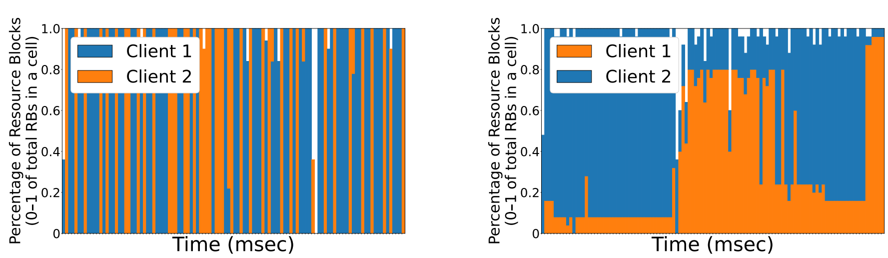

(a) All in TTI

(b) Split TTI

Fig. 1: Allocated PRBs over time for two different possible base station schedulers

Researchers have used observations of a specific commercial base station behavior to infer the assumptions that scheduling protocols make [4, 26, 5]. For example, some authors [4, 27] suggest that the base station scheduler assigns all PRBs to a single UE within one TTI, effectively generating a traffic burst (Figure 1a). Note in this figure how in each TTI the base station tends to assign resources to only one of the two UEs. Other studies [25, 26] have observed different resource allocation behaviors, such as in Figure 1b, where the base station usually shares resources of one TTI between two competing UEs. This difference in scheduling behavior can limit the benefits offered by proposed optimizations, e.g., BurstTracker [4] would trigger false positives in the case shown in Figure 1b.

2.2 Link Adaptation policy

Link Adaptation is the process by which the base station dynamically adjusts the Modulation and Coding Scheme (MCS) to encode channel quality information in each PRB, to maximize data rates and minimize loss. MCS selection and adaptation is vendor-specific; the 3GPP provides only recommendations. In practice, base stations generally rely on a combination of the Channel Quality Indicator (CQI) periodically sent by the UE and other radio KPIs measured at the base station. In general, higher CQI values enable the use of more advanced modulation schemes, such as 64-QAM, whereas lower CQI values trigger more robust coding and simpler modulation techniques, like QPSK.

As suggested in WiFi research, rate control algorithms can greatly affect network throughput and power consumption [18, 11, 9, 17]. How base station link adaptation is implemented, and how it impacts network performance, has been discussed in previous work [14, 15, 3]. However, those studies were either simulation-based or limited to a single vendor. None of them identified that CQI-related MCS selection strategies can vary across different vendors.

3 Methodology and Instrumentation

We describe our experimental environment and methodology, and demonstrate its efficacy for data acquisition and analysis of idle base stations.

3.1 Experimental Setup

Our goal is to delineate the differences between base station downlink schedulers implemented by different vendors. We performed all measurements with idle base stations, and generated downlink traffic by downloading from a controlled server to our UEs. Figure 2 illustrates our setup, with three distinct components.

Server Configuration. A server running Ubuntu 22.04 LTS with Linux kernel 5.15 LTS was deployed within the same geographic region as the base stations to ensure a round-trip time (RTT) under 40 milliseconds, thereby reducing the impact of end-to-end latency. The server was provisioned with sufficient egress bandwidth to avoid becoming a bottleneck. We generated network traffic using `iperf3 version 3.9` for TCP and UDP, `LSQUIC version 4.0.8` for QUIC, and `NGINX version 1.26.1` for HTTPS. We used the BBR congestion control algorithm for both TCP and QUIC in our tests.

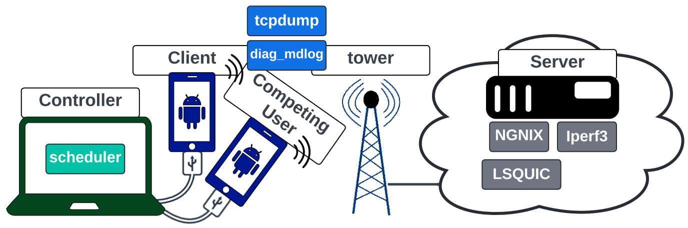

Fig. 2: Experimental setup for data collection.

Client Configuration. The client side consisted of two mobile phones (OnePlus Nord N30) running Android 14. The two phones competed for downlink resources by generating traffic using `iperf3`, `lsquic`, and `curl`. The connection between the phones and a laptop was maintained via USB, used solely for synchronization during code execution. To capture the base station's behavior at the Physical (PHY) layer, Qualcomm Modem Diagnostic Log (QMDL) files were recorded using `diag_mdlog` [13] and decoded with QXDM software [19, 23]. Additionally, `tcpdump` was employed to monitor and later analyze network traffic at the transport layer on the mobile devices.

Controller Functionality. The synchronization of the two phones was managed via a laptop by using thread barrier to guarantee a competitive scenario. We also release our code base on github [28].

3.2 Base Station Scheduler analysis

Downlink	**Ericsson**										**Samsung**						**Huawei**			**Nokia**
Config.	**Macro**						**Micro**				**Macro**				**Micro**		**Macro**			**Macro**
Provider	A	A	T	T	T	A	A	A	A	A	VZ	VZ	VZ	VZ	VZ	VZ	VO	VO	VO	A
Bandwidth	50	50	35	40	35	20	40	40	35	40	60	50	60	50	20	20	30	30	25	40
MIMO	2	2	4	4	4	4	2	2	2	2	2	4	4	4	4	4	2	2	2	4
Carrier Agg.	4	4	2	3	3	1	3	3	2	3	4	4	4	4	2	3	2	2	2	2

A=AT&T **T**=T-Mobile **VZ**=Verizon **VO**=Vodafone

Table 1: Details about the 20 diverse base stations we observed from 4 vendors

To understand the base station's scheduling behavior, we examine how it allocates bandwidth resources. This requires identifying an idle base station so that we can monitor all downlink traffic. An idle base station also allows us to reconstruct its behavior by aggregating the PHY-layer activity across all UEs.

To find idle base stations, we selected those located in non-residential areas during nocturnal hours to ensure a controlled environment. Most people remain at home during late night hours and typically use Wi-Fi even if they are awake accessing the Internet. The primary interference comes from passing pedestrians or drivers who briefly connect to the base station. To mitigate this risk, we performed an initial 5-second UDP download at the beginning of each experimental phase as metadata. When idle, all idle cells we tested allocated 98.8% of RBs to our single UE. We assume the remaining 1.2% of unallocated RBs were mostly reserved for control messages such as MIBs or SIBs, since they appear in subframes 0 and 5 [20]. We repeated each experiment multiple times (at least 3) to reduce the likelihood of interference from transient users.

During the experiments, we ensured that both phones were registered to the same Cell ID using NetMonster [21] before proceeding. We forced UEs to operate in LTE-only mode to observe LTE scheduling behavior, and enabled 5G mode when gathering 5G results. To study how vendors schedule heavy downlink traffic, we saturated the base station buffer for both User Equipments (UEs) by having each phone generate four parallel threads of 1 Gbps downlink UDP traffic from the server. By using UDP without congestion control and tolerating high packet loss rates, we ensured the base station buffer remained full for each UE throughout the experiment. Additionally, we used prepaid SIM cards to ensure consistent QoS and traffic conditions. In this way, we controlled as many factors as possible to observe vendor-specific scheduling differences.

Under high traffic load, carrier aggregation was triggered by the network. We identified cases during our experiments as carrier aggregation rather than dual connectivity by comparing physical cell IDs, RSSI, and frame number synchronization. The presence of consistent frame number and identical physical cell IDs indicated that multiple component carriers originated from the same base station (eNB) rather than from separate cells. Therefore, we treated these cases as carrier aggregation events and conducted our analysis across all aggregated carriers, under the assumption that they were managed by a single eNB with a common scheduler.

3.3 Data Overview

We collected downlink scheduler traces from 20 base stations across four of the top cellular vendors: Ericsson, Samsung, Nokia, and Huawei (Table 1), deployed by four major carriers: AT&T, T-Mobile, Verizon Wireless, and Vodafone. Our dataset includes 5 Ericsson Macro Cells, 4 Ericsson Micro Cells, 4 Samsung Macro Cells, 2 Samsung Micro Cells, 3 Huawei Macro Cells, and 1 Nokia Macro Cell. Operators typically deploy a vendor's equipment homogeneously within a given region (e.g., a city or state). To capture vendor diversity, we collected data from four U.S. cities—three on the west coast and one on the east coast. and from one city in Spain. Our data set represents most of the commonly deployed base stations worldwide, but is not exhaustive. In this preliminary study, we use this diverse collection of base stations to highlight differences in downlink schedulers, with the goal of emphasizing the need to account for such variation.

4 Results

We observed many differences in downlink scheduler behavior across vendors and base station types. Specifically, we identified significant differences along three dimensions of base station scheduling behavior: (1) the number of radio Resource Blocks (RBs) allocated to competing UEs per scheduling interval (TTI), (2) the link adaptation algorithm (i.e., how CQI is mapped to MCS), and (3) the behavior of UEs with diverse channel quality relative to the base station. We also provide a preliminary view of scheduling differences on 5G NR base stations.

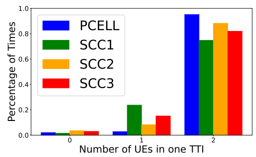

(a) Ericsson Macro BS

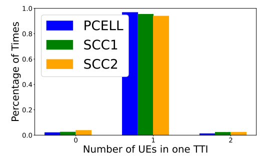

(b) Ericsson Micro BS

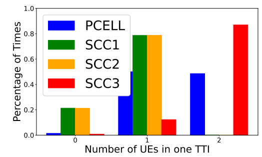

(c) Samsung Macro BS

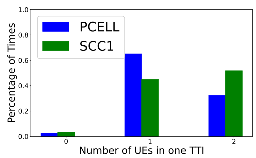

(d) Samsung Micro BS

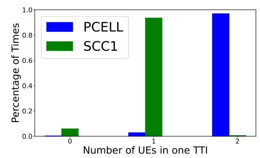

(e) Huawei Micro BS

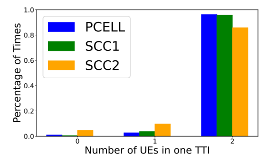

(f) Nokia Macro BS

Fig. 3: BS stands for Base Station. Base station schedulers across different vendors as well as macro and micro cell configurations.

4.1 Radio resource allocation policy

Figure 3 shows how often more than one of the two competing UEs are scheduled within a single TTI. Resources allocated to competing users varied across base station vendors, and even within each vendor's macro and micro cells. The differences between macro and micro base station reside in base station's power coverage. It can be easily differentiated by counting the number of sectors in one base station. Micro base stations are usually on the street lamp, while macro base stations can be found in base station towers. The behavior was consistent within each vendor and base station type, regardless of carrier. Therefore, We selected the most representative base stations to illustrate the behavior of those base stations in Figure 3 and we provide the remaining base station's result in Appendix.

Ericsson and Nokia *macro* base stations (Figure 3a), across all three carriers, consistently allocated resources to both UEs in every TTI. Three aggregated carriers transmitted traffic to the UEs simultaneously, including one primary and two secondary carriers. Within each carrier, resources were consistently distributed between the two competing phones. Resources were allocated in varying proportions each TTI, with ratios alternating every 5–6 TTIs. It is possible that *Nokia* base stations (Figure 3f) behave similarly to Ericsson's, as they are deployed by the same providers in different regions [10]. In contrast, Ericsson *micro* base stations (Figure 3b) generally allocated all RBs to a single UE per TTI, resulting in bursty resource distribution. However, across all aggregated carriers on the base station, both UEs received resources from at least one carrier 80% of the time. This suggests that Ericsson *micro* base stations primarily divide resources by assigning them to different users.

Samsung *macro* base stations (Figure 3c) displayed a different resource allocation pattern: 50% of the time, all RBs were assigned to one UE per TTI, while the other 50% of the time, both UEs shared resources within the same TTI on each carrier. Samsung *micro* base stations (Figure 3d) exhibited hybrid behavior, combining aspects of Ericsson's macro and micro cells. They frequently shared resources among UEs within the same TTI, like the macro cells, but occasionally allocated full RBs to one UE on each carrier, similar to the micro cells.

Huawei *micro* base stations (Figure 3e) behaved differently from Ericsson and Samsung, primarily in its use of carrier aggregation. The primary carrier consistently allocated resources to both UEs in each TTI, while the secondary carrier allocated all resources to only one UE per TTI. However, the secondary carrier alternated its allocations between the two UEs across TTIs, ensuring fairness.

Finally, we validated our UDP probing method by observing that similar scheduling policies emerged when using TCP downloads via iperf, as well as HTTPS and QUIC traffic (Section 5).

4.2 Link adaptation policy

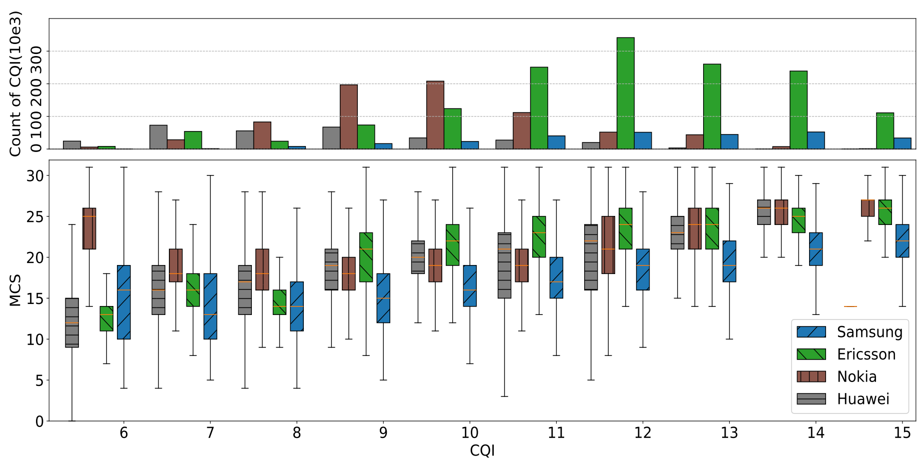

Fig. 4: Box-and-histogram plot showing MCS allocations by vendors across different CQI levels. The box plot indicates the range of MCS values observed for each CQI. The histogram shows the number of samples per CQI. *Note:* Results for CQI=6, particularly for Nokia, are based on limited data and should be interpreted with caution, as we rarely observed low-CQI scenarios (CQI 0–6). Such conditions likely triggered cell handoffs when better nearby cells were available.

Figure 4 shows the distribution of MCS rate control across various CQI levels (i.e., link adaptation). *Ericsson* is generally aggressive, assigning high MCS levels at high CQI. *Samsung* and *Huawei* exhibit a more cautious approach, with MCS allocation showing a linear decline in median MCS as CQI decreases, though *Huawei* generally assigns higher MCS values than *Samsung*. *Nokia* is as aggressive as *Ericsson* at high CQI but is more conservative at low CQI.

These observed differences reflect distinct vendor-specific strategies. Ericsson appears to prioritize speed and throughput under favorable signal conditions, potentially improving performance in high-CQI scenarios. By contrast, Samsung and Huawei's strategies appear more balanced, possibly to mitigate rapid throughput degradation under fluctuating signal strengths.

These results indicate that when a UE has high CQI, it will achieve higher end-to-end throughput from an Ericsson than from a Samsung base station. We validated this by testing the same UE in locations close to both Samsung and Ericsson base stations. We positioned the UE as close as possible to each base station and ensured that its CQI remained between 12–15 during the experiments (conducted at night when the base stations were idle). Using an iperf UDP flood, we measured bandwidth-normalized throughput of 8.75 bits/sec/Hz for Ericsson and 4.25 bits/sec/Hz for Samsung (both using 4x4 MIMO). This confirms that, under similar radio conditions, Ericsson base stations provided higher throughput. However, this performance gap likely narrowsin mid-range CQI levels, as Ericsson's aggressive strategy may lead to increased loss and retransmissions.

4.3 Policy for diverse channel quality

Next, we compared the resource allocation policies of macro base stations under varying network conditions. The focus was on observing how base stations from two vendors allocated resources to User Equipments (UEs) competing for the same radio resources when one UE had a much higher CQI than the other (e.g., when one was closer to the base station). We limited this experiment to *Ericsson* and *Samsung* to examine whether the divergent behaviors they demonstrated in earlier experiments also applied in this scenario (Figure 5). We compared

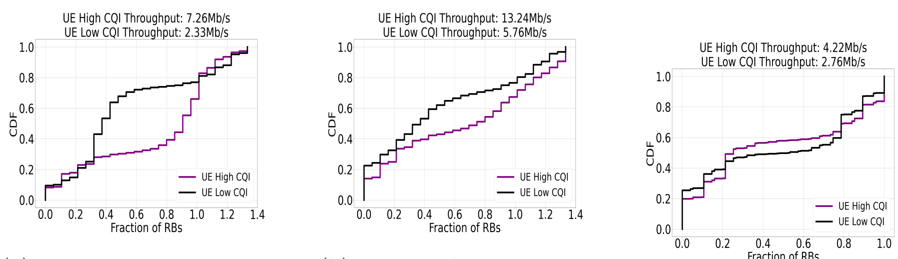

(a) Ericsson base station on campus.

(b) Ericsson base station in crowd area.

(c) Samsung base station.

Fig. 5: CDF of PRB usage for UEs with different CQIs contending for the same base station's downlink resources

the fraction of RBs allocated to each UE in each TTI, normalized by the total number of RBs available at the base station (i.e., its bandwidth). We discovered distinct vendor-specific allocation policies:

Ericsson: Ericsson exhibited resource distribution that varied according to signal quality differences between UEs:

- *Campus scenario (Figure 5a):* The UE closer to the base station had an average CQI of 11.7, while the more distant UE averaged 6.9. Surprisingly, the base station allocated more resources to the UE with poorer connectivity.

- *Crowd scenario (Figure 5b):* The closer UE averaged a CQI of 12.37, while the farther UE averaged 8.69. In this case, the base station allocated more resources to the UE with better channel quality.

Samsung: Samsung adopted a more balanced resource allocation strategy (Figure 5c); its macro base stations distributed RBs nearly equally between UEs, even when their channel qualities differed significantly. For the UE near the base station, the average CQI was 14.7, while the more distant UE averaged 7.19.

These divergent strategies highlight how Ericsson and Samsung base stations implement different approaches to resource management. In future work, we plan to extend this analysis to additional vendors.

5 Discussion

5.1 Scheduling policy diversity in 5G

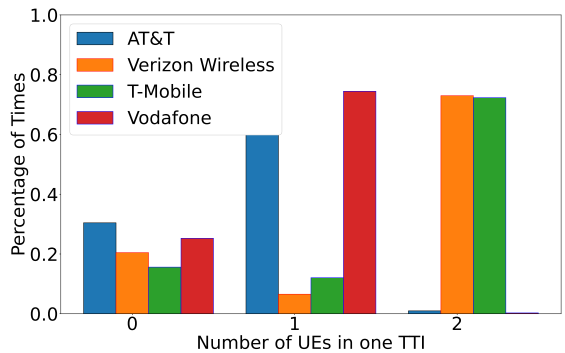

Fig. 6: UEs per TTI for 5G base station schedulers

This work focused on 4G LTE downlink scheduler differences, given its dominant deployment. However, we also compared radio resource scheduling on 5G base stations from Ericsson, Samsung, and Huawei. We conducted tests on AT&T, Verizon, and Vodafone who uses 5G NSA, and on T-Mobile, who use 5G Standalone. AT&T and T-Mobile use Ericsson base stations, Verizon uses Samsung, and Vodafone uses Huawei. Figure 6 displays the frequency of the number of UEs appearing per TTI. AT&T, Verizon, and Vodafone transmitted packets through both 5G NR and LTE radios simultaneously, as they use an NSA core network. AT&T and Vodafone scheduled one UE via the 5G radio and another via the LTE radio, which made it appear as though the 5G radio scheduled only one UE per TTI. However, synchronization between 5G and LTE transmissions is challenging because 5G uses Time Division Duplexing (TDD), whereas LTE uses Frequency Division Duplexing (FDD) in the regions where we collected data. Another issue is that 5G packet transmission is sensitive to CQI, since higher radio frequencies are more easily affected by noise. This makes observing behavior pattern between LTE and 5G harder. Conversely, T-Mobile, which operates on a 5G SA network, transmitted packets solely through 5G NR, and its base stations tend to share resources among UEs within the same TTI.

5.2 Policy effects on congestion control

Besides UDP tests, we conducted iperf TCP tests and file downloads via HTTPS and QUIC protocols with curl and lsquic to generate network traffic. We observed similar scheduling patterns as when downloading packets with UDP in each cell. Unfortunately, we are not able to provide further meaningful comparison between protocols, as the channel qualities between experiments are not under control.

We analyzed throughput fairness between UEs by evaluating 66 experiments conducted with two UEs located in the same position. We found that the throughput difference between the two UEs was less than 20% for over 85% of the time, indicating fair congestion management across all vendors despite differences in their scheduling algorithms. For the remaining 15% of cases, further investigation revealed two categories of anomalies. In five of these cases, the two UEs exhibited an average CQI difference greater than 2. We believe this may be related to the fact that even UEs placed in the same location can still experience slight channel differences. The next five cases were more interesting. In two of them—both involving Huawei base stations—we observed differences regarding how much each UE was using carrier aggregation. This raises questions about how base stations make carrier-aggregation decisions, which we plan to explore in future work. In the remaining three cases, one UE consistently received more than 25% additional resource blocks compared to the other. We also plan to investigate these cases further to determine whether they are caused by congestion-control algorithms or other underlying factors.

Our findings warrant reconsideration of previous studies that assume all base station scheduling behaves consistently. For example, BurstTracker [4] assumes that a base station schedules one UE in each TTI when managing traffic across competing UEs. Our experiments reveal this assumption is likely valid only for Ericsson Micro base stations. Applying the same assumption to Ericsson Macro base stations or Samsung or Nokia base stations would yield false positive detections of burst boundaries. In contrast, PBE-CC [26] assumes that resources are shared equally among multiple UEs within the same subframe or short runs of subframes. This model fits Ericsson Macro and Nokia base stations, where resources are shared among UEs for each carrier. However, Ericsson Micro base stations and occasionally Samsung base stations may allocate all RBs in a carrier to a single UE, which may limit the effectiveness of PBE-CC when implemented with these vendors. Deeper evaluation of these issues using tools such as NG-Scope [25] and NR-Scope [24] is a subject of ongoing work.

Overall, our findings highlight the importance of considering vendor-specific scheduling algorithms when evaluating or designing measurement and optimization tools for cellular networks. Ignoring these differences risks the integrity and generalizibility of research outcomes.

6 Conclusion

We analyzed downlink scheduling algorithms used by four base station vendors, Ericsson, Samsung, Huawei and Nokia, across five cities. Our experiments revealed vendor-specific strategies in resource allocation. Our findings also revealed significant differences in how these vendors allocate resources in LTE networks with contending users. Ericsson base stations tended to favor high MCS levels to users with better channel conditions, optimizing throughput in scenarios where signal quality was favorable. Samsung's approach was more conservative, maintaining a balance across varying signal conditions, which might help stabilize the user experience during fluctuating network quality. We hope these detailed insights into scheduling behavior and MCS allocation strategies will enhance future research on cellular network performance measurement and improvement.

7 Acknowledgements

We thank our shepherd and the anonymous reviewers for their insightful comments. This work was supported in part by National Science Foundation grants CNS-2223556, AST-2232457, OAC-2429485, ITE-2226460, ITE-2326928 and CNS-2213688. The views should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF.

References

1. 3GPP: 5G NR Medium Access Control (MAC) protocol specification. https://www.etsi.org/deliver/etsi_ts/138300_138399/138321/15.03.00_60/ts_138321v150300p.pdf
2. 3rd Generation Partnership Project (3GPP): Scheduling. https://www.3gpp.org/technologies/scheduling
3. Arshad, K.: Lte system level performance in the presence of cqi feedback uplink delay and mobility. In: 2015 International Conference on Communications, Signal Processing, and their Applications (ICCSPA'15). pp. 1–5 (2015). https://doi.org/10.1109/ICCSPA.2015.7081294
4. Balasingam, A., Bansal, M., Misra, R., Nagaraj, K., Tandra, R., Katti, S., Schulman, A.: Detecting if lte is the bottleneck with bursttracker. In: The 25th Annual International Conference on Mobile Computing and Networking. MobiCom '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3300061.3300140, https://doi.org/10.1145/3300061.3300140
5. Baranasuriya, N., Navda, V., Padmanabhan, V.N., Gilbert, S.: Qprobe: locating the bottleneck in cellular communication. In: Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies. CoNEXT '15, Association for Computing Machinery, New York, NY, USA (2015). https://doi.org/10.1145/2716281.2836118, https://doi.org/10.1145/2716281.2836118
6. Biernacki, A., Tutschku, K.: Comparative performance study of lte downlink schedulers. Wireless personal communications **74**, 585–599 (2014)
7. European Telecommunications Standards Institute: LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (3GPP TS 36.213 version 12.3.0 Release 12). Technical Specification TS 136 213 V12.3.0, ETSI, Sophia Antipolis, France (October 2014), available: https://www.etsi.org/
8. Ginsberg, M.: 4g shutdown timeline: When lte will end and how to future-proof with 5g and redcap. https://5gstore.com/blog/2025/06/25/4g-shutdown/ (2025), accessed: 2025-10-06
9. Gupta, V., Gutterman, C., Bejerano, Y., Zussman, G.: Experimental evaluation of large scale wifi multicast rate control. IEEE Transactions on Wireless Communications **17**(4), 2319–2332 (2018)
10. Hossain, B.: Case study: At&t's nokia-to-ericsson equipment swap in the u.s. https://www.linkedin.com/pulse/case-study-atts-nokia-to-ericsson-equipment-swap-us-bellal-hossain-hquve (2023), accessed: 2025-10-05
11. Huehn, T., Sengul, C.: Practical power and rate control for wifi. In: 2012 21st International Conference on Computer Communications and Networks (ICCCN). pp. 1–7. IEEE (2012)

12. Kumar, S., Sarkar, A., Sriram, S., Sur, A.: A three level lte downlink scheduling framework for rt vbr traffic. Computer Networks **91**, 654–674 (2015). https://doi.org/https://doi.org/10.1016/j.comnet.2015.08.027, https://www.sciencedirect.com/science/article/pii/S138912861500287X
13. Larrea, J., Shreedhar, T., Marina, M.K.: Biscay: Practical radio kpi driven congestion control for mobile networks. arXiv preprint arXiv:2509.02806 (2025)
14. Lu, F., Du, H., Jain, A., Voelker, G.M., Snoeren, A.C., Terzis, A.: Cqic: Revisiting cross-layer congestion control for cellular networks. In: Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications. p. 45–50. HotMobile '15, Association for Computing Machinery, New York, NY, USA (2015). https://doi.org/10.1145/2699343.2699345, https://doi.org/10.1145/2699343.2699345
15. Maattanen, H.L., Huovinen, T., Koivisto, T., Enescu, M., Tirkkonen, O., Valkama, M.: Performance evaluations for multiuser cqi enhancements for lte-advanced. In: 2011 IEEE 73rd Vehicular Technology Conference (VTC Spring). pp. 1–5 (2011). https://doi.org/10.1109/VETECS.2011.5956693
16. Marder, A., Larrea, J., Claffy, K., Kline, E., Jamieson, K., Huffaker, B., Thurlow, L., Luckie, M.: Reveal: Real-time evaluation and verification of external adversarial links. In: MILCOM 2024-2024 IEEE Military Communications Conference (MILCOM). pp. 1106–1111. IEEE (2024)
17. Murray, D., Koziniec, T., Dixon, M., Lee, K.: Measuring the reliability of 802.11 wifi networks. In: 2015 Internet Technologies and Applications (ITA). pp. 233–238. IEEE (2015)
18. Pal, S., Kundu, S.R., Basu, K., Das, S.K.: Ieee 802.11 rate control algorithms: Experimentation and performance evaluation in infrastructure mode. In: Passive and Active Measurement Conference. Citeseer (2006)
19. Qualcomm: eXtensible Diagnostic Monitor, https://tinyurl.com/yc4e9dcy
20. ShareTechnote: Sib scheduling. https://www.sharetechnote.com/html/BasicProcedure_LTE_SIB_Scheduling.html (2024)
21. Team, N.: Netmonster – advanced signal discovery. https://netmonster.app/ (2024)
22. Tech LTE World: Lte mac scheduler. https://techlteworld.com/lte-mac-scheduler
23. Vallina-Rodriguez, N., Auçinas, A., Almeida, M., Grunenberger, Y., Papagiannaki, K., Crowcroft, J.: Rilanalyzer: a comprehensive 3g monitor on your phone. In: Proceedings of the 2013 Conference on Internet Measurement Conference. p. 257–264. IMC '13, Association for Computing Machinery, New York, NY, USA (2013). https://doi.org/10.1145/2504730.2504764, https://doi.org/10.1145/2504730.2504764
24. Wan, H., Cao, X., Marder, A., Jamieson, K.: Nr-scope: A practical 5g standalone telemetry tool. In: Proceedings of the 20th International Conference on Emerging Networking EXperiments and Technologies. p. 73–80. CoNEXT '24, Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3680121.3697808, https://doi.org/10.1145/3680121.3697808
25. Xie, Y., Jamieson, K.: Ng-scope: Fine-grained telemetry for nextg cellular networks. Proc. ACM Meas. Anal. Comput. Syst. **6**(1) (feb 2022). https://doi.org/10.1145/3508032, https://doi.org/10.1145/3508032
26. Xie, Y., Yi, F., Jamieson, K.: Pbe-cc: Congestion control via endpoint-centric, physical-layer bandwidth measurements. In: Proceedings of the Annual Confer-

ence of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication. p. 451–464. SIGCOMM '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3387514.3405880, https://doi.org/10.1145/3387514.3405880
27. Xu, Y., Wang, Z., Leong, W.K., Leong, B.: An end-to-end measurement study of modern cellular data networks. In: Proceedings of the 15th International Conference on Passive and Active Measurement - Volume 8362. p. 34–45. PAM 2014, Springer-Verlag, Berlin, Heidelberg (2014). https://doi.org/10.1007/978-3-319-04918-2_4, https://doi.org/10.1007/978-3-319-04918-2_4
28. Zhang, Z.: Basestation scheduling test, https://github.com/ZSenZhang/Basestation_scheduling_test, gitHub repository

A Ethical Considerations

This study was conducted using researcher-owned devices on publicly deployed commercial cellular networks under non-invasive, observational conditions. All experiments were performed with our own SIM cards and equipment during low-traffic hours to avoid affecting normal users. No privileged network access, user data collection, or packet injection was involved, and only anonymized radio-level metadata and throughput statistics were analyzed. The work complies with institutional ethical standards and the ACM Code of Ethics, ensuring that all measurements respected user privacy and did not disrupt network operations or third parties.

B All the base stations' radio allocation behavior

In the following paragraph, we present all the other base stations' behaviors we have examined. Overall, for two phones downloading UDP packets at the same location at the same time, we conducted experiments on 16 base stations. Besides the 6 listed in Fig. 3, we list the remaining 10 in this appendix for reference.

B.1 Ericsson Macro Base Stations

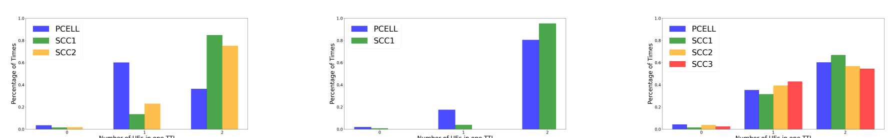

(a) Ericsson Macro BS 1 (b) Ericsson Macro BS 2 (c) Ericsson Macro BS 3

Fig. 7: Radio allocation for Ericsson macro base stations

B.2 Ericsson Micro Base Stations

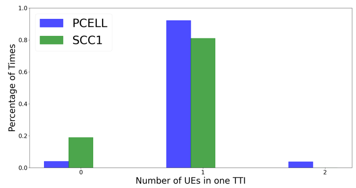

(a) Ericsson Micro BS 1

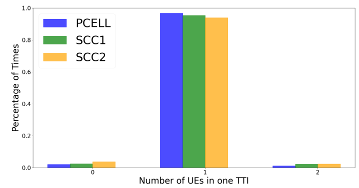

(b) Ericsson Micro BS 2

Fig. 8: Radio allocation for Ericsson micro base stations

B.3 Samsung Base Stations

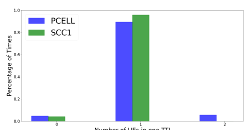

(a) Samsung Macro BS 1

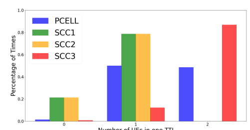

(b) Samsung Macro BS 2

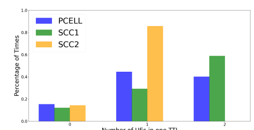

(c) Samsung Micro BS 1

Fig. 9: Radio allocation for Samsung base stations

B.4 Huawei Micro Base Stations

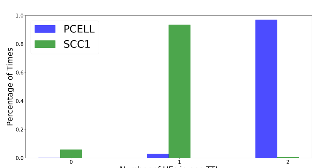

(a) Huawei Macro BS 1

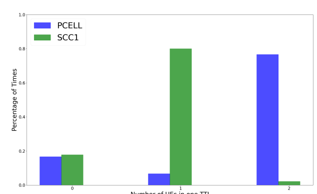

(b) Huawei Macro BS 2

Fig. 10: Radio allocation for Huawei micro base stations